

THE UNIVERSITY OF CHICAGO

GRAPH MATÉRN FIELDS WITH APPLICATIONS IN INVERSE PROBLEMS AND  
MACHINE LEARNING

A DISSERTATION SUBMITTED TO  
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES  
IN CANDIDACY FOR THE DEGREE OF  
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL AND APPLIED MATHEMATICS

BY  
RUIYI YANG

CHICAGO, ILLINOIS

JUNE 2022

Copyright © 2022 by Ruiyi Yang

All Rights Reserved

# TABLE OF CONTENTS

LIST OF FIGURES . . . . .	vii
LIST OF TABLES . . . . .	ix
ACKNOWLEDGMENTS . . . . .	x
ABSTRACT . . . . .	xi
<b>1 INTRODUCTION . . . . .</b>	<b>1</b>
1.1 Matérn Gaussian Fields—Beyond the Euclidean Setting . . . . .	1
1.2 Graph Matérn Fields . . . . .	3
1.2.1 The Series Representation . . . . .	3
1.2.2 The Graph-Laplacian . . . . .	6
1.2.3 Computation . . . . .	8
1.3 Applications . . . . .	9
1.3.1 The Bayesian Perspective . . . . .	9
1.3.2 An Elliptic Inverse Problem . . . . .	11
1.3.3 Semi-supervised Learning . . . . .	13
1.4 Continuum Limits . . . . .	14
1.4.1 Prior Level . . . . .	15
1.4.2 Posterior Level . . . . .	17
1.5 Outline and Main Contributions . . . . .	20
1.6 Notations . . . . .	21
<b>2 THE SPDE APPROACH TO MATÉRN GAUSSIAN FIELDS: GRAPH REPRESENTATIONS . . . . .</b>	<b>22</b>
2.1 Introduction . . . . .	22
2.1.1 Literature Review . . . . .	25
2.1.2 Main Contributions and Outline . . . . .	27
2.2 Matérn Models and the SPDE Approach . . . . .	28
2.2.1 Stationary Matérn Models . . . . .	28
2.2.2 Nonstationary Matérn Models . . . . .	30
2.3 GMRF Approximation with Graph Representations of SPDEs . . . . .	32
2.3.1 Graph Matérn Models . . . . .	34
2.3.2 Nonstationary Graph Matérn Models . . . . .	38
2.3.3 A Simulation Study . . . . .	40
2.4 Convergence of Graph Representations of Matérn Models . . . . .	44
2.4.1 Setup and Main Result . . . . .	45
2.4.2 Outline of Proof . . . . .	49
2.5 Numerical Examples . . . . .	52
2.5.1 A General Framework: Latent Gaussian Models . . . . .	53
2.5.2 Application in Bayesian Inverse Problems . . . . .	56

2.5.3	Application in Spatial Statistics . . . . .	60
2.5.4	Application in Machine Learning . . . . .	65
2.6	Conclusions and Open Directions . . . . .	68
2.7	Appendix . . . . .	70
2.7.1	Preliminaries . . . . .	70
2.7.2	Convergence of Spectrum with Rate . . . . .	73
2.7.3	Convergence of Eigenfunctions with Rate . . . . .	83
2.7.4	Convergence of Graph Matérn Field . . . . .	91
3	KERNEL METHODS FOR BAYESIAN ELLIPTIC INVERSE PROBLEMS ON MANIFOLDS . . . . .	98
3.1	Introduction . . . . .	98
3.2	Bayesian Inverse Problems on Manifolds and Point Clouds . . . . .	103
3.2.1	Bayesian Elliptic Inverse Problems on Manifolds . . . . .	103
3.2.2	Kernel Approximation of the Forward and Inverse Problem . . . . .	107
3.2.3	Kernel-Based Elliptic Inverse Problem on a Point Cloud . . . . .	109
3.3	Analysis of Kernel Approximation to the Forward and Inverse Problem . . . . .	114
3.3.1	Forward Map Approximation . . . . .	114
3.3.2	Posterior Approximation . . . . .	117
3.4	Numerical Experiments . . . . .	118
3.4.1	Implementation . . . . .	119
3.4.2	One-Dimensional Elliptic Problem on an Unknown Ellipse . . . . .	121
3.4.3	Two-Dimensional Elliptic Problem on an Unknown Torus . . . . .	124
3.4.4	Two-Dimensional Elliptic Problem on an Unknown Artificial Surface . . . . .	127
3.4.5	Hierarchical Bayesian Formulation . . . . .	129
3.5	Conclusions and Future Work . . . . .	134
3.6	Appendix . . . . .	135
3.6.1	Proof of Lemma 3.3.2 . . . . .	135
3.6.2	Proof of Lemma 3.3.3 . . . . .	136
3.6.3	Proof of Lemma 3.3.5 . . . . .	140
3.6.4	Proof of Lemma 3.3.7 . . . . .	143
4	UNLABELED DATA HELP IN GRAPH-BASED SEMI-SUPERVISED LEARNING: A BAYESIAN NONPARAMETRICS PERSPECTIVE . . . . .	146
4.1	Introduction . . . . .	146
4.1.1	Related Work . . . . .	149
4.1.2	Main Contributions and Scope . . . . .	150
4.2	Prior Construction and Main Result . . . . .	151
4.2.1	Graph-based Prior . . . . .	152
4.2.2	Main Result . . . . .	156
4.3	Posterior Contraction: Background and Set-up . . . . .	158
4.3.1	General Principles . . . . .	158
4.3.2	Application to Our Setting . . . . .	160
4.4	Proof of the Main Result . . . . .	161

4.4.1	Reduction via Interpolation . . . . .	162
4.4.2	Regularity and Contraction Properties of the Limiting Field . . . . .	164
4.4.3	Convergence of Gaussian Fields in $L^\infty(\mu)$ . . . . .	170
4.4.4	Putting Everything Together . . . . .	181
4.5	Discussion . . . . .	182
5	FINITE ELEMENT REPRESENTATIONS OF GAUSSIAN PROCESSES: BALANCING NUMERICAL AND STATISTICAL ACCURACY . . . . .	184
5.1	Introduction . . . . .	184
5.2	Background and Problem Setting . . . . .	188
5.2.1	The Matérn Covariance Function and SPDE Representations . . . . .	188
5.2.2	Finite Element Representations of Matérn-type Gaussian Processes . . . . .	190
5.2.3	Gaussian Process Regression and Classification: Finite Element Representations . . . . .	191
5.2.4	Our Criterion: Matching Posterior Contraction Rates . . . . .	193
5.3	Main Results . . . . .	196
5.3.1	FEM Construction and Spectral Error Bounds . . . . .	198
5.3.2	Balancing Numerical and Statistical Errors . . . . .	201
5.4	Simulation Study . . . . .	205
5.4.1	The One-dimensional Case . . . . .	207
5.4.2	The Two-dimensional Case . . . . .	209
5.5	Discussion and Open Directions . . . . .	211
5.6	Proof of Main Results . . . . .	214
6	LOCAL REGULARIZATION OF NOISY POINT CLOUDS: IMPROVED GLOBAL GEOMETRIC ESTIMATES AND DATA ANALYSIS . . . . .	224
6.1	Introduction . . . . .	224
6.1.1	Framework . . . . .	225
6.1.2	Main results . . . . .	228
6.1.3	Related and Future Work . . . . .	233
6.1.4	Outline . . . . .	236
6.2	Distance Approximation . . . . .	236
6.2.1	Geometric Preliminaries . . . . .	236
6.2.2	Local Distributions . . . . .	238
6.2.3	Bounding Expected Conditional Noise . . . . .	241
6.2.4	Bounding Difference in Geometric Bias . . . . .	243
6.2.5	Bounding Sampling Error . . . . .	256
6.3	From Local Regularization to Global Estimates . . . . .	258
6.4	Numerical Experiments . . . . .	260
6.4.1	Distance & Spectrum . . . . .	261
6.4.2	Classification . . . . .	263
6.4.3	Future Directions . . . . .	269
6.5	Appendix: Estimating $r_-$ and $r_+$ . . . . .	271
6.5.1	Estimating $r_-$ . . . . .	271

6.5.2 Estimating $r_+$ . . . . .	273
REFERENCES . . . . .	275

## LIST OF FIGURES

1.1	Graph approximations of the sphere. The edges encode information about closeness on the manifold and recover the geometry as the number of nodes increases to infinity. . . . .	2
1.2	Plots of samples of (1.5) for different $\tau$ 's and $s$ 's when the graph nodes are from the unit circle. The second row unfolds the plots in the first to the interval $[0, 2\pi]$ for better visualization of the fluctuations. . . . .	5
2.1	Random draws from nonstationary GFs on the unit circle; (a). $\kappa = 0.01; \tau = \exp(\cos(x))$ ; (b). $\kappa = \exp(\cos(x)), \tau = 100$ . . . . .	33
2.2	The three rows represent simulations for $n = 2500, 3600, 4900$ respectively. Within each row, the plots represent (1) Covariances: theoretical (blue) vs. approximation (red). (2) Spectra of $-\Delta$ (blue) and $\Delta_n$ (red). (3). Covariances (truncated): theoretical (blue) vs. approximation (red). (4). Sparsity pattern of $\Delta_n$ . . . . .	43
2.3	The three rows represent simulations for $n = 2892, 3612, 4842$ respectively. Within each row, the plots represent (1) Visualization of the triangulation. (2) Covariances: theoretical (blue), FEM-based (red), and graph-based (yellow). (3). Sparsity pattern of FEM-based precision. (4). Sparsity pattern of graph-based precision. . . . .	45
2.4	Posterior means and 95% credible intervals for (a) signal from nonstationary model; (b) length scale of nonstationary model; (c) signal from stationary model (constant $\tau_n$ ) when length scale is modeled through $\tau_n$ . . . . .	60
2.5	Posterior means and 95% credible intervals for (a) signal from nonstationary model; (b) length scale of nonstationary model; (c) signal from stationary model constant $\kappa_n$ ), when length scale is modeled through $\kappa_n$ . . . . .	60
3.1	Posterior means and 95% credible intervals for different $\sigma$ 's and $J$ 's. Here $\bar{\kappa}$ , $\kappa_{0.025}$ and $\kappa_{0.975}$ represent the posterior mean, 2.5% and 97.5% posterior quantiles respectively. . . . .	123
3.2	Posterior means and 95% credible intervals for $\sigma = 0.1$ and different $J$ 's. Here $\kappa_{0.025}$ and $\kappa_{0.975}$ represent the 2.5% and 97.5% posterior quantiles respectively. . . . .	124
3.3	Posterior means (first row) and standard deviations (second row) of posteriors for different $\sigma$ 's. horizontally. We have extended by interpolation the point cloud solution in order to ease visualization. . . . .	126
3.4	Reconstruction for the cow-shaped manifold. . . . .	127
3.5	Posterior means and 95% credible intervals for different truths. Figure are arranged so that the first two rows correspond to non-hierarchical and hierarchical respectively, and the third row shows the sample paths for $\tau$ . The three columns represent the truths $e^{\cos(\omega)}, e^{\cos(5\omega)}, e^{\cos(8\omega)}$ respectively. . . . .	131
3.6	Reconstruction of $e^{\cos(\omega)}$ and sample path for $\tau$ when $\sigma = 0.1$ . . . . .	133

5.1	The three columns represent simulations for $\kappa_0 = 1, 5$ and $25$ respectively with $s_0 = 2$ in all cases. The upper row shows plots of $f_0$ . The lower row compares the estimation error $\ \hat{\mathbf{f}} - \mathbf{f}_0\ _N$ between the covariance function (CF) approach and the finite element (FE) approach as $n_h$ increases, for three levels of data $N = 50, 500$ and $5000$ . . . . .	207
5.2	Comparison of estimation error $\ \hat{\mathbf{f}} - \mathbf{f}_0\ _N$ between the covariance function (CF) approach and the finite element (FE) approach on three data levels $N = 50, 500, 5000$ for (a) $s_0 = 1, \kappa_0 = 1$ and (b) $s_0 = 3, \kappa_0 = 25$ . . . . .	209
5.3	Comparison of the estimation error $\ \hat{\mathbf{f}} - \mathbf{f}_0\ _N$ as $n_h$ increases between the covariance function (CF) approach and two finite element approaches where FE1 is computed over $[-\rho, L + \rho]^2$ and FE2 is computed over $[-0.1\rho, L + 0.1\rho]^2$ . Simulation results for different combinations of $s_0$ and $\kappa_0$ are shown. . . . .	210
6.1	Comparison of spectra of continuum Laplacian, $\Delta_{\mathcal{X}_{n,\varepsilon}}, \Delta_{\mathcal{Y}_{n,\varepsilon}}$ and $\Delta_{\bar{\mathcal{Y}}_{n,\varepsilon}}$ for different values of $\sigma$ . . . . .	263
6.2	Classification error rates for $\Gamma_{\mathcal{X}_n}, \Gamma_{\mathcal{Y}_n}$ and $\Gamma_{\bar{\mathcal{Y}}_n}$ on two moons for different values of $\sigma$ . . . . .	266
6.3	Visualization of the point clouds $\mathcal{X}_n, \mathcal{Y}_n$ , and $\bar{\mathcal{Y}}_n$ . Each row contains scatter plots of the first two coordinates of the points in the datasets $\mathcal{X}_n, \mathcal{Y}_n$ , and $\bar{\mathcal{Y}}_n$ . . . . .	267
6.4	Visualization of the regularization effects. The second row is the regularized version of the corresponding image in the first row. While arguably more blurred, the digits in the second row are more homogeneous within each group, making classification easier. . . . .	269



## LIST OF TABLES

2.1	Comparison of the four models through RMSE, CRPS and LS. . . . .	63
2.2	Classification error rates with 2% labeled data for different pairs of digits with hyperprior $\log s \sim \mathcal{N}(\log 4, 1)$ . . . . .	68
2.3	Classification error rates with 2% labeled data for different pairs of digits with hyperprior $\log s \sim \mathcal{N}(\log 4, 0.01)$ . . . . .	68
3.1	Relative error of $\bar{\kappa}$ and $\bar{u}$ for different noise level, $\sigma$ 's and number of observations, $J$ , where $\bar{\kappa}$ and $\bar{u}$ are the posteriors means for $\kappa$ and $u$ respectively. In the last row, the relative noise level for each $\sigma$ is reported for diagnostic purposes. Particularly, note that the reconstruction error for $u^\dagger$ is much smaller than the relative noise level. . . . .	122
3.2	Relative error of $\kappa^\dagger$ and $\bar{u}$ for different noise level, $\sigma$ 's. In the last row, the relative noise level is reported for diagnostic purpose. Particularly, note that the reconstruction error for $u^\dagger$ is much smaller than the relative noise level. . . . .	125
3.3	Relative error of $\bar{\kappa}$ and $\bar{u}$ for different truths. Here "N" and "H" stand for non-hierarchical and hierarchical respectively. In the last row, the relative noise level for each $\sigma$ is reported for diagnostic purposes. . . . .	133
6.1	Entrywise $\infty$ -norm of $D\mathcal{X}_n - D\mathcal{Y}_n$ and $D\mathcal{X}_n - D\bar{\mathcal{Y}}_n$ on $\mathcal{S}$ for several $\sigma$ 's. . . . .	262
6.2	Entrywise $\infty$ -norm of $D\mathcal{X}_n - D\mathcal{Y}_n$ and $D\mathcal{X}_n - D\bar{\mathcal{Y}}_n$ on two moons for different values of $\sigma$ . . . . .	266
6.3	Classification error for different pairs of digits 3&8, 5&8, 4&9, and 7&9. . . . .	268
6.4	Classification error for 4&9 with different number of labels. . . . .	268
6.5	Comparison of classification errors with 4% labeled data. . . . .	271

## ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor Daniel Sanz-Alonso. I feel truly fortunate to be able to work with him and my Ph.D. work would be impossible without his continued support. He is always patient and generous with providing feedback and sharing valuable career advice. Under his guidance, I gradually learned about the academic world and decided my career path.

I would like to thank Bryon Aragam, with whom I learned a lot about research from the statistics side that complements my background. He is always passionate about sharing knowledge and career advice that helped pave my way. Thanks to my other wonderful collaborators Nicolás García Trillos and John Harlim for many insightful discussions.

I would like to thank all the faculty and staff in the Committee on Computational and Applied Mathematics and the Department of Statistics. Thanks to Mary Silber, our Program Director, for her continuous care and support throughout my whole Ph.D. study. Thanks to Guillaume Bal and Lek-Heng Lim for being in my thesis committee and for their help along the way.

I would like to thank my peers Zhisheng Xiao, Pinhan Chen, Qing Yan, and many others for the wonderful time that we have spent together. It is a great fortune to know these brilliant people, who set a standard for me. Thanks to my early roommates Yucheng Deng and Wenjun Cai for the time that we spent cooking and chatting together.

Finally, I would like to thank my parents for their everlasting love and support. Thanks to my buddy Shihao Chen who I knew from childhood for those Dota2 time.

## ABSTRACT

Matérn Gaussian fields are popular modeling choices in many aspects of Bayesian inverse problems, spatial statistics, machine learning, and numerous other scientific applications. In this thesis we investigate their generalizations to graphical domains, referred to as graph Matérn fields, by addressing their construction, application and theoretical properties. Graph Matérn fields share qualitatively similar features with the usual Matérn Gaussian fields on Euclidean spaces that are desirable for modeling purposes and enjoy a sparsity property that facilitates computation. Their wide applicability is demonstrated through several applications including precipitation modeling, an elliptic inverse problem, and semi-supervised classification, bridging together, and promoting exchange of ideas between spatial statistics, Bayesian inverse problems, and graph-based machine learning. Under an assumption that the graph nodes are sampled from a low-dimensional manifold, we show that our graph Matérn fields are consistent approximations of certain Matérn-type Gaussian fields defined over the underlying manifold. Study of the approximation error leads to new insights for the role of the unlabeled data in graph-based semi-supervised learning. Finally we complement the graph-based methods with a denoising algorithm that provably improves the performance when the graph nodes are noisy perturbations of manifold samples, which represents a more realistic scenario in many applications.

# CHAPTER 1

## INTRODUCTION

### 1.1 Matérn Gaussian Fields—Beyond the Euclidean Setting

A Gaussian field (GF)  $\{u(x)\}_{x \in \mathcal{I}}$  is a collection of random variables where the joint distribution of any finite subcollection is a multivariate Gaussian, which can be completely specified through a mean function  $\mu(\cdot)$  and a covariance function  $c(\cdot, \cdot)$ . GFs play a central role in modeling functions in many aspects of spatial statistics, Bayesian inverse problems, machine learning, and a variety of other scientific and engineering applications [Stein, 1999a, Williams and Rasmussen, 2006, Stuart, 2010, Sullivan, 2015]. For instance, they are used to interpolate data through a procedure called *kriging* in spatial statistics, or equivalently *Gaussian process regression* in the machine learning community, and they serve as standard priors in many statistical inverse problems including remote sensing and ground prospecting. A wide range of covariance functions can be adopted to reflect one’s prior belief on the fluctuations of the functions to be modeled, and such flexibility along with the analytical properties given by the Gaussianity makes GFs attractive and useful in practice.

Of particular interest is the family of the so-called *Matérn Gaussian fields* (MGFs), which consists of GFs with the Matérn covariance function, defined as

$$c_{\text{Matérn}}(x, \tilde{x}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} (\kappa|x - \tilde{x}|)^\nu K_\nu(\kappa|x - \tilde{x}|), \quad x, \tilde{x} \in \mathbb{R}^d. \quad (1.1)$$

Here  $|\cdot|$  denotes the Euclidean norm on  $\mathbb{R}^d$ ,  $\Gamma$  is the gamma function, and  $K_\nu$  is the modified Bessel function of the second kind. The three parameters  $\sigma, \nu, \kappa$  control respectively the marginal variance, sample path smoothness, and correlation lengthscale. In a more intuitive sense, they determine respectively the overall magnitude of the sample paths, number of times that they can be differentiated, and how oscillatory they look like. MGFs form a rich class of random fields that are able to model a wide range of phenomena. As noted in Stein

[1999a], one of the key properties that make MGFs attractive is their finite smoothness: the samples paths associated with (1.1) are (roughly speaking)  $\nu - 1$  times differentiable. In contrast, the squared exponential covariance function  $c_{\text{SE}}(x, \tilde{x}) = \exp\left(-\frac{|x-\tilde{x}|^2}{2\kappa^2}\right)$  leads to infinitely differentiable sample paths, which is less realistic for many practical applications.

Given the flexibility of MGFs, a natural question is whether one can construct similar Matérn type GFs on domains that are not necessarily Euclidean. In this thesis, we shall focus on graphical domains, motivated by two reasons. The first reason is that graphs are natural approximations of manifolds (see Figure 1.1 from García Trillos et al. [2020b]), which are important non-Euclidean domains in many applications. For instance, when interpolating




Figure 1.1: Graph approximations of the sphere. The edges encode information about closeness on the manifold and recover the geometry as the number of nodes increases to infinity.

spatial data along the Earth surface—such as precipitation or temperature—it is desirable to have GFs defined over the sphere. However, simply plugging the geodesic distances into the Euclidean covariance functions may no longer retain positive definiteness [Banerjee, 2005], and using the Euclidean distances fails to account for the geometry of manifold data. Hence generalizing the MGFs to manifolds is itself important and nontrivial. But in many cases one does not have access to the whole manifold but only unstructured samples from it. Therefore, it would be desirable to construct GFs over the graph built from the samples as approximations of GFs over the underlying manifold. The second reason is that when analyzing discrete objects such as social networks, the index set comes from a similarity

graph that is not easily embedded into a Euclidean space, limiting the use of the covariance functions. It would then be interesting to construct discretely indexed GFs that reflect the correlations between the nodes encoded by their pairwise similarities. Motivated by these questions, we shall study in this thesis the construction of Matérn type GFs over graphs, demonstrate their wide applicability in inverse problems and machine learning, and investigate their theoretical properties. The following Sections 1.2, 1.3 and 1.4 give an overview of these aspects.

## 1.2 Graph Matérn Fields

### 1.2.1 The Series Representation

Let's start by briefly discussing the main ideas behind their construction, motivated by the seminal work of Lindgren et al. [2011], which exploits the stochastic partial differential equation (SPDE) representation of MGFs on Euclidean spaces discovered by Whittle [1963]. Roughly speaking, a MGF  $u$  on  $\mathbb{R}^d$  is the unique stationary solution to

$$(\kappa^2 - \Delta)^{\nu/2+d/4}u(x) = \mathcal{W}(x), \quad x \in \mathbb{R}^d, \quad (1.2)$$

where  $\Delta$  is the usual Laplacian and  $\mathcal{W}$  is spatial white noise. Therefore, analysis and approximation of the original MGFs can be equivalently carried out on the above equation. Furthermore, as noted in Lindgren et al. [2011] the SPDE representation sidesteps the covariance function perspective and opens doors for various generalizations. In particular, one can define Matérn type GFs beyond the Euclidean setting by adjusting accordingly the SPDE (1.2), e.g., by lifting the equation to a manifold. Moreover, the usual MGFs are stationary in the sense that the covariance (1.1) between two points depend only on their distance and it is generally difficult to introduce nonstationarity by modifying the covariance function. However, the SPDE representation further allows seamless incorporation of nonstationar-

ity by letting the constant parameter  $\tau$  depend on  $x$  and replacing the Laplacian with an anisotropic elliptic operator such as  $\nabla \cdot (\mathbf{H}\nabla)$  for some matrix-valued function  $\mathbf{H}$  [Fuglstad et al., 2015a, Bolin and Kirchner, 2020]. Therefore, the SPDE representation creates a whole range of possibilities from a modeling perspective.

Motivated by these ideas, we shall construct a nonstationary Matérn type GF over a graph, which we refer to as a *graph Matérn field* in the following. Let  $(V, E)$  be an undirected graph with  $n$  nodes. Formally, we shall consider the following graphical analog of equation (1.2) (after redefining the parameters  $\tau, s > 0$ )

$$(\tau I_n + \Delta_n)^{s/2} u_n = \mathcal{W}_n, \quad (1.3)$$

where  $I_n$  is the  $n$ -dimensional identity matrix,  $\mathcal{W}_n \sim \mathcal{N}(0, I_n)$  and  $\Delta_n \in \mathbb{R}^{n \times n}$  is a graph analog of  $-\Delta$  to be determined. It turns out that tools from spectral graph theory are readily applicable here and what we need for  $\Delta_n$  is the so-called *graph-Laplacian*. We shall discuss it in more detail shortly in Section 1.2.2, but let's take for granted for the moment that a graph-Laplacian is a symmetric positive semidefinite matrix that shares many qualitative features with the usual Laplacian. In particular, it admits a spectral decomposition with nonnegative eigenvalues and the associated eigenvectors form an orthonormal basis of  $\mathbb{R}^n$ . The discrete equation (1.3) can then be equivalently seen as the multivariate Gaussian distribution

$$u_n \sim \mathcal{N}(0, (\tau I_n + \Delta_n)^{-s}), \quad (1.4)$$

which has a Karhunen-Loève expansion

$$u_n = \sum_{i=1}^n (\tau + \lambda_{n,i})^{-s/2} \xi_i \psi_{n,i}, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (1.5)$$

where  $(\lambda_{n,i}, \psi_{n,i})$ 's are the ordered eigenvalue-eigenvector pairs of  $\Delta_n$ . We shall treat (1.5)

as our definition for the graph Matérn field and remark that it retains the qualitative features of the MGFs on Euclidean domains. Indeed, the ordered eigenvector  $\psi_{n,i}$ 's are becoming more oscillatory as  $i$  increases, just as their counterparts for the usual Laplacian on bounded domains. Therefore, a larger  $s$  implies faster decay of the coefficients and a more regular sample path, whereas a larger  $\tau$  incorporates more essential frequencies and makes the sample path more oscillatory. Figures 1.2a, 1.2b and 1.2c demonstrate this behavior for three sets of parameters when the graph nodes are points on a unit circle. Nonstationary generalizations can be defined similarly by replacing  $\tau I_n$  in (1.4) with a diagonal matrix  $\tau_n$  and coming up with a graph analog of the elliptic operators. In this case, the different values of  $\tau_n$  can be interpreted as local lengthscales around each node, as shown in Figure 1.2d, where  $\tau_n$  increases from 1 (the left end) to 30 (the right end). More details of the construction will be discussed in Section 2 along with several applications.

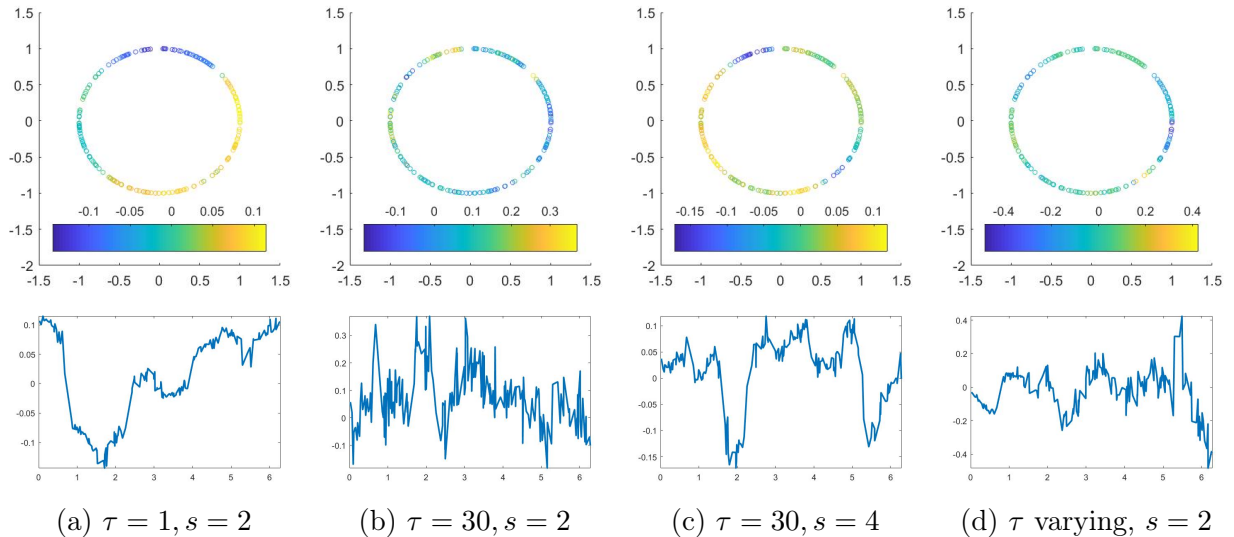


Figure 1.2: Plots of samples of (1.5) for different  $\tau$ 's and  $s$ 's when the graph nodes are from the unit circle. The second row unfolds the plots in the first to the interval  $[0, 2\pi]$  for better visualization of the fluctuations.



### 1.2.2 The Graph-Laplacian

We shall now introduce the key object in the above construction—the graph-Laplacian. Let  $(V, E)$  be an undirected weighted graph with  $n$  nodes whose edge weights are given by  $W \in \mathbb{R}^{n \times n}$ . We shall also refer to  $W$  as the similarity matrix since  $W_{ij}$  can be interpreted as the similarity between nodes  $i$  and  $j$ . The *unnormalized graph-Laplacian* is then defined as the matrix  $\Delta_n = D - W$ , where  $D$  is the diagonal matrix whose entries are  $D_{ii} = \sum_{j=1}^n W_{ij}$ . Several normalized versions can also be constructed (see e.g. Von Luxburg [2007] ), but for our purpose we shall focus on the unnormalized one. One readily sees that  $\Delta_n$  is positive semidefinite by the following relation

$$v^T \Delta_n v = \frac{1}{2} \sum_{i,j=1}^n W_{ij} |v(i) - v(j)|^2, \quad v \in \mathbb{R}^n. \quad (1.6)$$

Moreover, (1.6) implies that if we identify  $v$  with a function over the nodes  $V$ , then those  $v$ 's that change slowly with respect to the similarities lead to smaller values of  $v^T \Delta_n v$ . Therefore, Gaussian distributions of the form  $\mathcal{N}(0, \Delta_n^{-1})$  ( $\Delta_n^{-1}$  representing the pseudoinverse) would favor these “smooth”  $v$ 's since its negative log density is proportional to (1.6) (up to an additive constant), and our previous construction (1.4) can be seen as generalizing such smoothness constraint.

For many of the applications that we shall consider, a typical case is that we are given the graph nodes  $V$ , usually as a collection of data feature  $x_i$ 's, and we need to define the pairwise similarities by ourselves to reflect the structures of the specific problems. For instance when the  $x_i$ 's are Euclidean vectors, common choices are the  $\epsilon$ -graphs, where  $W_{ij} = \mathbf{1}\{|x_i - x_j| < \epsilon\}$ , and  $k$ -nearest neighbor ( $k$ -NN) graphs, where  $W_{ij} = \mathbf{1}\{x_i \text{ and } x_j \text{ are } k\text{-NN of each other}\}$ , among many others.

Of particular interest is the case where the  $x_i$ 's are further sitting on a manifold. Besides the modeling of the Earth surface mentioned above, such *manifold assumption* [Belkin and

[Niyogi, 2004] has been adopted frequently in the literature. The rationale is that one often encounters data features that live in high dimensional ambient spaces but carry certain low dimensional structures. An interesting example is the MNIST dataset consisting of images of hand-written digits from 0 to 9. Each image is a 784 (28-by-28 numerical matrix) dimensional vector, but one should expect that the intrinsic dimension of this dataset is much smaller since the shape of the digits (say annular for the digit 0) substantially limits the degrees of freedom they can take. Therefore, the manifold assumption is an idealization of such low-dimensional structure.

Under the manifold assumption, one can further establish certain theoretical properties of the graph-Laplacian. Suppose now we are given  $\{x_i\}_{i=1}^n$  that are sampled from a submanifold  $\mathcal{M}$  of  $\mathbb{R}^d$ . Define the similarities by

$$W_{ij} = \frac{2(m+2)}{n\nu_m h_n^{m+2}} \mathbf{1}\{|x_i - x_j| < h_n\}, \quad (1.7)$$

where  $m$  is the dimension of  $\mathcal{M}$ ,  $\nu_m$  is the volume of the  $m$ -dimensional unit ball, and  $h_n$  is a user-chosen graph connectivity parameter. Under suitable assumptions, it can be shown that the graph-Laplacian  $\Delta_n$  constructed using the similarities (1.7) converges to the Laplace-Beltrami operator  $\Delta_{\mathcal{M}}$  on  $\mathcal{M}$  (where we take the convention that  $\Delta_{\mathcal{M}}$  is positive semidefinite). Here by convergence we mean both in the pointwise sense [Hein et al., 2005, 2007], meaning that  $\Delta_n f(x) \xrightarrow{n \rightarrow \infty} \Delta_{\mathcal{M}} f(x)$  for smooth enough  $f$ 's, and in spectral sense [Burago et al., 2015, García Trillos et al., 2020a], meaning that the eigenvalues and eigenfunctions of  $\Delta_n$  converge (in an appropriate sense) to those of  $\Delta_{\mathcal{M}}$ . Recent work also shows that for a  $k$ -NN construction of the graph, these properties are retained [Calder and García Trillos, 2022]. Therefore in such cases, the graph-Laplacian is indeed a discrete analog of the usual Laplacian as we have assumed in Section 1.2. Furthermore, it can be expected that the discretely indexed GF (1.5) will converge to a continuum GF defined over  $\mathcal{M}$ . We shall come back to this point later in Section 1.4.

### 1.2.3 Computation

So far we have discussed the construction of our prior from a modeling perspective, by drawing connections with the MGFs on Euclidean spaces and the regularization power of the graph-Laplacian. We shall see now that our graph Matérn fields also facilitates computation. This lies in the crucial observation that the graph-Laplacians are usually sparse, as the weights matrices (e.g. the  $k$ -NN construction) usually are. In the case of manifold data, this can be easily seen from the definition (1.7). Indeed, it can be shown that for suitable choice of  $h_n$ , the number of nonzero entries of  $W$ , and hence  $\Delta_n$  is  $O(n^{3/2})$  while retaining an accurate approximation of  $\Delta_{\mathcal{M}}$ . Therefore for small integer  $s$ 's in (1.4), we are left with a sparse precision matrix for the Gaussian and numerical linear algebra methods can be employed to speed-up computation. For general  $s$ 's, a second layer of rational approximation proposed by Bolin and Kirchner [2020] can be adopted to retain sparsity.

In a broader sense, the graph Matérn fields constructed can be seen as an instance of the recent trend of leveraging sparsity for computational efficiency in GF methodologies. Despite their modeling flexibility, GFs suffer from high computational costs due to the need of factorizing the covariance matrix, which scales in general as  $O(n^3)$ . Therefore, several lines of works have been trying to alleviate this so called big- $n$  problem, including Vecchia approximations [Vecchia, 1988], screening effects [Stein, 2002], covariance tapering [Furrer et al., 2006], and low-rank approximations [Banerjee et al., 2008], among many others. The paper by Lindgren et al. [2011] shows that a finite element approximation to the SPDE (1.2) introduces sparsity and hence leads to a sparse approximation to the original MGF. In a similar sense, our graph Matérn fields can be seen as a sparse graph-based approximation to an SPDE that defines a continuum GF on some underlying space. Under the manifold assumption, our work can be seen as a generalization of theirs to the unstructured point cloud setting where finite element methods are not directly applicable.

On the other hand, GFs over graphs is still a developing area. Previous works in the

machine learning community mainly focus on objects like  $\mathcal{N}(0, (\tau I_n + \Delta_n)^{-1})$  by leveraging the regularization power of the graph-Laplacian [Zhu et al., 2003, Bertozzi et al., 2018, Ng et al., 2018]. Our construction generalizes these with nonstationarity and higher powers of graph-Laplacian by establishing a connection with the usual MGFs on Euclidean spaces, and provide an interpretation of the parameter  $\tau$  as local lengthscales instead of a nugget term to make the covariance positive definite as in most previous works. Furthermore, the connection with usual MGFs bridges the machine learning and spatial statistics communities and facilitates the exchange of ideas between them, as will be demonstrated through the latent Gaussian models in Section 2.

## 1.3 Applications

### 1.3.1 The Bayesian Perspective

Having talked about the construction and qualitative properties of the graph Matérn fields, we shall now illustrate some of their applications. Suppose we are given data pairs  $(x_i, y_i) \stackrel{i.i.d.}{\sim} \text{Law}(X, Y)$  for  $i = 1, \dots, n$ . We shall refer to the  $x_i$ 's as features and  $y_i$ 's as labels. Common examples of the label-generating mechanism include the regression and binary classification problems, where for some function  $u$

$$\begin{cases} y_i = u(x_i) + \eta_i, & \eta_i \sim \mathcal{N}(0, \sigma^2) & \text{regression} \\ \mathbb{P}(y_i = 1) = u(x_i) & & \text{classification.} \end{cases} \quad (1.8)$$

We shall interpret the function  $u$  as an underlying true parameter of the system that generates the data, and a typical goal is then to infer  $u$  based on the data pairs  $\{(x_i, y_i)\}_{i=1}^n$ . More sophisticated label-generating mechanism includes the solution operator of a partial differential equation (PDE) that appears in inverse problems as we shall discuss shortly in Section 1.3.2, where the function  $u$  will be the coefficient of the PDE.

Since the support of  $X$  is rarely known ahead, the more common goal would be to infer the vector  $u_n = (u(x_1), \dots, u(x_n))^T$  and we shall adopt a Bayesian approach to this problem. The starting point is to construct a *prior* for  $u_n$ , which is a probability distribution over  $\mathbb{R}^n$  (the space of all possible realizations of  $u_n$ ) that encodes our prior belief about how it might look like before seeing the label  $y_i$ 's. For instance we may have prior knowledge about the smoothness properties or the lengthscales of the underlying function  $u$ . Then by conditioning on the labels, we obtain through the Bayes formula the *posterior*, which is again a probability distribution over  $\mathbb{R}^n$ , satisfying

$$\underbrace{\mathbb{P}(u_n|y_1, \dots, y_n)}_{\text{posterior}} \propto \underbrace{\mathbb{P}(y_1, \dots, y_n|u_n)}_{\text{likelihood}} \underbrace{\mathbb{P}(u_n)}_{\text{prior}}. \quad (1.9)$$

Here  $\mathbb{P}(y_1, \dots, y_n|u_n)$  is the *likelihood* of the labels given that they are generated from  $u_n$  and factors into  $\prod_{i=1}^n \mathbb{P}(y_i|u_n)$  if independence of the  $y_i$ 's are assumed. In the setting of (1.8), the likelihoods take the form

$$\mathbb{P}(y_i|u_n) = \begin{cases} (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{|y_i - u_n(i)|^2}{2\sigma^2}} & \text{regression,} \\ u_n(i)^{y_i} [1 - u_n(i)]^{1-y_i} & \text{classification.} \end{cases}$$

The posterior represents our updated belief about  $u_n$  that is incurred by the labels. Roughly speaking, the posterior density will be large for those  $u_n$ 's that are consistent both with the labels and with our prior beliefs. Inference of the true  $u_n$  is then carried out based on point estimators such as the posterior mean or mode, and Bayesian confidence intervals computed based on the posterior.

A closely related alternative to the Bayesian perspective is the optimization approach,

which on a high level estimates the true  $u_n$  by solving the following problem

$$\hat{u}_n = \arg \max_{v \in \mathbb{R}^n} [\log \mathbb{P}(y_1, \dots, y_n | v) + R(v)]. \quad (1.10)$$

Here  $R(\cdot)$  is a *regularizer* to prevent overfitting, which plays a similar role as our prior. Indeed, the maximizer of (1.10) is conceptually equivalent to the posterior mode of (1.9) if the prior is taken to satisfy  $\mathbb{P}(v) \propto e^{R(v)}$ . A disadvantage of the optimization approach is that for complicated likelihoods such as those that involve solving a PDE, (1.10) is highly nonconvex and may be difficult to solve. However in the Bayesian perspective, Markov chain Monte Carlo methods can be employed to estimate the posterior mean instead of finding the mode, which alleviates such issue in certain cases, although the computational cost is generally higher due to the need of running a long Markov chain.

It turns out that the choice of prior is crucial in the Bayesian approach and this is where our graph Matérn fields come in. More precisely, we shall use our graph Matérn fields constructed in the previous section as prior models for  $u_n$ . In the following we shall briefly discuss two problems that are well-suited for their application, with more applications deferred to Section 2.

### 1.3.2 An Elliptic Inverse Problem

The first problem that we shall consider is an *elliptic inverse problem*. To be more concrete, consider the following diffusion equation

$$-\operatorname{div}(e^u \nabla p) = f, \quad x \in \mathcal{M}, \quad (1.11)$$

defined over a manifold  $\mathcal{M}$ , and we are given points  $\{x_i\}_{i=1}^n$  on  $\mathcal{M}$  with noisy observations

$$y_i = p(x_i) + \eta_i, \quad \eta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2), \quad i = 1, \dots, n$$

of the solution  $p$ . The goal is then to recover the coefficient function  $u$  based on the data pairs  $\{(x_i, y_i)\}_{i=1}^n$ . Possible application of the problem is in underground oil reservoir modeling, where  $p$  represents the pressure of the oil and the goal is to recover the permeability  $e^u$  of the rocks for better exploitation purposes. The manifold setup is to account for possible nonflat geometry underground, and because of this, a systematic representation of  $\mathcal{M}$  may not be available but only unstructured sample  $x_i$ 's are obtained which can be thought of as the locations of the underground sensors. Therefore, constructing a prior model for functions on all of  $\mathcal{M}$  may be infeasible while our graph Matérn fields then circumvent such issue with a graph-based approximation.

The key ingredient in solving an inverse problem is a solver for the forward problem, or in other words, given a function  $u$ , we need to solve or approximately solve the PDE (1.11). More precisely, let  $\mathcal{F} : \mathcal{B} \rightarrow L^2(\mathcal{M})$  denote the *forward map* that maps a coefficient function  $u$  in some Banach space  $\mathcal{B}$  to the corresponding solution  $p$  of the PDE. Since  $\mathcal{F}$  is rarely known analytically, approximations are necessary for practical computations. Due to the lack of a structured representation of  $\mathcal{M}$  such as a triangulation, traditional numerical solvers such as the finite element methods are not directly applicable here. Therefore in Section 3, a new graph-based discretization of the PDE (1.11) is introduced. The idea is based on the following simple observation that the diffusion operator can be written as

$$-\operatorname{div}(e^u \nabla p) = \sqrt{e^u} [\Delta_{\mathcal{M}}(p\sqrt{e^u}) - p\Delta_{\mathcal{M}}\sqrt{e^u}],$$

i.e., as a difference between two Laplacian-like terms. Since we know how to approximate  $\Delta_{\mathcal{M}}$  with  $\Delta_n$ , the approximation for the diffusion operator follows accordingly. We shall formalize these ideas in Section 3, together with numerical experiments to demonstrate their effectiveness in several examples including a complicated artificial manifold surface.

The Bayesian approach to inverse problems [Kaipo and Somersalo, 2006, Stuart, 2010] has received increasing attention over the past decades. Besides the advantage of having a

whole posterior that allows uncertainty quantification, the Bayesian approach further leads to a well-posed problem, meaning that a small perturbation in the data, the prior, or the forward map incurs a small perturbation in the posterior. In Section 3 we also establish such a result by bounding the error introduced from our approximate forward map. However, as mentioned above, the advantages of the Bayesian perspective comes with the generally high cost of the associated sampling algorithms. Moreover, theoretical properties such as consistency and convergence rates of the Bayesian approach to inverse problems are still emerging [Abraham and Nickl, 2020, Giordano and Nickl, 2020, Monard et al., 2021].

### 1.3.3 *Semi-supervised Learning*

The second line of problems that we shall consider lies in the area of *semi-supervised learning* (SSL). In short, SSL shares the same goal of inferring the underlying function  $u$  as above, but exploits further a collection of unlabeled data. In other words, we are given both labeled  $\{(x_i, y_i)\}_{i=1}^n$  and unlabeled data  $\{x_i\}_{i=n+1}^N$ , with the goal being to estimate the vector  $u_N = (u(x_1), \dots, u(x_N))^T$ . SSL arises with the fact that collecting unlabeled data can be done cheaply while assigning labels to them is a much more expensive process. For instance, categorizing websites and texts requires the time-consuming process of going through the contents, and labeling protein sequences would require extensive expert knowledge and experimentation. Therefore,  $N$  is usually much larger than  $n$  and the goal of SSL is to boost the learning performance by taking advantage of the abundance of unlabeled data.

The question of whether unlabeled data can indeed help has been widely debated, and different conclusions have been reached under different model assumptions and methodologies adopted. For our purpose, the manifold assumption mentioned above comes up again, in which case the graph-based Bayesian approach fits naturally. The intuition is that the additional unlabeled data add to our understanding of the underlying geometry and allow us to construct a more accurate model for  $u_N$ . Imagine in the limiting scenario with infinitely



many unlabeled data, we would be able to resolve the underlying manifold completely and construct instead a continuum GF over the manifold. We shall come back to these ideas in Section 1.4 and more rigorously in Section 4.

The graph-based methods is one of the most widely adopted approaches to SSL (see Van Engelen and Hoos [2020] for a review of the various methods). The overall idea is to regularize the learning problem based on the geometric information extracted from the unlabeled data and use the similarity information to propagate the labels. In the parallel optimization perspective, the regularizer is usually set to be  $R(v) = v^T \Delta_N v$ , which according to (1.6) would impose a penalty for vectors whose variation are inconsistent with the similarity patterns. Borrowing intuitions from the usual Laplacian, this penalty term can be further interpreted as penalizing certain discrete derivatives of the vector or equivalently a function over the graph nodes, just as the quadratic form  $\langle f, -\Delta f \rangle$  would do. As mentioned, the graph Matérn fields can be seen to generalize such regularization effect on the level of the prior model for  $u_N$ . Part of the novelty of our graph Matérn fields lies in the nonstationarity, i.e., Gaussian distributions of the form  $\mathcal{N}(0, (\tau_N + \Delta_N)^{-s})$  where  $\tau_N$  is a possibly nonconstant vector representing the local lengthscales at each node, as opposed to the usual graph-based prior  $\mathcal{N}(0, (\tau I_N + \Delta_N)^{-s})$  for a single uniform lengthscale  $\tau$ . We demonstrate in Section 2 that such a nonstationary modeling leads to improved performance over the standard stationary one for a semi-supervised classification problem, which shows their practical relevance for other potential machine learning applications.

## 1.4 Continuum Limits

Having discussed about the applications of our graph Matérn fields, we now turn to their theoretical properties. We shall make the manifold assumption throughout this section and investigate the asymptotic regime where the number of data features  $x_i$ 's increases to infinity, which we refer to as the continuum limit. In particular, we will see that the graph Matérn

fields converge to certain Matérn type GFs defined over the underlying manifold, and such convergence on the level of the prior in our Bayesian perspective leads to new insights about the corresponding posteriors in a SSL problem. We remark that although the manifold assumption facilitates the analysis, our graph Matérn fields are applicable in more general settings as we have seen above and as will be demonstrated in Section 2.

#### 1.4.1 Prior Level

Recall that our graph Matérn fields are constructed based on a graphical representation of the SPDE (1.2), which characterizes the usual MGFs on Euclidean spaces. It further admits a series representation (1.5) that we have used as the definition of our graph Matérn field. Therefore, a natural candidate for their continuum limit would be the manifold analog of the MGFs, which can be motivated similarly through the same SPDE by replacing the usual Laplacian with the Laplace-Beltrami operator. In particular, we shall consider formally the Gaussian measure  $\mathcal{N}(0, (\tau + \Delta_{\mathcal{M}})^{-s})$ . If we further assume  $\mathcal{M}$  is compact, then a similar series representation as (1.5) can be obtained

$$u = \sum_{i=1}^{\infty} (\tau + \lambda_i)^{-s/2} \xi_i \psi_i, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (1.12)$$

where  $(\lambda_i, \psi_i)$ 's are the eigenvalue-eigenfunction pairs of  $\Delta_{\mathcal{M}}$ . In a similar spirit as the graph Matérn field (1.5), the random field (1.12) has exactly the same qualitative features as a MGF defined over  $\mathcal{M}$ , with  $\tau$  and  $s$  controlling respectively the lengthscale and smoothness. In Section 2, we will show under suitable assumptions that

$$\mathbb{E} \|u_n \circ T_n - u\|_2^2 \xrightarrow{n \rightarrow \infty} 0, \quad (1.13)$$

where  $T_n : \mathcal{M} \rightarrow \{x_1, \dots, x_n\}$  is a sequence of *transport maps*. Roughly speaking, the transport maps can be interpreted as certain interpolation maps to reconcile the fact that  $u$

is a random field defined over all of  $\mathcal{M}$  while  $u_n$  is only defined over the  $x_i$ 's. The main novelty is a convergence result after the incorporation of nonstationarity, with the  $u$  defined (1.12) modified accordingly. An intuitive explanation of such convergence in the stationary case can be readily seen by comparing the series representations (1.5) and (1.12), and recalling that the construction of graph-Laplacians guarantees spectral convergence. The key point lies in a suitable choice of the graph connectivity parameter  $h_n$  in the definition of the weights (1.7), which needs to satisfy

$$\frac{(\log n)^{c_m}}{n^{1/m}} \ll h_n \ll \frac{1}{n^{1/2s}}, \quad (1.14)$$

with  $c_m = 3/4$  if  $m = 2$  and  $c_m = 1/m$  otherwise and  $m$  is the dimension of  $\mathcal{M}$ . The rationale for the choice (1.14) is that the lower bound characterizes the resolution of the  $x_i$ 's and  $h_n$  should be much larger than it to capture the local geometry, while the upper bound ensures that the higher frequencies of  $\Delta_n$  do not accumulate in the large  $n$  limit [García Trillos and Sanz-Alonso, 2018].

Our convergence result (1.13) can be seen as part of the recent line of work that studies discrete approximations of continuum GFs. Starting with the seminal work of Lindgren et al. [2011] that proposed a finite element approximation to the usual MGFs, several other authors have extended the ideas to study approximations of more general Matérn type GFs [Bolin et al., 2020, Bolin and Kirchner, 2020, Cox and Kirchner, 2020]. The sparsity structure introduced by the finite element approximation then facilitates computation. Our result shares a similar nature, but instead with a graph-based approximation, and the sparsity structure comes from that of the graph-Laplacian. Furthermore, our graph-based approximations complement and generalize the finite element approaches in high dimensional settings where the latter becomes less practical.

### 1.4.2 Posterior Level

The previous subsection considers the approximation properties of the graph Matérn fields, which can be seen as a consistency result on the level of the prior. A natural question would then be to understand the implications of such convergence on the corresponding posteriors when the graph Matérn fields are applied. It turns out that in the SSL setting that we have described above, we do obtain some new insights: the unlabeled data indeed help to improve the learning performance.

To illustrate the main idea, let's introduce some notations and concepts. Recall that in the SSL problem, we are given both labeled  $\{(x_i, y_i)\}_{i=1}^n$  and unlabeled data  $\{x_i\}_{i=n+1}^{N_n}$ , with the goal being to estimate the conditional mean function  $u_0(x) = \mathbb{E}[Y|X = x]$ . Notice that we have introduced a subscript for  $N$  since we are interested in understanding its dependency with respect to  $n$ . We adopt a Bayesian perspective by putting a graph Matérn prior  $\pi_{N_n}$  over  $u_0|_{\{x_1, \dots, x_{N_n}\}}$  that is constructed with all the  $x_i$ 's, which after conditioning on the  $y_i$ 's gives us a posterior that we denote as  $\mu_{N_n}^y$ . The quality of the posterior determines the quality of various estimators that we construct, and for good estimation performance, we would ideally want our posterior to be nearly supported around the true parameter  $u_0$ . The following concept, the so-called *posterior contraction rate* introduced by Ghosal et al. [2000], gives a quantitative measure of such concentration. We say that the sequence of posteriors  $\mu_{N_n}^y$  contracts around  $u_0$  with rate  $\epsilon_n$  if, for all sufficiently large  $M > 0$ ,

$$\mu_{N_n}^y(u \in \mathbb{R}^{N_n} : \|u - u_0\|_n \leq M\epsilon_n) \xrightarrow{n \rightarrow \infty} 1 \quad (1.15)$$

in probability, where

$$\|u - u_0\|_n^2 := \frac{1}{n} \sum_{i=1}^n |u(x_i) - u_0(x_i)|^2.$$

Here again we identify a vector in  $\mathbb{R}^{N_n}$  with a function over  $\{x_i\}_{i=1}^{N_n}$ . The convergence

(1.15) implies that asymptotically the sequence of posteriors  $\mu_{N_n}^y$  will be nearly supported on a ball of radius  $O(\epsilon_n)$  around  $u_0|_{\{x_1, \dots, x_{N_n}\}}$ . Therefore,  $\epsilon_n$  characterizes the rate at which the posterior “contracts” around  $u_0$  and can be intuitively interpreted as the convergence rate of the posterior distribution towards the truth. As a consequence of the convergence (1.15), the point estimator defined as

$$\hat{u}_n := \arg \min_{u \in \mathbb{R}^{N_n}} [\mu_{N_n}^y(v \in \mathbb{R}^{N_n} : \|v - u\|_n \leq M\epsilon_n)]$$

converges (in probability) to  $u_0$  with the same rate  $\epsilon_n$ . Therefore, the posterior contraction rates can be compared with the minimax theory of statistical inference to quantify the performance of Bayesian procedures.

In Section 4 we will show in the regression and binary classification setting that optimal contraction rates can be achieved if  $N_n \asymp n^{2m}$ , which suggests that unlabeled data help in this case. To be slightly more precise, we will show that if the truth  $u_0$  belongs to a “ $\beta$ -regular function class”, then the minimax optimal rate  $n^{-\beta/(2\beta+m)}$  can be achieved, provided that  $N$  grows sufficiently fast with  $n$ . Here by  $\beta$ -regular we mean a Besov-type function space with parameter  $\beta$  defined over the manifold that is introduced by Coulhon et al. [2012], Castillo et al. [2014], which roughly speaking represents functions that are  $\beta$ -times differentiable.

The rationale of our result lies in the connection between graph Matérn fields and their continuum limits. First, it can be shown that optimal contraction rates can be achieved if the continuum Matérn type GF (1.12) is used as the prior. Of course, this is an idealized setting when the underlying manifold is known exactly, but points us to a promising direction since the continuum limit of our graph Matérn field is a “good” prior for the learning problem. Then it comes the second important ingredient of the analysis, which is a result from [van der Vaart and van Zanten, 2008a, Theorem 2.2] that says if  $\Pi_n$  is a sequence of Gaussian priors approximating “sufficiently well” a limiting Gaussian prior  $\Pi$ , then the posterior contraction rates with respect to the  $\Pi_n$ ’s will be the same as that of  $\Pi$ . In our context,  $\Pi_n$  would be the

graph Matérn prior  $\pi_{N_n}$ , whose limit  $\Pi$  is given by (1.12), and the condition of sufficiently good approximation translates to the requirement that

$$\mathbb{E}\|u_{N_n} \circ T_{N_n} - u\|_2^2 \lesssim n^{-1} \quad (1.16)$$

following the framework developed in Section 1.4.1. A refined analysis of the consistency result (1.13) gives a rate

$$\mathbb{E}\|u_{N_n} \circ T_{N_n} - u\|_2^2 \lesssim N_n^{-\frac{1}{2m}}, \quad (1.17)$$

and therefore the required sample size for  $N_n$  is obtained by matching (1.16) and (1.17). We see that the convergence rate of the graph Matérn fields suffers from the curse of dimensionality, which is not too surprising since the resolution of the point cloud  $x_i$ 's scales as  $N_n^{-1/m}$ . Therefore, the overall intuition is that the slow convergence rate of the graph Matérn fields can be compensated by the abundance of the unlabeled data to still retain an accurate approximation of the continuum GFs, which leads to optimal estimation performance.

The question of whether unlabeled data help in SSL has been studied under different model assumptions and methodologies, leading to different conclusions [Wasserman and Lafferty, 2007, Singh et al., 2008, Niyogi, 2013]. Our result then suggests the effectiveness of unlabeled data in the specific setting of graph-based Bayesian approach under a manifold assumption. As we discussed, an important step in our analysis is to understand the continuum limits of the graph-based methods. This has been an active research area, including both a Bayesian perspective [García Trillos and Sanz-Alonso, 2018, García Trillos et al., 2020b] with a similar set up as ours, and an optimization perspective [Calder et al., 2020, Calder and Slepčev, 2020, Bungert et al., 2021] that draws connection with certain limiting PDEs over the underlying manifold. On the other hand, our analysis uses tools from the theory Bayesian nonparametrics and complements recent works on analyzing Bayesian estimation

approaches over graphs [Kirichenko and van Zanten, 2017] and manifolds [Castillo et al., 2014, Yang and Dunson, 2016].

## 1.5 Outline and Main Contributions

We now give an outline of the following sections and summarize their main contributions.

- Chapter 2 is based on Sanz-Alonso and Yang [2022a]. We study the construction of the graph Matérn fields and formalize a rigorous framework to study their convergence towards certain continuum GFs. Several applications demonstrate their wide applicability in spatial statistics and machine learning problems. Our construction generalizes the finite element-based approximations of MGFs proposed by Lindgren et al. [2011] and the Laplacian regularization in the graph-based learning literature, while at the same time retains sparsity structure that facilitates computation in practice.
- Chapter 3 is based on Harlim et al. [2020]. We apply the graph Matérn fields to solve an elliptic inverse problem in Bayesian perspective. A new kernel type approximate solver of the elliptic PDE is proposed, which can be easily discretized using graph-based techniques. A consistency result of the proposed approximate solving scheme is established as the kernel bandwidth parameter goes to zero. Several numerical examples demonstrate the effectiveness of our approach.
- Chapter 4 is based on Sanz-Alonso and Yang [2022b]. We study the effect of unlabeled data in SSL problems based on the convergence framework developed in Section 2. We show that when there are sufficiently many unlabeled data, the Bayesian approach with graph Matérn fields as the prior can achieve optimal learning performance under a manifold assumption, suggesting that the unlabeled data help. The analysis has a conceptually simple interpretation in terms of the connection between graph Matérn fields and their continuum limits and a quantitative bound on the required number of

unlabeled data is provided.

- Chapter 5 is based on Sanz-Alonso and Yang [2022c]. We adopt the framework developed in Section 4 to give some new insights about the related finite element-based methods proposed by Lindgren et al. [2011]. We show that besides the computational saving provided by the sparsity of the finite element approximations, a second layer of speed-up is possible by reducing the dimension of the problem under mild smoothness assumptions. Our theory is then complemented by simulation studies in a wide range of practical scenarios that demonstrate its relevance.
- Chapter 6 is based on García Trillos et al. [2019b]. We extend the manifold assumption that we have adopted for many parts of the thesis. We assume instead the more realistic setting where the data features no longer lie exactly on a manifold, but are noisy perturbations of points on a manifold. A local regularization scheme is proposed that provably leads to a better behaved graph-Laplacian. Several numerical experiments demonstrate improved learning performance with the regularization scheme on both synthetic and real data, suggesting its potential practical impact.

## 1.6 Notations

We close this section by introducing some notation. We denote by  $\mathcal{L}(Z)$  the law of the random variable  $Z$ . For  $a, b$  two real numbers, we denote  $a \wedge b = \min\{a, b\}$  and  $a \vee b = \max\{a, b\}$ . The symbol  $\lesssim$  will denote less than or equal to up to a universal constant. For two real sequences  $(a_n)$  and  $(b_n)$ , we denote (i)  $a_n \ll b_n$  if  $\lim_n(a_n/b_n) = 0$ ; (ii)  $a_n = O(b_n)$  if  $\limsup_n(a_n/b_n) \leq C$  for some positive constant  $C$ ; and (iii)  $a_n \asymp b_n$  if  $c_1 \leq \liminf_n(a_n/b_n) \leq \limsup_n(a_n/b_n) \leq c_2$  for some positive constants  $c_1, c_2$ . Finally we let  $\gamma$  denote the Lebesgue measure on  $\mathbb{R}$ .



## CHAPTER 2

# THE SPDE APPROACH TO MATÉRN GAUSSIAN FIELDS: GRAPH REPRESENTATIONS

### 2.1 Introduction

The stochastic partial differential equation (SPDE) approach to Gaussian fields (GFs) has been one of the key developments in spatial statistics over the last decade [Lindgren et al., 2011]. The main idea is to represent GFs as finite element solutions to SPDEs, reducing the computational cost of inference and sampling by invoking a Gaussian Markov random field (GMRF) approximation [Rue and Held, 2005]. This chapter investigates *graph representations* of stationary and nonstationary Matérn fields following the SPDE perspective, contributing to and unifying the extant theoretical, computational and methodological literature on GFs in Bayesian inverse problems, spatial statistics and graph-based machine learning. We demonstrate through transparent mathematical reasoning that, under a manifold assumption, graph representations give GMRF approximations to the Matérn model with error guarantees. In addition, we show that graph representations *generalize* the Matérn model to unstructured point clouds and graphs where existing finite element representations are not applicable.

Recall that a random function  $u(x)$ ,  $x \in \mathbb{R}^d$ , is a GF if all finite collections  $\{u(x_i)\}_{i=1}^n$  have self-consistent multivariate Gaussian distributions [Abrahamsen, 1997, Adler, 2010]. A GF can be specified using a mean function  $\mu(\cdot)$  and a covariance function  $c(\cdot, \cdot)$ , so that the mean vector and covariance matrix of the finite dimensional distributions are  $\{\mu(x_i)\} \in \mathbb{R}^n$  and  $\Sigma = \{c(x_i, x_j)\} \in \mathbb{R}^{n \times n}$ . GFs are natural models for spatial, temporal and spatio-temporal data. They have desirable analytic properties, including an explicit normalizing constant and closed formulae when conditioning on Gaussian data. However, in practice GFs have two main caveats. First, it is crucial and non-trivial to find *flexible* covariance functions with

few but interpretable parameters that can be learned from data. Second, inference of these parameters from Gaussian data of size  $n$  —or sampling the field at  $n$  locations— involves factorizing a kernel matrix  $\Sigma \in \mathbb{R}^{n \times n}$ , leading to a  $\mathcal{O}(n^3)$  computational cost and  $\mathcal{O}(n^2)$  memory cost unless further structure is assumed or imposed on the covariance model. For this reason, many recent works have investigated novel ways to deal with large datasets, some of which are reviewed in Heaton et al. [2019].

The SPDE approach tackles the big  $n$  problem by replacing the GF with a GMRF approximation. A GMRF is a discretely indexed GF  $u_n(i)$ ,  $i \in \{1, \dots, n\}$ , such that the full conditional distribution at each site  $1 \leq i \leq n$  depends only on a (small) set of neighbors  $\partial i$  to site  $i$ . This conditional independence structure is fully encoded in the precision matrix  $Q$  of the multivariate Gaussian distribution of  $u_n \in \mathbb{R}^n$ : it holds that  $Q_{ij} \neq 0$  iff  $i \in \partial j$ . Computationally, the main advantage comes from using numerical linear algebra techniques and Markov chain Monte Carlo algorithms that exploit, respectively, the sparsity of  $Q$  and the characterization of the GMRF in terms of its full conditionals. The speed-up can be dramatic, with a typical computational cost  $\mathcal{O}(n)$ ,  $\mathcal{O}(n^{3/2})$ , and  $\mathcal{O}(n^2)$  for GMRFs in time, space, and space-time in two spatial dimensions, see Rue and Held [2005]. In addition to alleviating the computational burden of GF methods, the SPDE approach also alleviates the modeling challenges by suggesting nonstationary generalizations of Matérn fields and extensions beyond Euclidean settings.

In this chapter we employ graph-based discretizations of SPDEs to represent stationary and nonstationary Matérn models. With few exceptions e.g. García Trillos and Sanz-Alonso [2018], Bertozzi et al. [2018], García Trillos et al. [2020b], Harlim et al. [2020], previous work stemming from the SPDE approach considered representations based on finite element or finite difference discretizations [Lindgren et al., 2011, Bolin and Kirchner, 2020, Bolin et al., 2020, 2018, Roininen et al., 2019, Wiens et al., 2020, Bolin, 2014]. Graph representations provide a way to *generalize* the Matérn model to discrete and unstructured point clouds, and

thus to settings of practical interest in statistics and machine learning where only similarity relationships between abstract features may be available. Moreover, in contrast to finite elements, graph representations require minimal pre-processing cost: there is no need to compute triangulations and finite element basis or to define ghost domains as in Lindgren et al. [2011]. This is an essential advantage when interpolating manifold data living in a high dimensional ambient space, particularly so when the underlying manifold or its dimension are unknown. Finally, a wide range of problems in Bayesian inversion, spatial statistics and graph-based machine learning can be formulated as latent Gaussian models, and using graph representations of Matérn fields as priors allows us to unify and contribute to the exchange of ideas across these disciplines.

A disadvantage of the graph-based approach is that error guarantees are weaker than for finite element or finite difference representations. Our belief is that this is due to the generality of the graph-based approach, and also to the underdevelopment of existing theory. Here we provide an up-to-date perspective of spectral convergence of graph Laplacians which overviews and generalizes some of the recent literature [Burago et al., 2015, García Trillos et al., 2019a,b] and we further show how these results can be used to establish the convergence of GMRFs to GFs. We view graph representations as being complementary to, rather than a replacement for, finite element representations. If the underlying domain is known and a suitable mesh can be obtained, then finite element representations would be recommended on the grounds of better error guarantees and sparsity, see Section 2.3.3.1. However, graph-based methods are more broadly applicable, and in particular generalize the Matérn model to unstructured point clouds as will be demonstrated in our numerical examples in Sections 2.5.3 and 2.5.4.

### 2.1.1 Literature Review

The ubiquity of GFs in statistics, applied mathematics and machine learning has led, unsurprisingly, to the reinvention and relabeling of many algorithms and ideas. GFs play a central role in spatial statistics [Gelfand et al., 2010, Heaton et al., 2019], especially in the subfield of geostatistics [Stein, 1999a], where they are used to interpolate data in a procedure called *kriging* and as a building block of modern hierarchical spatial models [Banerjee et al., 2014]. In machine learning, GFs are called Gaussian processes and kriging is known as Gaussian process regression [Williams and Rasmussen, 2006]. Gaussian processes are one of the main tools in Bayesian non-parametric inference [Williams and Rasmussen, 1996, van der Vaart and van Zanten, 2008a, García Trillos et al., 2020b] and are an alternative to neural networks for supervised and semi-supervised regression [MacKay, 1997, García Trillos et al., 2020b]. They are also related to, or used within, other machine learning algorithms including splines, support vector machines and Bayesian optimization [Sollich, 2002, Seeger, 2000b, Brochu et al., 2010, Frazier, 2018]. GFs are standard prior models for statistical Bayesian inverse problems [Kaipo and Somersalo, 2006, Calvetti and Somersalo, 2007, Stuart, 2010, Sanz-Alonso et al., 2019] with applications in medical imaging, remote sensing and ground prospecting [Bardsley and Kaipo, 2013, Dunlop et al., 2017, Somersalo et al., 1992, Dunlop and Stuart, 2016, García Trillos and Sanz-Alonso, 2017]. Within Bayesian inversion, GFs are also employed as surrogates for the likelihood function [Stuart and Teckentrup, 2018]. GFs have found numerous applications, allowing for *uncertainty quantification* [Sullivan, 2015] in astrophysics [Bardeen et al., 1986], biology [Taylor and Worsley, 2007, Stathopoulos et al., 2014], calibration of computer models [Kennedy and O’Hagan, 2001, Martin and Simpson, 2005], data-driven learning of partial differential equations [Raissi et al., 2018, 2017], geophysics [Isaac et al., 2015, Bui-Thanh et al., 2013], hydrology [Sanchez-Vila et al., 2006], image processing and medical imaging [Cohen et al., 1991, Somersalo et al., 1992, Roininen et al., 2014], meteorology [Bolin and Lindgren, 2011, Lindgren et al., 2011] and probabilistic

numerics [Hennig et al., 2015, Kersting and Hennig, 2016], among others.

The emphasis of this chapter is on Matérn models [Matérn, 2013] and generalizations thereof. Matérn models are widely used in spatial statistics [Stein, 1999a, Gelfand et al., 2010], machine learning [Williams and Rasmussen, 2006] and uncertainty quantification [Sullivan, 2015], with applications in various scientific fields [Guttorp and Gneiting, 2006, Cameletti et al., 2013]. The SPDE approach to construct GMRF approximations to GFs was proposed in the seminal paper [Lindgren et al., 2011] and was further popularized through the software R-INLA [Bakka et al., 2018]. GMRFs in statistics were pioneered by Besag [1974, 1975] and their computational benefits and applications are overviewed in the monograph Rue and Held [2005].

In an independent line of work, the desire to define positive semi-definite kernels using only similarity relationships between features motivated the introduction of diffusion kernels [Kondor and Lafferty, 2002], which can be interpreted as limiting cases of Matérn models. The main idea underlying the construction of diffusion kernels is to exploit that graph Laplacians [Chung, 1997, Von Luxburg, 2007] and their powers satisfy the positive semi-definiteness requirement. This observation has permeated the construction of graph-based regularizations in manifold learning and machine learning applications, as well as the design of model reduction techniques e.g. [Zhu et al., 2003, Ng et al., 2018, Li et al., 2018, Liu et al., 2014, Belkin et al., 2004, Belkin and Niyogi, 2008, 2004]. Our work aims to demonstrate that a wide family of graph-based kernels in machine learning may be interpreted, in a rigorous way, as discrete approximations of standard GF models in spatial statistics.

Large sample limits of graph Laplacians have been widely studied. Most results concern either pointwise convergence [Hein et al., 2005, Belkin and Niyogi, 2008, Giné and Koltchinskii, 2006, Hein, 2006, Singer, 2006, Ting et al., 2010] or variational and spectral convergence [Belkin and Niyogi, 2007, Singer and Wu, 2017, Tao and Shi, 2020, Burago et al., 2015, García Trillos et al., 2019b], with Calder and García Trillos [2022] reconciling both perspec-

tives to obtain improved rates. This chapter builds on and generalizes spectral convergence theory—that is, convergence of eigenvalues and eigenfunctions of the graph-based operators to those defined in the continuum—to study GMRF approximations of GFs. Unsurprisingly, we shall see that optimal transport ideas are key to linking discrete and continuum objects.

### *2.1.2 Main Contributions and Outline*

Further to providing a unified narrative of existing literature, this chapter contains some original contributions. We introduce GMRF approximations of nonstationary GFs defined on manifolds through graph representations of the corresponding SPDEs and generalize the constructions to arbitrary point clouds. Our main theoretical result, Theorem 2.4.2, covers nonstationary models and, to our knowledge, is the first to give rates of convergence of graph-based representations of GFs. We also demonstrate through numerical examples that the mathematical unity that comes from viewing graph-based methods as discretizations of continuum ones facilitates the transfer of methodology and theory across Bayesian inverse problems, spatial statistics and graph-based machine learning. In particular, we introduce nonstationary models for graph-based classification problems, which to our best knowledge has not been considered before, and empirically observe an improvement of performance that deserves further research.

This chapter is organized as follows. Section 2.2 introduces the SPDE formulation of the Matérn model and extends it to incorporate nonstationarity. Section 2.3 introduces the graph-based approach and constructs graph approximations of the Matérn fields. Section 2.4 presents the main result on the convergence of the graph Matérn model towards its continuum counterpart and discusses the ideas of the proof. Section 2.5 illustrates the application of graph Matérn models in Bayesian inverse problems, spatial statistics and graph-based machine learning. Section 2.6 discusses further research directions. Our aim is to provide a digestible narrative and for this reason we postpone all proofs and most of the

technical material to an appendix.

## 2.2 Matérn Models and the SPDE Approach

In this section we provide some background on GFs and the SPDE approach. We introduce the Matérn family in Section 2.2.1 and a nonstationary generalization in Section 2.2.2. All fields will be assumed to be centered and we focus our attention on their covariance structure.

### 2.2.1 Stationary Matérn Models

Recall that a GF in  $\mathbb{R}^m$  belongs to the Matérn class if its covariance function can be written in the form

$$c_{\sigma,\nu,\ell}(x,y) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left( \frac{|x-y|}{\ell} \right)^\nu K_\nu \left( \frac{|x-y|}{\ell} \right), \quad x,y \in \mathbb{R}^m, \quad (2.1)$$

where  $|\cdot|$  is the Euclidean distance in  $\mathbb{R}^m$ ,  $\Gamma$  denotes the Gamma function and  $K_\nu$  denotes the modified Bessel function of the second kind. The parameters  $\sigma, \nu$  and  $\ell$  control, respectively, the marginal variance (magnitude), regularity and correlation length scale of the field. While being defined in terms of three interpretable parameters, the modeling flexibility afforded by the Matérn covariance (2.1) is limited by its *stationarity* (the value of the covariance function depends only on the difference between its arguments) and *isotropy* (it depends only on their Euclidean distance).

An important characterization by Whittle [1954, 1963] is that Matérn fields can be defined as the solution to certain fractional order stochastic partial differential equation (SPDE). Precisely, setting  $\tau := \ell^{-2}$ ,  $s := \nu + \frac{m}{2}$ , a Gaussian field with covariance function (2.1) is the unique stationary solution to the SPDE

$$(\tau I - \Delta)^{\frac{s}{2}} u(x) = \mathcal{W}(x), \quad x \in \mathbb{R}^m, \quad (2.2)$$

where the marginal variance of  $u$  is

$$\sigma^2 = \frac{\Gamma(s - \frac{m}{2})}{(4\pi)^{\frac{m}{2}} \Gamma(s) \tau^{s - \frac{m}{2}}}. \quad (2.3)$$

Throughout this chapter, fractional power operators such as  $(\tau I - \Delta)^{\frac{s}{2}}$  will be defined spectrally [Lischke et al., 2020] and  $\mathcal{W}$  denotes spatial Gaussian white noise with unit variance.

As discussed in Lindgren et al. [2011], the SPDE formulation of Matérn GFs has several advantages. First, it allows to approximate the solution to (2.2) by a GMRF and thereby to reduce the computational cost of inference and sampling [Rue and Held, 2005, Simpson et al., 2012]. Second, it suggests natural nonstationary and anisotropic generalizations of the Matérn model by letting  $\tau$  depend on the spatial variable [Roininen et al., 2019] or by replacing the Laplacian with an elliptic operator with spatially varying coefficients [Fuglstad et al., 2015a,b]. Third, it allows to *define* Matérn models in manifolds and in bounded spatial, temporal and spatio-temporal domains by using modifications of the SPDE (2.2), possibly supplemented with appropriate boundary conditions [Khristenko et al., 2019]. In order to gain theoretical understanding, in subsequent sections we will work under a manifold assumption and analyze the convergence of graph representations of Matérn fields defined on manifolds. This setting is motivated by manifold learning theory [Belkin et al., 2004, García Trillos et al., 2019b] and will allow us to build on the rich literature on GFs on manifolds [Adler, 2010].

In more mathematical terms, the SPDE characterization shifts attention from the covariance function (or spectral density) description of Gaussian measures to the covariance (or precision) operator description [Bogachev, 1998]: keeping only the  $\tau$  term in the marginal variance given by equation (2.3), we see that the law of the field  $u(x)$  defined by equation (2.2) is —up to a scaling factor independent of  $\tau$  that we drop in what follows— the Gaussian



measure

$$\mathcal{N}(0, \mathcal{C}), \quad \mathcal{C} = \tau^{s-\frac{m}{2}} (\tau I - \Delta)^{-s}, \quad (2.4)$$

where the factor  $\tau^{s-\frac{m}{2}}$  can be interpreted as a normalizing constant. This observation motivates our definition of the nonstationary Matérn field in Section 2.2.2, which facilitates the theory.

### 2.2.2 Nonstationary Matérn Models

In this subsection we introduce a family of nonstationary Matérn fields by modifying the SPDE (2.2). We consider a manifold setting which does not hinder the understanding of the modeling and will later allow us to frame the analysis in a concrete setting of applied significance. To that end, we let  $\mathcal{M}$  be an  $m$ -dimensional smooth, connected, compact Riemannian manifold without boundary that is embedded in  $\mathbb{R}^d$ . We will let  $\tau$  depend on the spatial variable and replace the Laplacian by an elliptic operator  $\nabla \cdot (\kappa(x)\nabla)$ , where differentiation is defined on  $\mathcal{M}$ . Formally, we consider the SPDE

$$\left[ \tau(x)I - \nabla \cdot (\kappa(x)\nabla) \right]^{\frac{s}{2}} u(x) = \mathcal{W}(x), \quad x \in \mathcal{M}, \quad (2.5)$$

where  $\mathcal{W}$  is a spatial Gaussian white noise with unit variance on  $\mathcal{M}$ . The additional  $\kappa$  acts as a change of coordinate  $\tilde{x} = \sqrt{\kappa(x)}x$  and introduces a factor of  $\kappa(x)^{-\frac{m}{2}}$  for the marginal variance, whence the field  $u(x)$  in equation (2.5) has marginal variance proportional to  $\tau(x)^{\frac{m}{2}-s}\kappa(x)^{-\frac{m}{2}}$  at each location. If  $\mathcal{M} = \mathbb{R}^d$ , the solution  $u$  to (2.5) defines a nonstationary field, and in analogy we will use the term nonstationary for fields defined by the SPDE (2.5), or approximations thereof, in manifold and more abstract settings. In such settings, stationarity or “shift-invariance” is not well-defined without introducing an algebraic action, and nonstationarity should be understood as nonhomogeneity.

Following again the covariance operator viewpoint, we formally consider the Gaussian measure  $\mathcal{N}(0, [\mathcal{L}^{\tau, \kappa}]^{-s})$ , where  $\mathcal{L}^{\tau, \kappa} := \tau I - \nabla \cdot (\kappa \nabla)$ , with a proper normalization to be made precise below. Assuming sufficient regularity,  $\mathcal{L}^{\tau, \kappa}$  is self-adjoint with respect to the  $L^2(\mathcal{M})$  inner product and admits a spectral decomposition. Therefore we shall *define* our nonstationary Matérn field through the following Karhunen-Loève expansion,

$$u(x) := \tau(x)^{\frac{s}{2} - \frac{m}{4}} \kappa(x)^{\frac{m}{4}} \sum_{i=1}^{\infty} \left[ \lambda^{(i)} \right]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)}(x), \quad (2.6)$$

where  $\xi^{(i)} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$  and  $\{(\lambda^{(i)}, \psi^{(i)})\}_{i=1}^{\infty}$  are the eigenpairs of  $\mathcal{L}^{\tau, \kappa} := \tau I - \nabla \cdot (\kappa \nabla)$ . The factor  $\tau(x)^{\frac{s}{2} - \frac{m}{4}} \kappa(x)^{\frac{m}{4}}$  serves as a normalizing constant for the marginal variance at each point. For the theory outlined in Section 2.4 we will assume that  $\tau$  is Lipschitz,  $\kappa$  is continuously differentiable and both are bounded from below by a positive constant, whence Weyl's law [Canzani, 2013, Theorem 72] implies that  $\lambda^{(i)} \asymp i^{\frac{2}{m}}$ . Therefore by setting  $s > \frac{m}{2}$ , we have  $\mathbb{E}\|u\|_{L^2(\mathcal{M})}^2 < \infty$  and the series (2.6) converges in  $L^2(\mathcal{M})$  almost surely. The idea of viewing the functions  $\tau$  or  $\kappa$  as hyperparameters and learning them from data has been investigated in Fuglstad et al. [2015a], Roininen et al. [2019], Monterrubio-Gómez et al. [2020], Fuglstad et al. [2015b], Wiens et al. [2020] and has motivated the need to penalize the complexity of priors [Fuglstad et al., 2019]. We note that other approaches to introduce nonstationarity that do not stem directly from the SPDE formulation have been considered in the literature (e.g. Anderes and Stein [2008], Gramacy and Lee [2008], Kim et al. [2005], Montagna and Tokdar [2016], Sampson et al. [2001]).

**Remark 2.2.1.** *The normalizing factors  $\tau(x)^{\frac{s}{2} - \frac{m}{4}} \kappa(x)^{\frac{m}{4}}$  are crucial for hierarchical models in that they balance the marginal variances at different locations. To gain more intuition on the powers, consider the case where both  $\tau$  and  $\kappa$  are constant. Weyl's law then implies that the eigenvalues of  $\mathcal{L}^{\tau, \kappa}$  satisfy*

$$\lambda^{(i)} \asymp \tau + C \kappa i^{\frac{2}{m}},$$

and therefore

$$\begin{aligned}
\mathbb{E} \left\| \sum_{i=1}^{\infty} [\lambda^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)} \right\|_{L^2(\mathcal{M})}^2 &\asymp \sum_{i=1}^{\infty} \left[ \tau + C\kappa i^{\frac{2}{m}} \right]^{-s} \\
&\asymp \sum_{i:\tau \gtrsim \kappa i^{2/m}} \tau^{-s} + \sum_{i:\tau \lesssim \kappa i^{2/m}} \kappa^{-s} i^{-\frac{2s}{m}} \\
&\asymp \tau^{-s} \left( \frac{\tau}{\kappa} \right)^{\frac{m}{2}} + \kappa^{-s} \int_{\left(\frac{\tau}{\kappa}\right)^{\frac{m}{2}}}^{\infty} x^{-\frac{2s}{m}} \\
&\asymp \tau^{\frac{m}{2}-s} \kappa^{-\frac{m}{2}}.
\end{aligned}$$

We thus see that the normalizing factor above balances the expected norm.  $\square$

**Remark 2.2.2.** Both parameters  $\tau$  and  $\kappa$  control the local length scales of the sample paths.

To see this, note that when both  $\tau$  and  $\kappa$  are constant (2.6) simplifies to

$$u = \left( \frac{\tau}{\kappa} \right)^{\frac{s}{2} - \frac{m}{4}} \sum_{i=1}^{\infty} \left[ \frac{\tau}{\kappa} + \lambda^{(i)} \right]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)},$$

where  $\{(\lambda^{(i)}, \psi^{(i)})\}_{i=1}^{\infty}$  are eigenpairs of  $-\Delta$ . Therefore  $\frac{\tau}{\kappa}$  acts as a threshold on the essential frequencies of the samples, where those frequencies with corresponding eigenvalue on the same order of  $\frac{\tau}{\kappa}$  have effective contributions. Hence a large  $\tau$  (or a smaller  $\kappa$ ) incorporates higher frequencies and gives sample paths with small length scale. Their opposite role in controlling the local length scale can be seen in Figure 2.1, which represent two random draws from Gaussian fields defined on the unit circle with different choices of  $\kappa$  and  $\tau$ .  $\square$

### 2.3 GMRF Approximation with Graph Representations of SPDEs

In this section we study GMRF approximations of the Matérn models introduced in Section 2.2. Since the work of Lindgren et al. [2011], a burgeoning literature has been devoted to linking GFs and GMRFs, doing the modeling with the former and computations with the




Figure 2.1: Random draws from nonstationary GFs on the unit circle; (a).  $\kappa = 0.01; \tau = \exp(\cos(x))$ ; (b).  $\kappa = \exp(\cos(x)), \tau = 100$ .

latter [Bakka et al., 2018]. The main idea of Lindgren et al. [2011] is to introduce a stochastic weak formulation of the SPDE (2.2):

$$\langle (\tau I - \Delta)^{\frac{s}{2}} u, \phi_i \rangle_{L^2} \stackrel{\mathcal{D}}{=} \langle \mathcal{W}, \phi_i \rangle_{L^2}, \quad i = 1, \dots, M,$$

where  $\{\phi_i\}_{i=1}^M$  is a set of test functions and  $\stackrel{\mathcal{D}}{=}$  denotes equal in distribution. Then one constructs a finite element (FEM) *representation* of the solution

$$u(x) = \sum_{k=1}^n w_k \varphi_k(x),$$

where  $n$  is the number of vertices in the triangulation,  $\{\varphi_k\}$  are interpolating piecewise linear hat functions and  $\{w_k\}$  are Gaussian distributed weights. Importantly, these finite dimensional representations allow to obtain a GMRF precision matrix with computational cost  $\mathcal{O}(n)$ . The convergence of the FEM representation to the GF has been studied in Lindgren et al. [2011] and in more generality in Bolin and Kirchner [2020], Bolin et al. [2020, 2018].

The FEM representation requires triangulation of the domain, possibly adding artificial nodes to obtain a suitable mesh, and in practice it is rarely implementable in dimension higher than 3. However for many applications e.g. in machine learning, interest lies in

interpolating or classifying input data in high dimensional ambient space with moderate but unknown intrinsic dimension, making FEM representations of GFs impractical. Graph Laplacians, discussed next, provide a canonical way to construct GMRF approximations in the given point cloud.

### 2.3.1 Graph Matérn Models

Let  $\mathbb{X} = \{x_1, \dots, x_n\}$  be a given point cloud, over which we put a graph structure by considering a symmetric weight matrix  $W \in \mathbb{R}^{n \times n}$  whose entries  $W_{ij} \geq 0$  prescribe the closeness between points. In applications including classification and regression, each  $x_i$  will represent either a feature or an auxiliary point used to improve the accuracy of the GMRF approximations described in this subsection. The graph structure encodes the geometry of the point cloud and can be exploited through the graph Laplacian.

Several definitions of graph Laplacians co-exist in the literature. Defining  $D := \text{diag}(d_1, \dots, d_n)$  the degree matrix with  $d_i := \sum_{j=1}^n W_{ij}$ , three popular choices are unnormalized  $\Delta_n^{\text{un}} := D - W$ , symmetric  $\Delta_n^{\text{sym}} := D^{-1/2} \Delta_n^{\text{un}} D^{-1/2}$  and random-walk  $\Delta_n^{\text{rw}} := D^{-1} \Delta_n^{\text{un}}$  graph Laplacians, see Von Luxburg [2007]. To streamline the presentation, we use  $\Delta_n \in \mathbb{R}^{n \times n}$  as placeholder for a graph Laplacian with  $n$  data points; its choice will be made explicit whenever it is relevant to the problem at hand.

To gain some intuition, let us consider the unnormalized graph Laplacian, whose positive semi-definiteness is verified by the relation

$$u_n^T (D - W) u_n = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n W_{ij} |u_n(i) - u_n(j)|^2 \geq 0. \quad (2.7)$$

Here  $u_n = [u_n(1), \dots, u_n(n)]^T \in \mathbb{R}^n$  is an arbitrary vector in  $\mathbb{R}^n$ , interpreted as a function on  $\mathbb{X}$  with the identification  $u_n(x_i) \equiv u_n(i)$ . Note that  $\Delta_n^{\text{un}} = D - W$  annihilates constant vectors (in agreement with the intuition that the Laplacian annihilates constant functions)

and 0 is always an eigenvalue. For a fully connected graph, one can see that the eigenvalue 0 has multiplicity 1, with the constant vectors as its only eigenspace. If we consider  $\mathcal{N}(0, \Delta_n^{-1})$  (with  $\Delta_n^{-1}$  representing the Moore-Penrose inverse) as a degenerate Gaussian distribution in  $\mathbb{R}^n$  with support on the orthogonal complement of the constant vectors, then (2.7) is the negative log-density of this distribution (up to an additive constant), which suggests that functions that take similar values on close nodes are favored, with closeness quantified by the weight matrix  $W$ . Moreover, it can be shown that the second eigenvector  $\psi_n^{(2)}$  of  $\Delta_n$  solves a relaxed graph cut problem [Von Luxburg, 2007], so that  $\psi_n^{(2)}$  encodes crucial information about partition of the points  $x_i$ 's. Hence  $\mathcal{N}(0, \Delta_n^{-1})$  naturally serves as a prior for clustering and classification [Bertozzi et al., 2018]. Various choices of the weight matrix have been considered in the literature, including  $\varepsilon$ -graphs and  $k$ -NN graphs, which set  $W_{ij}$  to be zero if  $d(x_i, x_j) > \varepsilon$  and if  $x_i$  is not among the  $k$ -nearest neighbors of  $x_j$  (or vice versa) respectively for some distance function  $d$ . Both of them introduce sparsity in the weight matrix, which is inherited by the graph Laplacian. Under such circumstances, the graph Laplacian can be viewed as a sparse precision matrix, which gives rise to a GMRF.

*It is important to note that the preceding discussion makes no assumption on the points  $x_i$  or how their closeness is defined.* Therefore, the graph-based viewpoint allows to *generalize* the Matérn model to unstructured point clouds, and thus to settings of practical interest in statistics and machine learning where only similarity relationships between abstract features may be available. For instance, the points may represent books and their closeness may be based on a reader's perception of similarity between them. However, an important example in which we will frame our theoretical investigations arises from making a *manifold assumption*.

**Assumption 2.3.1** (Manifold Assumption). *The points  $x_i$  are independently sampled from the uniform distribution  $\gamma$  on an  $m$ -dimensional smooth, connected, compact manifold  $\mathcal{M}$  without boundary that is embedded in Euclidean space  $\mathbb{R}^d$ , with bounded sectional curvature and Riemannian metric inherited from  $\mathbb{R}^d$ . Assume further that  $\mathcal{M}$  is normalized so that*

$vol(\mathcal{M})=1$ .

To emphasize the stronger structure imposed by the manifold assumption we denote the point cloud by  $\mathbb{X} \equiv \mathcal{M}_n = \{x_1, \dots, x_n\} \subset \mathcal{M}$ . For many applications, the manifold assumption is an idealization of the fact that the point cloud has low dimensional structure despite living in a high dimensional ambient space, e.g., the MNIST dataset that we will study in Section 2.5.4. From a theoretical viewpoint, Assumption 2.3.1 allows us to establish a precise link between graph Laplacians and their continuum counterparts, as we now describe heuristically. Define the weight matrix  $W$  on  $\mathcal{M}_n$  by

$$W_{ij} := \frac{2(m+2)}{n\nu_m h_n^{m+2}} \mathbf{1}\{|x_i - x_j| < h_n\}, \quad (2.8)$$

where  $|\cdot|$  is the Euclidean distance in  $\mathbb{R}^d$ ,  $h_n$  is the graph connectivity and  $\nu_m$  is the volume of the  $m$ -dimensional unit ball. Then the unnormalized graph Laplacian  $\Delta_n^{\text{un}}$  is a discrete approximation of the Laplace-Beltrami operator  $-\Delta$  on  $\mathcal{M}$ .

Indeed, for a smooth function  $f : \mathbb{R}^m \rightarrow \mathbb{R}$  we have by Taylor expansion of  $f$  around  $X$

$$\begin{aligned} \int_{B_{h_n}(X)} [f(Y) - f(X)] dY &\approx \int_{B_{h_n}(X)} \nabla f(X)^T (Y - X) dY \\ &\quad + \frac{1}{2} \int_{B_{h_n}(X)} (Y - X)^T \nabla^2 f(X) (Y - X) dY, \end{aligned}$$

where  $B_{h_n}(X)$  is the Euclidean ball centered at  $X$  with radius  $h_n$ . By symmetry, the first integral is zero and the second integral reduces to (after a change of variable  $Z = Y - X$ )

$$\frac{1}{2} \sum_{i=1}^m \frac{\partial^2 f}{\partial X_i^2}(X) \int_{B_{h_n}(0)} Z_i^2 dZ = \frac{\nu_m h_n^{m+2}}{2(m+2)} \Delta f(X),$$

where  $X_i$  and  $Z_i$  represent the  $i$ -th coordinates of  $X$  and  $Z$ . This gives

$$-\Delta f(x_i) \approx \frac{2(m+2)}{\nu_m h_n^{m+2}} \int_{B_{h_n}(x_i)} [f(x_i) - f(Y)] dY \approx \sum_{j=1}^n W_{ij} [f(x_i) - f(x_j)], \quad (2.9)$$

which is exactly the way  $\Delta_n$  is defined. Since  $\mathcal{M}$  is locally homeomorphic to  $\mathbb{R}^m$  and the geodesic distance between any two points is well approximated by the Euclidean distance, the heuristic argument above can be formalized to show *point-wise* convergence of  $\Delta_n$  to  $-\Delta$  in the manifold case. A rigorous result on *spectral* convergence will be given in more generality in Section 2.4.

The previous discussion suggests to introduce the following discrete analog to the Gaussian measure (2.4)

$$\mathcal{N}(0, \mathcal{C}_n), \quad \mathcal{C}_n = \tau^{s-\frac{m}{2}} (\tau I_n + \Delta_n)^{-s},$$

whose samples admit a Karhunen-Loève expansion

$$u_n = \tau^{\frac{s}{2}-\frac{m}{4}} \sum_{i=1}^n \left[ \tau + \lambda_n^{(i)} \right]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)}, \quad (2.10)$$

where  $\{\xi^{(i)}\}_{i=1}^n$  are independent standard normal random variables and  $\{(\lambda_n^{(i)}, \psi_n^{(i)})\}_{i=1}^n$  are eigenpairs of  $\Delta_n$ . This will be our definition of the stationary graph Matérn field.

**Remark 2.3.2.** *We note once again that the model (2.10) can be used in wide generality: it only presupposes that the practitioner is given a weight matrix associated with an abstract point cloud, and it only requires to specify a graph-Laplacian. We will show that (2.10) generalizes the stationary Matérn model in the sense that if the point cloud is sampled from a manifold, the weights are defined through an appropriate  $\varepsilon$ -graph, and an unnormalized graph-Laplacian is used, then the graph-based model approximates the Matérn model on the manifold. Similar convergence results could be established with  $k$ -nearest neighbor graphs and*



other choices of graph Laplacian. Our numerical examples in Section 2.5 will illustrate the application of graph-based Matérn models both in manifold and abstract settings, and using a variety of graph Laplacians.  $\square$

**Remark 2.3.3.** The above construction can be adapted when the points  $x_i$  are distributed according to a Lipschitz density  $q$  that is bounded below and above by positive constants. In this case, (2.9) should take the form

$$-\Delta f(x_i) \approx \sum_{j=1}^n W_{ij}[f(x_i) - f(x_j)]q(x_j)^{-1} \approx \frac{1}{2} \sum_{j=1}^n W_{ij}[f(x_i) - f(x_j)][q(x_i)^{-1} + q(x_j)^{-1}],$$

where the last step follows from the Lipschitzness of  $q$  and the fact that  $q^{-1}$  is bounded away from 0 and is needed to ensure symmetry of the new weights. Setting  $f = q$  in (2.9) we have

$$q(x_i) \approx \frac{1}{\nu_m h_n^m} \int_{B_{h_n}(x_i)} q(y) dy - h_n^2 \Delta q(x_i) \approx \frac{1}{n \nu_m h_n^m} \sum_{j=1}^n \mathbf{1}\{|x_i - x_j| < h_n\} := q_{h_n}(x_i)$$

where we have dropped  $h_n^2 \Delta q$  since it is of lower order. Hence the new weights should be adjusted as

$$W_{ij} = \frac{m+2}{n \nu_m h_n^{m+2}} \mathbf{1}\{|x_i - x_j| < h_n\} [q_{h_n}(x_i)^{-1} + q_{h_n}(x_j)^{-1}].$$

$\square$

### 2.3.2 Nonstationary Graph Matérn Models

Now we are ready to construct nonstationary graph Matérn fields that approximate the nonstationary Matérn field in Section 2.2.2. In analogy with the previous subsection, the crucial step is to obtain a graph discretization of the operator  $\mathcal{L}^{\tau, \kappa} = \tau I - \nabla \cdot (\kappa \nabla)$  with

spatially varying  $\tau$  and  $\kappa$ . Notice that we have

$$\nabla \cdot (\kappa \nabla f) = \sqrt{\kappa} [\Delta(\sqrt{\kappa} f) - f \Delta \sqrt{\kappa}].$$

Applying (2.9) to  $\Delta(\sqrt{\kappa} f)$  and  $f \Delta \sqrt{\kappa}$  gives

$$-\nabla \cdot (\kappa \nabla f) \approx \int_{B_{h_n}(x)} \sqrt{\kappa(x)\kappa(y)} [f(x) - f(y)] \approx \sum_{i=1}^n W_{ij} \sqrt{\kappa(x_i)\kappa(x_j)} [f(x_i) - f(x_j)].$$

Hence  $-\nabla \cdot (\kappa \nabla \cdot)$  can be approximated by  $\Delta_n^\kappa = \tilde{D} - \tilde{W}$ , where

$$\tilde{W}_{ij} = W_{ij} \sqrt{\kappa(x_i)\kappa(x_j)} = \frac{2(m+2)}{n\nu_m h_n^{m+2}} \mathbf{1}\{|x_i - x_j| < h_n\} \sqrt{\kappa(x_i)\kappa(x_j)}, \quad (2.11)$$

$$\tilde{D}_{ii} = \sum_{j=1}^n \tilde{W}_{ij}. \quad (2.12)$$

Denoting  $\tau_n = \text{diag}(\tau(x_1), \dots, \tau(x_n))$  and  $\kappa_n = \text{diag}(\kappa(x_1), \dots, \kappa(x_n))$ , we define—similarly as in Section 2.2.2—the nonstationary graph Matérn field through the Karhunen-Loève expansion

$$u_n := \tau_n^{\frac{s}{2} - \frac{m}{4}} \kappa_n^{\frac{m}{4}} \sum_{i=1}^n [\lambda_n^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)}, \quad (2.13)$$

where  $\{\xi^{(i)}\}_{i=1}^n$  are independent standard normal random variables and  $\{(\lambda_n^{(i)}, \psi_n^{(i)})\}_{i=1}^n$  are eigenpairs of  $L_n^{\tau, \kappa} := \tau_n + \Delta_n^\kappa$ . Equation (2.13) is a natural finite dimensional approximation of (2.6) and one should expect that spectral convergence of  $\Delta_n$  towards  $-\Delta$  will translate into convergence of (2.13) towards (2.6) in the large  $n$  limit. This will be rigorously shown in Section 2.4.

In the covariance operator view,  $u_n$  follows a Gaussian distribution  $\mathcal{N}(0, C_n^{\tau, \kappa})$  with

$$C_n^{\tau, \kappa} := \tau_n^{\frac{s}{2} - \frac{m}{4}} \kappa_n^{\frac{m}{4}} [L_n^{\tau, \kappa}]^{-s} \kappa_n^{\frac{m}{4}} \tau_n^{\frac{s}{2} - \frac{m}{4}}. \quad (2.14)$$

Therefore samples can be generated by solving

$$[L_n^{\tau, \kappa}]^{\frac{s}{2}} u_n = \xi_n, \quad \xi_n \sim \mathcal{N}(0, I_n) \quad (2.15)$$

and then multiplying with the diagonal matrix  $\tau_n^{s/2-m/4} \kappa_n^{m/4}$ . For  $0 < s < 2$ , (2.15) can be solved with a sparse approximation as in Harizanov et al. [2018], Bolin and Kirchner [2020] and for  $s \geq 2$  an iterative scheme can be employed. We remark that (2.15) can be solved exactly by performing a spectral decomposition of  $L_n^{\tau, \kappa}$ , which is computationally more expensive and not recommended for large  $n$ 's.

**Remark 2.3.4.** *As discussed in Remark 2.2.2, both  $\tau_n$  and  $\kappa_n$  control the local length scale. For many applications e.g. in machine learning, we shall focus on the modeling choice with  $\tau_n$  only, because the operator  $\nabla \cdot (\kappa \nabla)$  is less motivated for general  $x_i$ 's that do not come from a manifold. In such cases, one can construct a nonstationary Matérn field similarly as above, by using a graph Laplacian built with the  $x_i$ 's, e.g. with a  $k$ -NN graph. Indeed, the only key step is to normalize properly the marginal variances, which are largely determined by the growth of the spectrum as in Remark 2.2.1. Hence one can find an integer  $m$  so that the first several  $\lambda_n^{(i)}$ 's grow roughly as  $i^{\frac{2}{m}}$  and use  $m$  as an effective dimension of the problem for normalization. Moreover, both the  $k$ -NN and  $\varepsilon$ -graphs result in sparsity in  $\Delta_n$ , and numerical linear algebra techniques for sparse systems can be employed to attain speed-up.  $\square$*

### 2.3.3 A Simulation Study

In this subsection we perform a simulation study on the unit sphere to demonstrate the graph approximation of Matérn fields and its sparsity. Let  $\mathcal{M}$  be the two-dimensional unit sphere embedded in  $\mathbb{R}^3$  and formally consider the Matérn model specified by the SPDE

$$(I - \Delta)^{-\frac{s}{2}} u(x) = \mathcal{W}(x),$$

where  $\Delta$  is the Laplace-Beltrami operator on  $\mathcal{M}$  and  $\mathcal{W}$  is spatial white noise with unit variance. More precisely, we consider the Matérn field defined by the Karhunen-Loève expansion

$$u(x) = \sum_{i=1}^{\infty} \left[1 + \lambda^{(i)}\right]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)}, \quad \xi^{(i)} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1),$$

where  $\{(\lambda^{(i)}, \psi^{(i)})\}_{i=1}^{\infty}$  are the eigenpairs of  $-\Delta$ . It is known that the eigenvalues are  $\ell(\ell-1)$  with multiplicity  $2\ell-1$  for  $\ell = 1, 2, \dots$ , and the eigenfunctions are the spherical harmonics. The covariance function associated with this field is

$$c(x, y) = \sum_{i=1}^{\infty} \left[1 + \lambda^{(i)}\right]^{-s} \psi^{(i)}(x) \psi^{(i)}(y). \quad (2.16)$$

We shall investigate the approximation of this covariance function by the covariance function of a graph Matérn field. In Section 2.3.3.1 we consider the case where only unstructured samples from the sphere are available to demonstrate the generality of the graph-based method and in Section 2.3.3.2 we restrict ourselves to triangulations of the sphere for comparison with the FEM-based approximation.

### 2.3.3.1 Unstructured Grids

In this subsection we consider “pseudo-unstructured” point clouds generated as follows. The idea is to parametrize points on the sphere in polar coordinates  $(\theta, \phi) \in [0, \pi] \times [0, 2\pi]$ . so that the uniform distribution on the sphere can be generated with the formula

$$\theta = \arccos(1 - 2U), \quad \phi = 2\pi V,$$

where  $U, V$  are independent  $\text{unif}(0, 1)$  random variables. Now instead of generating  $n$  i.i.d. pairs of  $(U, V)$ , we will partition the domain  $[0, \pi] \times [0, 2\pi]$  uniformly into subgrids of size  $\frac{\pi}{M}$  by  $\frac{2\pi}{M}$  for an integer  $M$  and then pick one point from each subgrid randomly and uniformly.

The reason is that due to the rotational symmetry of the spherical harmonics, the computed graph eigenfunctions may be out-of-phase versions of the true eigenfunctions and hence we need some structure from the point cloud in order to make them well aligned. Therefore the generated point cloud is only close to being unstructured.

Let  $\{x_i\}_{i=1}^{n=M^2}$  denote the generated point cloud. We construct an  $\varepsilon$ -graph over the  $x_i$ 's by setting the weights as in (2.8). Precisely, we define

$$W_{ij} = \frac{2(m+2)\text{vol}(\mathcal{M})}{n\nu_m h^{m+2}} \mathbf{1}\{|x_i - x_j| < h_n\} = \frac{32}{nh_n^4} \mathbf{1}\{|x_i - x_j| < h_n\}, \quad (2.17)$$

where  $m = 2$  in this case and the additional factor  $\text{vol}(\mathcal{M})$  is needed to account for the fact that  $\text{vol}(\mathcal{M}) \neq 1$ . Let  $\Delta_n = D - W$  be the unnormalized graph Laplacian. Then (2.16) is approximated by

$$c_n(x_j, x_k) = \sum_{i=1}^n \left[1 + \lambda_n^{(i)}\right]^{-s} \psi_n^{(i)}(x_j) \psi_n^{(i)}(x_k), \quad (2.18)$$

where the  $\psi_n^{(i)}$ 's are suitably normalized eigenfunctions of  $\Delta_n$ .

To demonstrate the approximation, we pick  $M$  points  $x_1, \dots, x_M$  with similar longitude ranging from the north pole to the south pole. That way, the distance between  $x_1$  and  $x_k$  is increasing, with  $x_M$  being the farthest and we then compare the values of  $c_n(x_1, x_k)$  and  $c(x_1, x_k)$  for  $k = 1, \dots, M$ . We shall focus on the  $s = 2$  case and set the connectivity as  $h_n = 1.5 \times n^{-1/4}$ , which is motivated by (2.23) below. Figure 2.2 shows three simulations with  $M = 50, 60, 70$  respectively in each row. The three plots in the first column show that we get reasonable approximations except for the first entry. The reason lies in the poor spectral approximation after certain threshold (Theorems 2.4.6 and 2.4.7) as demonstrated in the second column. The spectrum of  $\Delta_n$  becomes almost flat after some threshold and hence the tails of (2.18) have a nonnegligible contribution that worsens the approximation. Such effect is most prominent for  $c_n(x_1, x_1)$ , since for  $j \neq k$  the vectors  $\{\psi_n^{(i)}(x_j)\}_{i=1}^n$

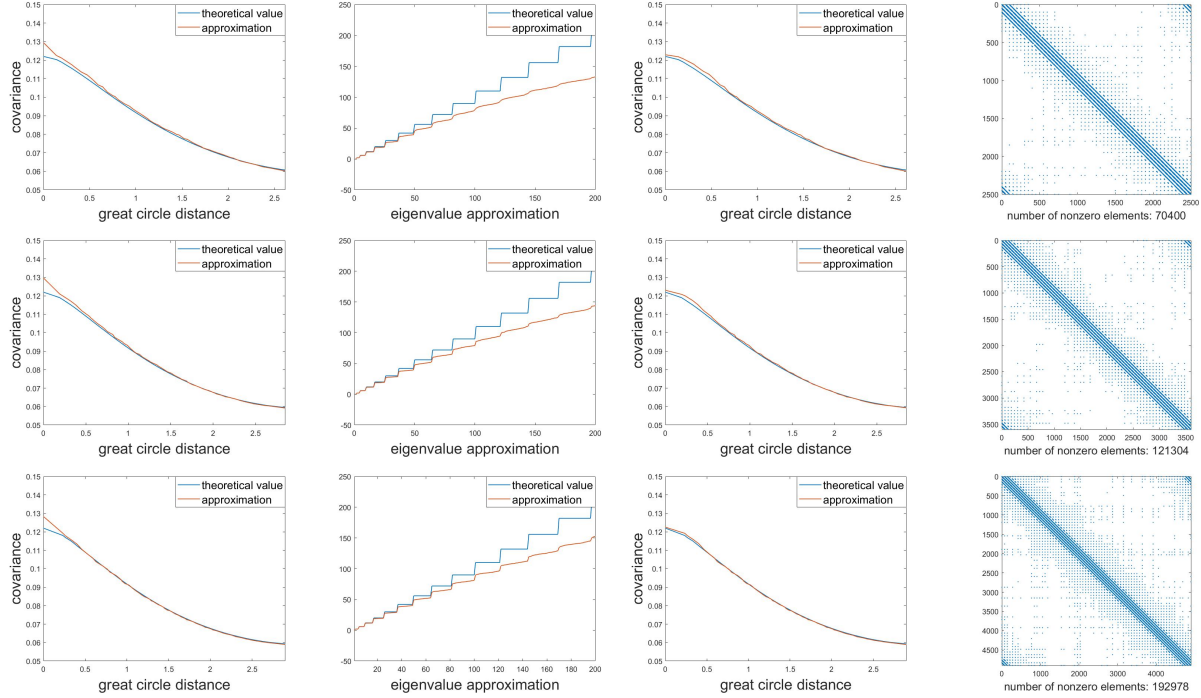


Figure 2.2: The three rows represent simulations for  $n = 2500$ ,  $3600$ ,  $4900$  respectively. Within each row, the plots represent (1) Covariances: theoretical (blue) vs. approximation (red). (2) Spectra of  $-\Delta$  (blue) and  $\Delta_n$  (red). (3). Covariances (truncated): theoretical (blue) vs. approximation (red). (4). Sparsity pattern of  $\Delta_n$ .

and  $\{\psi_n^{(i)}(x_k)\}_{i=1}^n$  are less correlated (and in fact orthogonal because they are the rows of the eigenvector matrix from singular value decomposition), so that cancellations reduce the contribution of their tails in (2.18) even if the spectrum gets flat. In the third column of Figure 2.2 we compute the truncated version of (2.18) by keeping only the first  $\sqrt{n}$  terms, and we see that  $c_n(x_1, x_1)$  is improved substantially with the rest being almost the same as before. The truncation level  $\sqrt{n}$  is motivated by Theorem 2.4.6 so that the error  $h_n \sqrt{\lambda^{(\sqrt{n})}} = O(1)$  in this case. Finally the last column of Figure 2.2 shows the sparsity patterns of the  $\Delta_n$ 's, which have decreasing percentages of nonzero entries 1.13%, 0.94% and 0.8% as  $n$  increases. The sparsity indeed leads to computational speed-up as for instance the Matlab built-in function `chol` takes in our machine 0.1713s to factorize  $(I_n + \Delta_n)^S$  but 0.6082s to factorize a random positive semi-definite  $n$  by  $n$  dense matrix when  $n = 4900$ .

### 2.3.3.2 Structured Grids

The unstructured point clouds in the previous subsection do not necessarily form triangulations of the sphere and FEM-based approximation would require additional nodes. Therefore to compare the two approximations, we assume instead to be given a triangulation as shown in the first column of Figure 2.3. The three triangulations are generated from the R-INLA function `inla.mesh.create(globe=k)` with  $k = 17, 19, 22$ , which consist of  $n = 2892, 3612, 4842$  points respectively. Graph-based approximations are constructed as in (2.17) with  $h_n = 2 \times n^{-1/4}$  and FEM-based approximations are computed using the R-INLA package. The second column of Figure 2.3 shows that the FEM-based approximations are almost indistinguishable from the truth and the truncated graph-based approximations (as in Section 2.3.3.1 with truncation level  $\sqrt{n}$ ) are also reasonably accurate. The last two columns of Figures 2.3 demonstrate the corresponding sparsity patterns of the precision matrices, i.e. the matrices  $(I_n + \Delta_n)^s$  for the graph-based case. The comparisons suggest that FEM-based approximations outperform the graph-based ones, which is not unexpected since the graph-based approach does not fully exploit the structure of the triangulation. Therefore the graph-based approximation is especially valuable when a triangulation of the domain is unavailable and difficult to obtain, as will be demonstrated in our numerical examples in Section 2.5.

## 2.4 Convergence of Graph Representations of Matérn Models

In this section we study the convergence of graph representations of GFs under the manifold Assumption 2.3.1. The analysis will generalize existing literature to cover the nonstationary models introduced in Section 2.3.2, obtaining new rates of convergence.

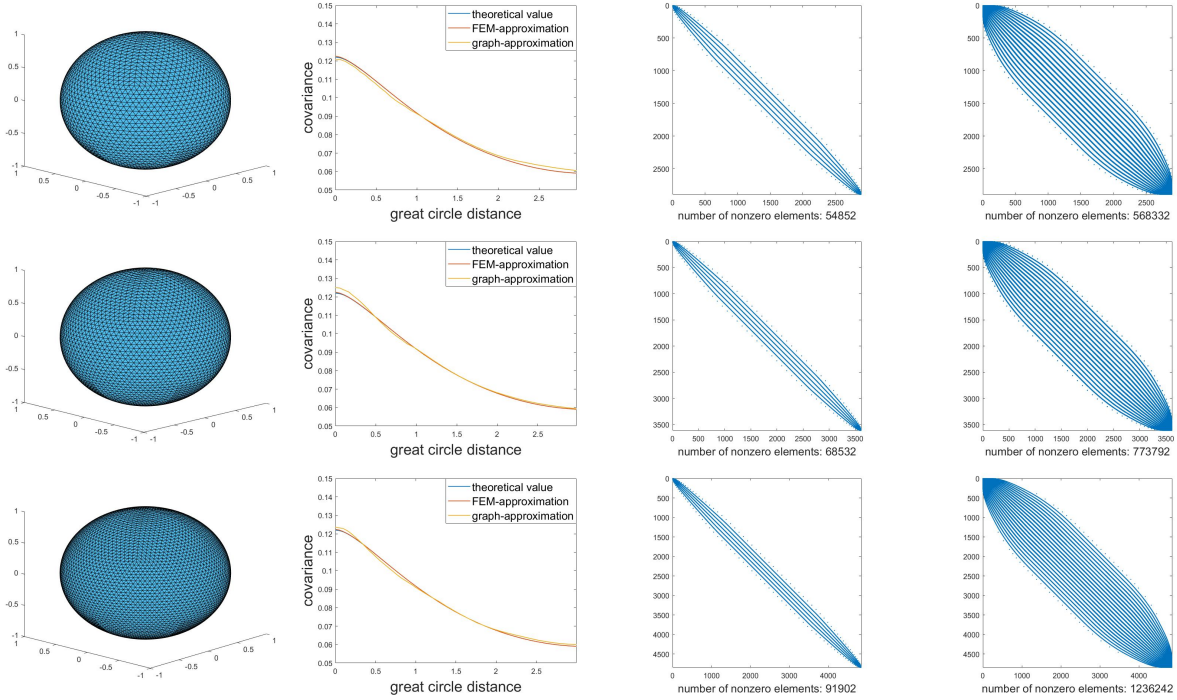


Figure 2.3: The three rows represent simulations for  $n = 2892, 3612, 4842$  respectively. Within each row, the plots represent (1) Visualization of the triangulation. (2) Covariances: theoretical (blue), FEM-based (red), and graph-based (yellow). (3). Sparsity pattern of FEM-based precision. (4). Sparsity pattern of graph-based precision.

### 2.4.1 Setup and Main Result

Recall from (2.6) and (2.13) that draws from the continuum and graph Matérn fields are defined by the series

$$u = \tau^{\frac{s}{2} - \frac{m}{4}} \kappa^{\frac{m}{2}} \sum_{i=1}^{\infty} \left[ \lambda^{(i)} \right]^{\frac{s}{2}} \xi^{(i)} \psi^{(i)}, \quad (2.19)$$

$$u_n = \tau_n^{\frac{s}{2} - \frac{m}{4}} \kappa_n^{\frac{m}{2}} \sum_{i=1}^n \left[ \lambda_n^{(i)} \right]^{\frac{s}{2}} \xi^{(i)} \psi_n^{(i)}. \quad (2.20)$$

We seek to establish convergence of  $u_n$  towards  $u$ . But notice that  $u_n$  is only defined on the point cloud  $\mathcal{M}_n$ , while  $u$  is defined on  $\mathcal{M}$ . Therefore a natural route is to introduce an interpolation scheme, using ideas from optimal transport theory.



Recall the manifold assumption 2.3.1 that  $\{x_i\}_{i=1}^\infty$  is a sequence of independent samples from the uniform distribution  $\gamma$  on  $\mathcal{M}$  and denote by  $\gamma_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$  the empirical measure of  $\{x_i\}_{i=1}^n$ . It should be intuitively clear that for the graph Matérn fields to approximate well their continuum counterparts, the point cloud  $\{x_i\}_{i=1}^n$  needs to first approximate  $\mathcal{M}$  well. The following result ([García Trillos et al., 2019a, Theorem 2] together with Borel-Cantelli) measures such approximation quantitatively.

**Proposition 2.4.1.** *There is a constant  $C$  such that, with probability one, there exists a sequence of transport maps  $T_n : \mathcal{M} \rightarrow \mathcal{M}_n$  so that  $\gamma_n = T_{n\#}\gamma$  and*

$$\limsup_{n \rightarrow \infty} \frac{n^{1/m} \sup_{x \in \mathcal{M}} d_{\mathcal{M}}(x, T_n(x))}{(\log n)^{c_m}} \leq C, \quad (2.21)$$

where  $c_m = 3/4$  if  $m = 2$  and  $c_m = 1/m$  otherwise.

Here  $d_{\mathcal{M}}$  denotes the geodesic distance on  $\mathcal{M}$  and the notation  $\gamma_n = T_{n\#}\gamma$  means that  $\gamma(T_n^{-1}(U)) = \gamma_n(U)$  for all measurable  $U$ , so that  $T_n$  is a measure preserving map. Intuitively  $T_n$  transports the mass of  $\mathcal{M}$  to the points  $\{x_i\}_{i=1}^n$ , so that the preimage of each singleton gets  $1/n$  of the mass, i.e.,  $\gamma(T_n^{-1}(\{x_i\})) = 1/n$ . Furthermore, the sets  $U_i := T_n^{-1}(\{x_i\})$  form a partition of  $\mathcal{M}$  and by Proposition 2.4.1 we have  $U_i \subset B_{\mathcal{M}}(x_i, \varepsilon_n)$ , where

$$\varepsilon_n := \sup_{x \in \mathcal{M}} d_{\mathcal{M}}(x, T_n(x)) \lesssim \frac{(\log n)^{c_m}}{n^{1/m}} \quad (2.22)$$

and  $B_{\mathcal{M}}(x, r)$  refers to the geodesic ball centered at  $x$  with radius  $r$ . Therefore each  $U_i$  is “centered around”  $x_i$  and the function

$$u_n \circ T_n(x) = \sum_{i=1}^n u_n(x_i) \mathbf{1}_{U_i}(x), \quad x \in \mathcal{M},$$

can be thought of as a locally constant interpolation of  $u_n$  to a function on  $\mathcal{M}$ . This motivates us to quantify the convergence of graph Matérn fields by the expected  $L^2 := L^2(\gamma)$ -norm

between  $u_n \circ T_n$  and  $u$ . Recall that  $h_n$  is the graph connectivity that is crucial in defining  $u_n$ .

**Theorem 2.4.2.** *Suppose  $\tau$  is Lipschitz,  $\kappa \in C^1(\mathcal{M})$  and both are bounded below by positive constants. Let  $s > m$  and*

$$\frac{(\log n)^{c_m}}{n^{1/m}} \ll h_n \ll \frac{1}{n^{1/2s}}, \quad (2.23)$$

where  $c_m = 3/4$  if  $m = 2$  and  $c_m = 1/m$  otherwise. Then, with probability one,

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} \xrightarrow{n \rightarrow \infty} 0,$$

where  $\{T_n\}_{n=1}^\infty$  is a sequence of transport maps as in Proposition 2.4.1. If further  $s > (5m + 1)/2$  and

$$h_n \asymp \sqrt{\frac{(\log n)^{c_m}}{n^{1/m}}}, \quad (2.24)$$

then, with probability one,

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} = O\left(\sqrt{h_n}\right) = O\left(\frac{(\log n)^{c_m/4}}{n^{1/4m}}\right).$$

**Remark 2.4.3.** *Notice that  $\varepsilon_n$  as defined in (2.22) represents the finest scale of variations that the point cloud can resolve and hence  $h_n$  needs to be much larger than  $\varepsilon_n$  to capture local geometry, which is reflected in the lower bound of (2.23). We will see below that the scaling (2.24) gives the optimal convergence rate. In such cases the geodesic ball of radius  $h_n$  has volume  $O(n^{-1/2})$  up to logarithmic factors and hence the average degree of the graph is  $O(n^{1/2})$  up to logarithmic factors as the  $x_i$ 's are uniformly distributed. Therefore, the number of nonzero elements in the weight matrix (and hence the graph Laplacian) is  $O(n^{3/2})$  up to*

logarithmic factors.

Although not as sparse as the stiffness matrix in FEM approach which has  $O(n)$  nonzero entries, the graph Laplacian still gives computational speed-up. Since we only assume a random design for the point cloud  $x_i$ 's, we are essentially analyzing the worst case scenario. Therefore we expect that the convergence rate and the resulting sparsity of the graph Laplacian can be improved if the point cloud is structured. However this will require a different analysis which is beyond the scope of this chapter.  $\square$

In previous related work, the  $T_n$ 's are taken to be the optimal transport maps, whose computation can be quite challenging, especially in high dimensions [Peyré and Cuturi, 2019]. Therefore we consider an alternate interpolation map below that can be computed efficiently. Consider the Voronoi cells

$$V_i = \{x \in \mathcal{M} : |x - x_i| = \min_{j=1, \dots, n} |x - x_j|\}.$$

Up to a set of ambiguity of  $\gamma$ -measure 0, the  $V_i$ 's form a partition of  $\mathcal{M}$  and we shall assume they are disjoint by assigning points in their intersections to only one of them. We then define a map  $\mathcal{T}_n : \mathcal{M} \rightarrow \mathcal{M}_n$  by  $\mathcal{T}_n(x) = x_i$  for  $x \in V_i$  and consider the nearest-neighbor interpolation

$$u_n \circ \mathcal{T}_n(x) = \sum_{i=1}^n u_n(x_i) \mathbf{1}_{V_i}(x), \quad x \in \mathcal{M}.$$

Note its resemblance with  $u_n \circ T_n$ . Indeed, the sets  $V_i$  and  $U_i$ 's are comparable and we have a similar result as Theorem 2.4.2.

**Theorem 2.4.4.** *Suppose  $\tau$  is Lipschitz,  $\kappa \in C^1(\mathcal{M})$  and both are bounded below by positive*

constants. Let  $s > (5m + 1)/2$  and

$$h_n \asymp \sqrt{\frac{(\log n)^{c_m}}{n^{1/m}}}.$$

Then, with probability one,

$$\mathbb{E}\|u_n \circ \mathcal{T}_n - u\|_{L^2} = O\left(\frac{(\log n)^{(2m+1)c_m/4}}{n^{1/4m}}\right).$$

**Remark 2.4.5.** We see that the rate has an additional factor  $(\log n)^{mc_m/2}$ , which comes from the bound for eigenfunction approximation using the map  $\mathcal{T}_n$ . In fact, due to this additional logarithmic term, we were unable to show convergence for general  $s > m$  using the same strategy as before.  $\square$

### 2.4.2 Outline of Proof

From the Karhunen-Loève expansions (2.19) and (2.20) we can see that the fundamental step in analyzing the convergence of the graph representations is to obtain bounds for both the eigenvalue and eigenfunction approximations.

Since  $L_n^{\tau, \kappa}$  and  $\mathcal{L}^{\tau, \kappa}$  are self-adjoint with respect to the  $L^2(\gamma_n)$  and  $L^2(\gamma)$  inner products, respectively, and are both positive definite, their spectra are real and positive. Denote by  $\{(\lambda_n^{(i)}, \psi_n^{(i)})\}_{i=1}^n$  and  $\{(\lambda^{(i)}, \psi^{(i)})\}_{i=1}^\infty$  the eigenpairs of  $L_n^{\tau, \kappa}$  and  $\mathcal{L}^{\tau, \kappa}$ , with the eigenvalues in increasing order. Suppose we are in a realization where the conclusion of Proposition 2.4.1 holds, and recall the definition of  $\varepsilon_n$  in (2.22). We formally have the following results on spectral approximations. (Rigorous statements can be found in Appendices 2.7.2 and 2.7.3.)

**Theorem 2.4.6** (Eigenvalue Approximation cf. Theorem 2.7.8). *Suppose  $\varepsilon_n \ll h_n$  and*

$h_n \sqrt{\lambda^{(k_n)}} \ll 1$ . Then

$$\frac{|\lambda_n^{(i)} - \lambda^{(i)}|}{\lambda^{(i)}} = O\left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}}\right),$$

for  $i = 1, \dots, k_n$ .

**Theorem 2.4.7** (Eigenfunction Approximation cf. Theorem 2.7.12). *Suppose  $\varepsilon_n \ll h_n$  and  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$ . Then there exists orthonormalized eigenfunctions  $\{\psi_n^{(i)}\}_{i=1}^n$  and  $\{\psi^{(i)}\}_{i=1}^\infty$  so that*

$$\begin{aligned} \|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} &= O\left(i^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}}}\right), \\ \|\psi_n^{(i)} \circ \mathcal{T}_n - \psi^{(i)}\|_{L^2} &= O\left((\log n)^{\frac{mcm}{2}} i^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}}}\right), \end{aligned}$$

for  $i = 1, \dots, k_n$ .

Theorems 2.4.6 and 2.4.7 generalize existing results in Burago et al. [2015], García Trillos et al. [2019a] where  $\tau$  and  $\kappa$  are constant. Theorem 2.4.2 and 2.4.4 are shown in Appendix 2.7.4 based on these two results.

Here we briefly illustrate the main idea for  $u_n \circ T_n$ . The assumption that  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$  in Theorems 2.4.6 and 2.4.7 is crucial in that the spectral approximations are only provably accurate up to the  $k_n$ -th eigenvalue and eigenfunction. Therefore to bound the difference between  $u_n \circ T_n$  and  $u$ , we need to consider the truncated series

$$u_n^{k_n} := \tau_n^{\frac{s}{2} - \frac{m}{4}} \kappa_n^{\frac{m}{2}} \sum_{i=1}^{k_n} \left[\lambda_n^{(i)}\right]^{\frac{s}{2}} \xi^{(i)} \psi_n^{(i)}, \quad (2.25)$$

and such truncation introduces an error of order  $\sqrt{n} k_n^{-s/m}$ . Therefore  $k_n$  needs to satisfy  $n^{m/2s} \ll k_n \ll h_n^{-m}$  and this explains the upper bound on  $h_n$  in (2.23). By repeated application of the triangle inequality, we can show that  $\mathbb{E}\|u_n \circ T_n - u\|_{L^2}$  is dominated by

the error coming from the truncation and eigenfunction approximation:

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} \lesssim \sqrt{n}k_n^{-\frac{s}{m}} + \sum_{i=1}^{k_n} \left[\lambda^{(i)}\right]^{-\frac{s}{2}} \|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2}. \quad (2.26)$$

If we are only interested in showing convergence without a rate, then we can first fix an  $\ell \in \mathbb{N}$  and use the fact that  $\|\psi_n^{(i)} \circ T_n\|_{L^2}^2 = n^{-1} \sum_{i=1}^n \psi_n^{(i)}(x_i)^2 = \|\psi_n^{(i)}\|_{L^2(\gamma_n)}^2 = 1$  to get

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} \lesssim \sqrt{n}k_n^{-\frac{s}{m}} + \sum_{i=1}^{\ell} \left[\lambda^{(i)}\right]^{-\frac{s}{2}} \|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} + \sum_{i=\ell+1}^{k_n} \left[\lambda^{(i)}\right]^{-\frac{s}{2}}.$$

Then letting  $n \rightarrow \infty$ , we have  $\|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} \rightarrow 0$  for  $i = 1, \dots, \ell$  and

$$\limsup_{n \rightarrow \infty} \mathbb{E}\|u_n \circ T_n - u\|_{L^2} \lesssim \sum_{i=\ell+1}^{\infty} \left[\lambda^{(i)}\right]^{-\frac{s}{2}}.$$

The last expression goes to 0 if we let  $\ell \rightarrow \infty$  given that  $s > m$ .

However if we want to derive a rate of convergence, then  $k_n$  needs to be chosen carefully in (2.26):  $k_n$  should be small so that the spectral approximations up to level  $k_n$  are sufficiently accurate, but at the same time not be too small to leave a large truncation error. In particular, by Theorem 2.4.7 and Weyl's law an upper bound on the second term (2.26) is

$$k_n [\lambda^{(k_n)}]^{-s/2} \max_{i=1, \dots, k_n} \|\psi_n^{(k_n)} - \psi^{(k_n)}\|_{L^2} \lesssim k_n^{-\frac{s}{m} + \frac{5}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n k_n^{\frac{1}{m}}}. \quad (2.27)$$

If we only have  $s > m$ , then this does not go to zero for any choice of  $h_n$  given the constraint that  $n^{m/2s} \ll k_n$ . Hence we need the additional assumption that  $s > \frac{5}{2}m + \frac{1}{2}$ , which allows us to bound (2.27) by  $\sqrt{\frac{\varepsilon_n}{h_n} + h_n}$  and then

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} \lesssim \sqrt{n}k_n^{-\frac{s}{m}} + \sqrt{\frac{\varepsilon_n}{h_n} + h_n}.$$

Therefore optimal rates are obtained by setting  $h_n \asymp \sqrt{\varepsilon_n}$  and  $k_n$  correspondingly, which together with (2.22) gives the scaling (2.24).

**Remark 2.4.8.** *The paper García Trillos et al. [2020b] proposed to consider directly the truncated series (2.25) as the graph field and this will allow us to obtain rates of convergence in Theorem 2.4.2 for all  $s > m$ . But in this way the precision matrix representation as in (2.14) no longer holds and one will need to perform spectral decomposition on  $L_n^{\tau, \kappa}$  for sampling, which generally is more costly than Cholesky factorization.*  $\square$

## 2.5 Numerical Examples

In this section we demonstrate the use of the graph Matérn models introduced in Section 2.3 by considering applications in Bayesian inverse problems, spatial statistics and graph-based machine learning.

For the three examples we employ graph Matérn models as priors within the general framework of latent Gaussian models, briefly overviewed in Section 2.5.1. Section 2.5.2 studies a toy Bayesian inverse problem on a manifold setting. Our aim is to compare the modeling of length scale through  $\tau_n$  and  $\kappa_n$ ; we further show that the accuracy of the reconstruction with the graph-based approach is satisfactory and that adding nonstationarity may help to overcome the poor performance of more naive hierarchical approaches in large noise regimes. In Section 2.5.3 we investigate the use of graph Matérn fields for interpolating U.S. county-level precipitation data, assuming to only have access to pairwise distances between counties and precipitation data for some of them. Contrary to finite element representations, the graph-based approach is applicable in this discrete setting without the need of performing multidimensional scaling to reconstruct the configuration of the point cloud. In addition, the graph approach does not require to introduce any artificial nodes. We also compare the performance of stationary and nonstationary graph Matérn models. In Section 2.5.4, a semi-supervised classification problem in machine learning is studied, where the low dimensional

structure of the data naturally motivates the graph-based approach; we further show that nonstationary models may improve the classification accuracy over stationary ones.

### 2.5.1 A General Framework: Latent Gaussian Models

Latent Gaussian models are a flexible subclass of structured additive regression models defined in terms of a likelihood function, a latent process and hyperparameters. Let  $\{x_i\}_{i=1}^n$  be a collection of features that we identify with graph nodes. The observation variable  $y$  is modeled as a (possibly noisy) transformation of the latent process  $u_n := [u_n(x_1), \dots, u_n(x_n)]^T$ , which conditioned on the hyperparameters follows a Gaussian distribution. Finally, a prior is placed on the hyperparameters. More precisely, we have

$$\begin{aligned} y|u_n, \mu &\sim \pi(y|u_n, \mu), \\ u_n|\theta &\sim \mathcal{N}\left(0, Q(\theta)^{-1}\right), \\ (\mu, \theta) &\sim \pi(\mu, \theta), \end{aligned}$$

where  $Q(\theta)$  is the precision matrix of the latent process and  $\mu, \theta$  are hyperparameters. Markov Chain Monte Carlo inference methods are standard in Bayesian inverse problems with complex likelihood functions, but less computationally expensive deterministic approximations are often preferred in other applications. In particular, the integrated nested Laplace approximations proposed by Rue et al. [2009] and the corresponding R-INLA package has greatly facilitated inference of such models.

The sparsity of the precision matrix  $Q(\theta)$  is crucial for efficient likelihood evaluations and sampling of the latent process. For the problems that we consider, the latent process will be modeled as a graph Matérn field, i.e.,

$$u_n|\tau_n \sim \mathcal{N}\left(0, Q(\tau_n, s)^{-1}\right), \quad Q(\tau_n, s) = \tau_n^{\frac{m}{4} - \frac{s}{2}} (\tau_n + \Delta_n)^s \tau_n^{\frac{m}{4} - \frac{s}{2}},$$



where  $\Delta_n$  is a graph Laplacian constructed with the  $x_i$ 's and  $\tau_n$  is a diagonal matrix modeling the length scale at each node. We note that the graph Laplacian is often sparse and its sparsity is inherited by  $Q(\tau_n, s)$  for small or moderate integer  $s$ .

A constant length scale graph Matérn field hyperprior is then placed on  $\log(\tau_n)$ :

$$\log \tau_n \sim \mathcal{N} \left( 0, \nu^{s_0 - \frac{m}{2}} (\nu I + \Delta_n)^{-s_0} \right), \quad (2.28)$$

where  $\nu$  and  $s_0$  are chosen by prior belief on the length scale. However, when  $n$  is large, learning  $\tau_n$  as an  $n$ -dimensional vector is computationally demanding. We instead adopt a truncated Karhunen-Loève approximation for  $\tau_n$ . Recall that  $\log \tau_n$  has the characterization

$$\log \tau_n = \nu^{\frac{s_0}{2} - \frac{m}{4}} \sum_{i=1}^n \left[ \nu + \lambda_n^{(i)} \right]^{-\frac{s_0}{2}} \xi^{(i)} \psi_n^{(i)},$$

where  $\xi^{(i)} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$  and  $\{(\lambda_n^{(i)}, \psi_n^{(i)})\}_{i=1}^n$  are the eigenpairs of  $\Delta_n$ . Since the  $\lambda_n^{(i)}$ 's are increasing, the contribution of the higher frequencies is less significant. Hence we consider a truncated expansion and model  $\log \tau_n$  as

$$\log \tau_n = \nu^{\frac{s_0}{2} - \frac{m}{4}} \sum_{i=1}^{n_0} \left[ \nu + \lambda_n^{(i)} \right]^{-\frac{s_0}{2}} \theta^{(i)} \psi_n^{(i)},$$

where  $n_0 \ll n$  is chosen based on the spectral growth and prior belief on  $\tau_n$ . Now the hyperparameters are the  $\theta_i$ 's, which are only  $n_0$ -dimensional, and the hyperprior for each is naturally taken to be the standard normal. Therefore a complete model of our interest in the following subsections can be summarized as

$$\begin{aligned} y|u, \mu &\sim \pi(y|u, \mu), \\ u|\theta, s &\sim \mathcal{N} \left( 0, Q(\theta, s)^{-1} \right), \\ (\mu, s) &\sim \pi(\mu, s), \quad \theta \sim \mathcal{N}(0, I_{n_0}), \end{aligned}$$

where

$$Q(\theta, s) = [\tau_n(\theta)]^{\frac{m}{4} - \frac{s}{2}} [\tau_n(\theta) + \Delta_n]^s [\tau_n(\theta)]^{\frac{m}{4} - \frac{s}{2}}, \quad (2.29)$$

$$\log \tau_n(\theta) = \nu^{\frac{s_0}{2} - \frac{m}{4}} \sum_{i=1}^{n_0} \left[ \nu + \lambda_n^{(i)} \right]^{-\frac{s_0}{2}} \theta^{(i)} \psi_n^{(i)}. \quad (2.30)$$

**Remark 2.5.1.** *We remark that (2.30) can be viewed as representing  $\tau_n$  as a combination of the first several eigenfunctions of  $\Delta_n$ , which are a natural basis for functions over the point cloud. A disadvantage is that (2.30) requires knowledge of the  $\psi_n^{(i)}$ 's. However since  $\Delta_n$  is sparse and we only need to know the first  $n_0 \ll n$  eigenfunctions, the computational cost for  $\tau_n$  is still better than  $O(n^3)$ .  $\square$*

**Remark 2.5.2.** *The marginal variance of the latent process  $u_n$  can be tuned and fixed easily by matching the scales of  $u_n$  and the data  $y$ , and thus we do not include a marginal variance parameter. Indeed, the normalizing factors  $\tau_n^{m/4 - s/2}$  guarantee that  $\mathbb{E}|u_n|^2$  are roughly the same for different  $\tau_n$ 's. Hence one can for instance estimate  $\mathbb{E}|u_n|^2$  by setting  $\tau_n \equiv 1$  and normalize the observations  $y$  by  $\sqrt{\mathbb{E}|u_n|^2/|y|^2}$ .*

*Similarly, we need to tune for the marginal variance of the hyperparameters as in (2.28). As  $\tau_n$  essentially acts as a cut off on the significant frequencies, this can be done by matching the scale of  $\tau_n$  with the eigenvalues of  $\Delta_n$  based on one's prior belief.  $\square$*

Suppose for illustration that we are interested in the simple regression problem of inferring a Matérn field  $u(x)$  based on data  $y$  comprising Gaussian measurement of  $u$  at given locations/features  $x_1, \dots, x_n$ . As noted in the introduction, the computation cost scales as  $\mathcal{O}(n^3)$ . However, by modeling  $u_n$  using a graph Matérn model we obtain a GMRF approximation, with sparse precision matrix, dramatically reducing the computational cost. Thus, one could introduce further auxiliary nodes  $x_{n+1}, \dots, x_N$  with  $N \gg n$  to improve the prior GMRF approximation of the original Matérn model and still reduce the computational cost over formulations based on GF priors. Such ideas arise naturally in semi-supervised appli-

cations in machine learning where most features are unlabeled, but can also be of interest in applications in spatial statistics, as discussed in [Rue and Held, 2005, Chapter 5] and grant further investigation in Bayesian inverse problems.

### 2.5.2 Application in Bayesian Inverse Problems

In this subsection we investigate the use of nonstationary graph Matérn models to define prior distributions in Bayesian inverse problems. For simplicity of exposition and to avoid distraction from our main purpose of illustrating the modeling of the nonstationarity, we consider a toy example taken from the inverse problem literature [Roininen et al., 2019]. The ideas presented here apply immediately to Bayesian inverse problems with more involved likelihood functions, defined for instance in terms of the solution operator of a differential equation [Harlim et al., 2020, Bigoni et al., 2020].

We study the reconstruction of a signal function given noisy but direct point-wise observations. The domain of the problem is taken to be the unit circle, where the hidden signal  $u^\dagger$  is parametrized by  $t \in [0, 2\pi)$  as

$$u^\dagger(t) = \begin{cases} \exp\left(4 - \frac{\pi^2}{t(\pi-t)}\right), & t \in (0, \pi), \\ 1, & t \in [\pi + 0.5, 1.5\pi], \\ -1, & t \in (1.5\pi, 2\pi - 0.5], \\ 0, & \text{otherwise.} \end{cases}$$

Hence if  $x = (\cos(t), \sin(t))$  for  $t \in [0, 2\pi)$ , then  $u^\dagger(x)$  is understood as  $u^\dagger(t)$ . Such signal is considered in Roininen et al. [2019] for its varying length scale, where the domain is the interval  $[0, 10] \subset \mathbb{R}$  and a uniform grid finite difference discretization is used to define a Matérn prior following the SPDE approach. Here we suppose instead to have only indirect access to the domain through  $n = 1000$  points  $x_i$ 's that are drawn independently from the

uniform distribution on the circle, and use a graph Matérn model.

We assume to be given noisy observations of the signal at  $J = n/2$  points:

$$y(x_i) = u^\dagger(x_i) + \eta_i, \quad \eta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2), \quad i = 1, \dots, J, \quad (2.31)$$

where  $\sigma$  is set to be 0.1 and we have observations at every other node. To recover the signal function at the nodes  $x_i$ 's, we adopt a hierarchical Bayesian approach which we cast into the framework of latent Gaussian models. More precisely, the observation equation (2.31) gives the likelihood model

$$y|u_n \sim \mathcal{N}(Su_n, \sigma^2 I_J),$$

where  $S \in \mathbb{R}^{J \times n}$  is a matrix of 0's and 1's that indicates the location of the observations. The latent process  $u_n := [u^\dagger(x_1), \dots, u^\dagger(x_n)]^T$  and the hyperparameters are modeled as in Section 2.5.1, where the smoothness is fixed as  $s = 2$  and the other parameters are chosen as  $s_0 = 1, \nu = 10, n_0 = 21$ . We note that by setting  $s_0 = 1$ , the hyperprior is actually an approximation of a zero-mean Gaussian process with exponential covariance function, where the sample paths can undergo sudden changes. The choice  $n_0 = 21$  is motivated by the fact that the Laplacian on the circle has eigenvalues  $\{i^2\}_{i=0}^\infty$ , where any non-zero eigenvalue has multiplicity 2. Therefore the cutoff is at the eigenvalue 100, which is an order of magnitude larger than  $\nu$ , and higher frequencies are less consequential. The graph Laplacian is constructed as in Section 2.3.1 with connectivity  $\varepsilon_n = 4 \times n^{-1/1.8}$ .

We will follow a similar MCMC sampling as in Roininen et al. [2019] detailed in Algorithm 1 for inferring the signal function  $u_n$  together with the length scale  $\tau_n$ . To illustrate the idea, we notice that the observation equation and the graph Matérn model for  $u_n$  translate into

---

**Algorithm 1** Posterior sampling  $u, \tau|y$ 


---

Initialize  $u_n^0, \theta^0$  and set step size  $\beta_\theta$ .

**for**  $i = 1, \dots, M$  **do**

(a) Generate  $\xi \sim \mathcal{N}(0, I_{J+n})$  and update

$$u_n^i = \begin{bmatrix} \sigma^{-1}S \\ L(\theta^{i-1}) \end{bmatrix}^\dagger \left( \begin{bmatrix} \sigma^{-1}y \\ 0 \end{bmatrix} + \xi \right),$$

where  $L(\theta)$  is the Cholesky factor of  $Q(\theta)$  and  $\dagger$  denotes the matrix pseudoinverse.

(b)

**for**  $j = 1, \dots, n_0$  **do**

(i) Generate  $\xi^{(j)} \sim \mathcal{N}(0, 1)$  and propose  $\tilde{\theta}^{i,(j)} = \theta^{i-1,(j)} + \beta_\theta \xi^{(j)}$ .

(ii) Denote  $\theta^{i,(-j)} = (\theta^{i,(1)}, \dots, \theta^{i,(j-1)}, \theta^{i-1,(j+1)}, \dots, \theta^{i-1,(n_0)})$ . Accept  $\tilde{\theta}^{i,(j)}$

with probability

$$p = \min \left\{ 1, \frac{\pi(\tilde{\theta}^{i,(j)}, \theta^{i,(-j)} | u_n^i, y)}{\pi(\theta^{i-1,(j)}, \theta^{i,(-j)} | u_n^i, y)} \right\},$$

where  $\pi(\theta | u_n, y)$  is given in (2.32).

**end for**

**end for**

---

the equations

$$\sigma^{-1}S u_n = \sigma^{-1}y + \xi_1,$$

$$L(\theta)u_n = \xi_2,$$

where  $L(\theta)$  is the Cholesky factor of  $Q(\theta)$  and  $\xi_1 \sim \mathcal{N}(0, I_J), \xi_2 \sim \mathcal{N}(0, I_n)$ . The above pair of equations motivate the update for  $u_n$  as

$$\begin{bmatrix} \sigma^{-1}S \\ L(\theta) \end{bmatrix}^\dagger \left( \begin{bmatrix} \sigma^{-1}y \\ 0 \end{bmatrix} + \xi \right),$$

where  $\xi \sim \mathcal{N}(0, I_{J+n})$ . The hyperparameters  $\theta$  are updated with a Metropolis-within-Gibbs

sampling scheme, where the full posterior  $\theta|u_n, y$  has the form

$$\pi(\theta|u_n, y) \propto \sqrt{\det(Q(\theta))} \exp\left(-\frac{1}{2}u_n^T Q(\theta)u_n - \frac{1}{2}|\theta|^2\right). \quad (2.32)$$

In Figures 2.4a and 2.4b we plot the posterior means of  $u_n$  and  $\tau_n$  and 95% credible intervals for each of their coordinates. The oscillatory paths of the reconstructions are due to the graph approximation, where the eigenfunctions  $\psi_n^{(i)}$  of the graph Laplacian are in general very ragged. We notice that  $\tau_n$  varies rapidly in the region where the signal is piecewise constant, indicating a change of length scale. Moreover, the sudden jump of the signal from 1 to -1 suggests a small local length scale which leads to a larger  $\tau_n$ , as predicted in Remark 2.2.2.

To further understand the effect of modeling the length scale through  $\kappa_n$ , we choose a different prior for the latent process by setting

$$Q(\theta) = \kappa_n(\theta)^{-\frac{m}{4}} (I + \Delta_n^\kappa)^s \kappa_n(\theta)^{-\frac{m}{4}},$$

$$\log \kappa_n(\theta) = \nu^{\frac{s_0}{2} - \frac{m}{4}} \sum_{i=1}^{n_0} \left[\nu + \lambda_n^{(i)}\right]^{-\frac{s_0}{2}} \theta^{(i)} \psi_n^{(i)},$$

so that the length scale is controlled by  $\kappa_n$  instead. Here  $\Delta_n^\kappa$  is the discrete approximation of  $\nabla \cdot (\kappa \nabla)$  introduced in Section 2.3.2 and we adopt the same hyperparameter modeling as for  $\tau_n$ . In Figures 2.5a and 2.5b, we plot the posterior means and 95% credible intervals for each coordinate of  $u_n$  and  $\kappa_n$ , as a comparison with their counterparts for  $\tau_n$ . The figures show that the two approaches give similar reconstructions for the signal and, in agreement with the intuition given in Remark 2.2.2,  $\tau_n$  and  $\kappa_n$  are almost inversely proportional to each other.

**Remark 2.5.3.** *As noted in [Harlim et al., 2020, Section 4.5], the hierarchical approach performs poorly if the noise in the observations is large and the latent process is modeled*

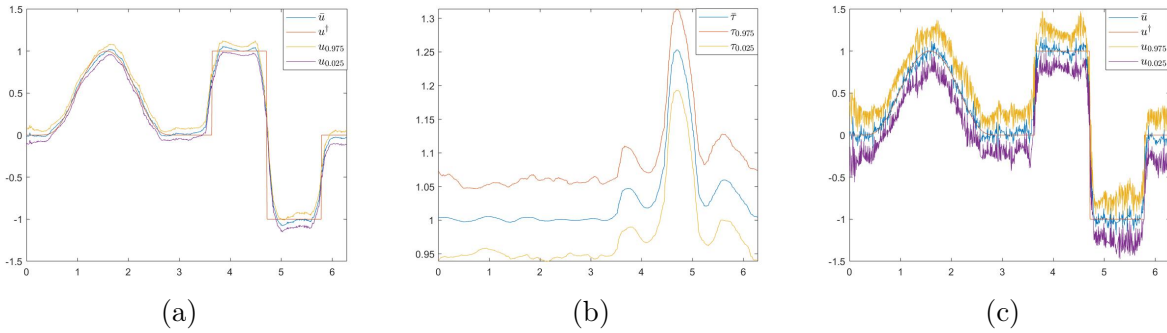


Figure 2.4: Posterior means and 95% credible intervals for (a) signal from nonstationary model; (b) length scale of nonstationary model; (c) signal from stationary model (constant  $\tau_n$ ) when length scale is modeled through  $\tau_n$ .




Figure 2.5: Posterior means and 95% credible intervals for (a) signal from nonstationary model; (b) length scale of nonstationary model; (c) signal from stationary model constant  $\kappa_n$ ), when length scale is modeled through  $\kappa_n$ .

with a constant length scale; the single length scale is blurred by the noise and the model essentially fits the noisy observations, as can be seen from the oscillatory reconstructions in Figures 2.4c and 2.5c. An important observation stemming from the above example is that adding nonstationarity into the length scale may help alleviate such issue.  $\square$

### 2.5.3 Application in Spatial Statistics

In this subsection we consider interpolation of county-level precipitation data in the U.S. for January 1981, available from <https://www.ncdc.noaa.gov/cag/county/mapping>. Similar problems have been studied in Fuglstad et al. [2015b], Bolin and Kirchner [2020] using

the SPDE formulation and finite element representations, adding nodes for triangulation of the space. Here we shall assume, however, that *only* pairwise distances between counties are available. In particular, in contrast to Fuglstad et al. [2015b], Bolin and Kirchner [2020] we do not assume to have access to the spatial domain or to the locations of the observations, and we will perform inference without adding artificial nodes. The inter-county distances are available as great circle distances from <https://data.nber.org/data/county-distance-database.html> and are only recorded for each pair of counties that are closer than a certain threshold distance apart, which naturally suggests a graph representation. Let  $\{x_i\}_{i=1}^{n=3107}$  denote the  $n$  counties (excluding Alaska, Hawaii and several other counties that we do not have precipitation data for). We model the precipitation  $y$  with a latent Gaussian model, where the  $J$  observations are given as noisy perturbations of the latent process  $u_n$ :

$$y|u, \sigma \sim \mathcal{N}(Su, \sigma^2 I_J),$$

where  $S \in \mathbb{R}^{J \times n}$  is a matrix of 0's and 1's that specifies the observation locations. Notice that we have included the noise size  $\sigma$  as a hyperparameter to be inferred. The latent process  $u_n$  will be modeled in four different ways for comparison purposes as in Bolin and Kirchner [2020]. The idea is to consider a graph Matérn prior for  $u_n$ , where

$$u_n|\theta, s \sim \mathcal{N}\left(0, Q(\theta, s)^{-1}\right), \quad Q(\theta, s) = \tau_n(\theta)^{\frac{m}{4} - \frac{s}{2}} (\tau_n(\theta) + \Delta_n)^s \tau_n(\theta)^{\frac{m}{4} - \frac{s}{2}}, \quad m = 2,$$

and consider  $s$  to be possibly a hyperparameter, while at the same time allow  $\tau_n$  to be identically equal to a constant. In other words, we will model  $u_n$  as a stationary/nonstationary graph Matérn field with possibly fractional smoothness parameter. In the most general case, the length scale parameter  $\tau_n$  is modeled as in (2.30), where the parameters are chosen as



$s_0 = 2, \nu = 0.1, n_0 = 10$  and the hyperpriors for  $\sigma, s$  are chosen as

$$\log \sigma \sim \mathcal{N}(\log(0.01), 1),$$

$$\log s \sim \mathcal{N}(\log 2, 1).$$

The marginal variance of  $\tau_n$  is then tuned to be on the same order as the 10th eigenvalue of  $\Delta_n$ . For the stationary case, the modeling for the latent process simplifies to

$$u_n | \tau, s \sim \mathcal{N}(0, Q(\tau, s)^{-1}), \quad Q(\tau, s) = \tau^{\frac{m}{2}-s} (\tau I_n + \Delta_n)^s,$$

and a log-normal hyperprior is placed on  $\tau$ .

For this problem  $\Delta_n$  is an unnormalized graph Laplacian  $D - W$ , with weights

$$W_{ij} = \begin{cases} \exp\left(-\frac{d_{ij}^2}{2\underline{d}^2}\right), & \text{if } d_{ij} \text{ is recorded,} \\ 0, & \text{otherwise,} \end{cases}$$

where  $d_{ij}$  is the distance between two counties and  $\underline{d}$  is the mean of all the pairwise distances. As mentioned above, the pairwise distances are only recorded for counties that are less than 100 miles apart, which implies that  $W$  is sparse (with the percentage of nonzero entries being 1.6%). Instead of using an MCMC sampling scheme, we adopt an evidence maximization approach, where we first compute the optimal hyperparameter  $(\sigma^*, s^*, \theta^*)$  (for the most general case) by maximizing the marginal posterior  $(\sigma, s, \theta) | y$ , and then compute the posterior  $u_n | y, \sigma^*, s^*, \theta^*$ . For the most general case, the marginal posterior of  $(\sigma, s, \theta)$  is equal, up to a constant, to

$$\begin{aligned} \log \pi(\sigma, s, \theta | y) &= \log \pi(\sigma, s, \theta) - J \log \sigma + \frac{1}{2\sigma^4} y^T S \tilde{Q}(\sigma, \theta, s)^{-1} S^T y - \frac{\|y\|_2^2}{2\sigma^2} \\ &+ \frac{1}{2} \log \left[ \det(Q(\theta, s)) - \det(\tilde{Q}(\sigma, \theta, s)) \right], \end{aligned}$$

where  $\tilde{Q}(\sigma, \theta, s) = \sigma^{-2}S^T S + Q(\theta, s)$  and the posterior  $u_n|y, \sigma^*, s^*, \theta^*$  is a Gaussian  $\mathcal{N}(\mu, \Sigma)$ , where

$$\mu = [\sigma^*]^{-2}\tilde{Q}(\sigma^*, \theta^*, s^*)^{-1}S^T y, \quad \Sigma = \tilde{Q}(\sigma^*, \theta^*, s^*)^{-1}. \quad (2.33)$$

The predictive distribution is then the restriction of  $\mathcal{N}(\mu, \Sigma)$  to the unobserved nodes, denoted by  $\pi_{\text{pred}}$ , and techniques for computing partial inverse of a sparse matrix can be applied. We numerically optimize  $\log \pi(\sigma, s, \theta|y)$  with the `fminunc` function in Matlab.

	RMSE	CRPS	LS
Stationary & $s = 2$	0.0394	0.0199	-578.9
Stationary & Inferred $s$	0.0399	0.0201	-521.7
Nonstationary & $s = 2$	0.0408	0.0185	-648.5
Nonstationary & Inferred $s$	0.0414	0.0186	-644.6

Table 2.1: Comparison of the four models through RMSE, CRPS and LS.

To perform inference, we first normalize the data  $y$  as described in Remark 2.5.2 so that it has mean-zero and has magnitude at the same level of  $u_n$  from the graph Matérn field, in which case we are only interested in the variations of  $u_n$ . We then adopt a pseudo-crossvalidation by randomly selecting 90% of the data as observations and make predictions for the remaining ones. The process is repeated 20 times and we evaluate the predictions through the root mean square error (RMSE), the continuous rank probability score (CRPS), and the logarithmic scoring rule (LS) as shown in Table 2.1. The three criteria are considered in Fuglstad et al. [2015b], Bolin and Kirchner [2020] for similar comparisons, with each

defined as

$$\begin{aligned}
\text{RMSE} &:= \frac{\|y_{\text{test}} - \mathbb{E}\pi_{\text{pred}}\|_2}{\sqrt{n_{\text{test}}}}, \\
\text{CRPS} &:= \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \text{crps}(\pi_{\text{pred}}^{(i)}, y_{\text{test}}^{(i)}) \\
&= \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \int_{-\infty}^{\infty} \left[ \pi_{\text{pred}}^{(i)}(y_{\text{test}}^{(i)}) - \mathbf{1}_{\{y_{\text{test}}^{(i)} \leq t\}} \right]^2 dt, \\
\text{LS} &:= -\log \pi_{\text{pred}}(y_{\text{test}}),
\end{aligned}$$

where  $y_{\text{test}}$  is the test data with size  $n_{\text{test}}$  and  $\pi_{\text{pred}}^{(i)}$  is the marginal distribution for each unobserved node. Since  $\pi_{\text{pred}}$  is Gaussian (and hence its marginals), the above quantities are computable, and the CRPS can be calculated with its representation for Gaussians Gneiting et al. [2005] :

$$\text{crps} \left( \mathcal{N}(a, b^2), y \right) = b \left\{ \frac{y-a}{b} \left[ 2\Phi \left( \frac{y-a}{b} \right) - 1 \right] + 2\phi \left( \frac{y-a}{b} \right) - \frac{1}{\sqrt{\pi}} \right\},$$

where  $\Phi$  and  $\phi$  are the c.d.f. and p.d.f. of the standard normal respectively.

We notice in Table 2.1 that the nonstationary model improves CRPS and LS but not RMSE over the stationary one, as was observed in Fuglstad et al. [2015b], Bolin and Kirchner [2020] using finite element representations of GFs. This suggests that adding nonstationarity improves the prediction in a distribution sense. However, inferring the smoothness  $s$  from data appears not to improve the predictions, in contrast to the results in Bolin and Kirchner [2020]. We believe this is due to the different formulations that we are taking, where Bolin and Kirchner [2020] adopts an SPDE approach and a rational approximation for fractional smoothness. It is also possible that this is in general a feature of the graph representation as similar observations are made in Section 2.5.4. We shall leave more in-depth investigations for future studies.

**Remark 2.5.4.** *Prediction at new locations can be done by augmenting the graph with new nodes and refitting the model. Interpolation is not readily applicable in this example since we assume to only be given pairwise distances between counties and to not have a continuous representation of the underlying manifold (here the map of the contiguous states of the U.S.). However, in examples where the underlying manifold is known and the graph nodes form a triangulation, we could construct a finite-element basis and carry out interpolation using the basis.*

#### 2.5.4 Application in Machine Learning

In this subsection we illustrate the use of graph Matérn priors in a Bayesian formulation of semi-supervised binary classification [Bertozzi et al., 2018]. We seek to classify images  $\{x_i\}_{i=1}^n$  of two different digits of the MNIST dataset given  $J \ll n$  noisy labels. Similarly as above, the problem is cast into a latent Gaussian model, where the labels are assumed to be a probit transform of the latent process  $u_n$ :

$$y_i = \text{sign}(u_n(x_i) + \eta_i), \quad i = 1, \dots, J,$$

where  $\eta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$ . The likelihood model associated with the above equation is

$$\pi(y|u_n, \sigma) = \prod_{i=1}^J \Phi\left(\frac{y_i u_n(x_i)}{\sigma}\right),$$

where  $\Phi$  is the c.d.f. of the standard normal. As in Section 2.5.3, the latent process  $u_n$  will be modeled as a graph Matérn field in four different ways, by considering stationary/nonstationary length scale and fixed/inferred smoothness  $s$ . For the most general case,

the latent process is modeled as

$$u_n | \theta, s \sim \mathcal{N}(0, Q(\theta, s)^{-1}), \quad Q(\theta, s) = \tau_n(\theta)^{\frac{m}{4} - \frac{s}{2}} (\tau_n(\theta) + \Delta_n)^s \tau_n(\theta)^{\frac{m}{4} - \frac{s}{2}},$$

and  $\tau_n$  is modeled as in (2.30):

$$\log \tau_n(\theta) = \nu^{\frac{s_0}{2} - \frac{m}{4}} \sum_{i=1}^{n_0} \left[ \nu + \lambda_n^{(i)} \right]^{-\frac{s_0}{2}} \theta^{(i)} \psi_n^{(i)}$$

with standard normal hyperprior on each of the  $\theta_i$ 's and a log-normal prior for  $s$ . As in Remark 2.3.4, the effective dimension  $m$  is about 4 and the other parameters are chosen as  $s_0 = 4, \nu = 0.1, n_0 = 10$ , with marginal variance of  $\tau_n$  tuned empirically. For this problem  $\Delta_n$  is taken to be a symmetric  $k$ -nearest neighbor graph Laplacian  $\Delta_n = I - D^{-1/2} W D^{-1/2}$ , with self-tuning weights proposed by Zelnik-Manor and Perona [2005]:

$$W_{ij} = \exp \left( \frac{-|x_i - x_j|^2}{2\delta(i)\delta(j)} \right),$$

where the images  $x_i$ 's are viewed as vectors in  $\mathbb{R}^{784}$  and  $\delta(i)$  is the Euclidean distance between  $x_i$  and its  $k$ -th nearest neighbor. The sparsity of  $\Delta_n$  follows from the  $k$ -nearest neighbor construction.

Similarly as in Section 2.5.3, we adopt an evidence maximization approach for inferring the optimal hyperparameters, which are then used to find the MAP estimator for  $u_n$ . However, since the likelihood is non Gaussian, there is no closed form formula for the marginal posterior of the hyperparameters  $\pi(\sigma, \theta, s | y)$ , and we then apply a Laplace approximation [Rue et al., 2009]. More precisely, denoting all the hyperparameters by  $z$ ,  $\pi(z | y)$  is approximated by

$$\pi(z | y) \propto \frac{\pi(u_n, z, y)}{\pi(u_n | z, y)} \Big|_{u_n = u_n^*} \approx \frac{\pi(u_n, z, y)}{\tilde{\pi}(u_n | z, y)} \Big|_{u_n = u_n^*}, \quad (2.34)$$

where  $u_n^*$  is the mode of  $\pi(u_n|z, y)$  and  $\tilde{\pi}(u_n|z, y)$  is its Laplace approximation at  $u_n^*$ . The log density for  $u_n|z, y$  has form

$$\log \pi(u_n|z, y) \propto \sum_{i=1}^J \Phi\left(\frac{y_i u_n(x_i)}{\sigma}\right) - \frac{1}{2} u_n^T Q(\theta, s) u_n, \quad (2.35)$$

and the mode  $u_n^*$  is found numerically with the Newton's method, where the gradient and Hessian of (2.35) are available analytically. The logarithm of the last expression in (2.34) is equal up to a constant to

$$\begin{aligned} & \log \pi(\sigma, \theta, s) - \frac{1}{2} [u_n^*]^T Q(\theta, s) u_n^* \\ & + \frac{1}{2} \log \left[ \det(Q(\theta, s)) - \det(\tilde{Q}(\sigma, \theta, s)) \right] + \sum_{i=1}^J \log \Phi\left(\frac{y_i u_n^*(x_i)}{\sigma}\right), \end{aligned}$$

where  $\tilde{Q}(\theta, s) = Q(\sigma, \theta, s) + H$  and  $H$  is diagonal with entries

$$H_{ii} = \frac{y_i u_n^*(x_i) \phi(\sigma^{-1} u_n^*(x_i))}{\sigma^2 \Phi(\sigma^{-1} y_i u_n^*(x_i))} + \left[ \frac{\phi(\sigma^{-1} u_n^*(x_i))}{\sigma \Phi(\sigma^{-1} y_i u_n^*(x_i))} \right]^2, \quad i = 1, \dots, J,$$

and zero otherwise. The priors on  $\sigma$  and  $s$  are taken to be

$$\log \sigma \sim \mathcal{N}(\log(0.1), 1),$$

$$\log s \sim \mathcal{N}(\log(4), 1).$$

Table 2.2 shows the classification error rates of the four different models for four pairs of digits, with  $n = 1000$  and  $J = 20$ , where each experiment is repeated 100 times. We see that the nonstationary model improves slightly the performance while the model with inferred smoothness does the opposite. We believe this may be due to the fact that  $s = 4$  is already a near optimal choice for this problem, or it may also be an intrinsic characteristic of the graph representations as mentioned in Section 2.5.3. Table 2.3 shows the classification error

rates for the inferred  $s$  case when the prior is taken to be narrower:

$$\log s \sim \mathcal{N}(\log 4, 0.01),$$

in which case the nonstationary model with inferred  $s$  also improves the prediction.

	3&8	5&8	4&9	7&9
Stationary & $s = 4$	8.90%	8.64%	17.67%	10.13%
Stationary & Inferred $s$	9.61%	9.51%	18.33%	11.00%
Nonstationary & $s = 4$	8.44%	7.23%	17.38%	9.37%
Nonstationary & Inferred $s$	8.77%	8.12%	19.92%	10.70%

Table 2.2: Classification error rates with 2% labeled data for different pairs of digits with hyperprior  $\log s \sim \mathcal{N}(\log 4, 1)$ .

	3&8	5&8	4&9	7&9
Stationary & Inferred $s$	8.91%	8.66%	17.69%	10.14%
Nonstationary & Inferred $s$	8.67%	7.06%	17.54%	9.71%

Table 2.3: Classification error rates with 2% labeled data for different pairs of digits with hyperprior  $\log s \sim \mathcal{N}(\log 4, 0.01)$ .

**Remark 2.5.5.** *Label prediction at new images can be done by augmenting the graph with new nodes and refitting the model, or by using a  $k$ -NN interpolation as described in García Trillos et al. [2020b].*

## 2.6 Conclusions and Open Directions

This chapter introduces graph representations of Matérn fields motivated by the SPDE approach. We have shown through rigorous analysis that graph Matérn fields approximate the Matérn model under a manifold assumption, and we have established an explicit rate of convergence. We have emphasized that graph Matérn models can be used in a wide range of settings, as they generalize the Matérn model to abstract point clouds beyond Euclidean

or manifold settings. In addition, graph Matérn fields are GMRFs and therefore numerical linear algebra techniques can be applied to gain speed up by exploiting sparsity.

We have illustrated through numerical examples the application of graph Matérn fields in Bayesian inverse problems, spatial statistics and graph-based machine learning, bringing these fields together and transferring ideas among them. The graph Matérn models can be directly implemented on the given point cloud, without any additional pre-processing such as adding nodes for triangulation for FEM methods. We demonstrate through comparisons certain benefits of the nonstationary models, where in particular nonstationarity improves classification accuracy. We believe that adding nonstationarity for graph-based learning problems had not been considered before and we hope that our empirical findings will stimulate further research in this area.

The nonstationarity introduced through  $\tau$  had been well-studied in Euclidean settings, while comparatively less has been said about  $\kappa$ . We hope to investigate its modeling effects beyond the role as a length scale described in Remark 2.2.2 and to consider graph representations of anisotropic models where the Laplacian is replaced by  $\nabla \cdot (H(x)\nabla)$ . The case where  $H$  is a constant positive definite matrix can be easily dealt with by introducing a coordinate transformation by  $H^{-1/2}$ . However, the general case where  $H$  is a function of the spatial variable is more involved and further research is needed.

Another direction for further research is to investigate in more detail the case where the points are distributed according to a non-uniform density. As noted in Remark 2.3.3, one can normalize the weights to remove the effects of the density, aiming at recovering the Laplacian. A more interesting question is whether the density can be incorporated as part of the continuum operator that will lead to meaningful Matérn type field. Especially for machine learning applications, we wonder if the density of the point cloud can be exploited in the construction of prior distributions.



## 2.7 Appendix

In this Appendix we prove the main theorems in Section 2.4. Section 2.7.1 makes precise our setting and assumptions, defines several quantities of interest and presents other necessary preliminaries. Sections 2.7.2 and 2.7.3 contain proofs for the spectral convergence of  $L_n^{\tau,\kappa}$  towards  $\mathcal{L}^{\tau,\kappa}$ . Finally, Section 2.7.4 gives the proof of Theorem 2.4.2.

### 2.7.1 Preliminaries

Suppose  $\mathcal{M}$  is an  $m$ -dimensional smooth, connected, compact manifold without boundary embedded in  $\mathbb{R}^d$ , with the absolute value of sectional curvature bounded by  $K$  and Riemannian metric inherited from  $\mathbb{R}^d$ . Let  $\{x_n\}_{n=1}^\infty$  be a sequence of independent samples from the uniform distribution  $\gamma$  on  $\mathcal{M}$ , i.e. the normalized volume measure. Denote  $\mathcal{M}_n := \{x_1, \dots, x_n\}$  and let  $\gamma_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$  be the empirical distribution of the point cloud  $\mathcal{M}_n$ . Throughout  $\tau$  and  $\kappa$  will denote Lipschitz continuous functions on  $\mathcal{M}$  with  $\kappa$  being continuously differentiable. We assume that both functions are bounded from below by positive constants, so there exist  $\alpha, \beta > 0$  with  $\frac{1}{\beta} \leq \tau \leq \beta$ ,  $\frac{1}{\alpha} \leq \kappa \leq \alpha$ . We will analyze the following operators

$$\mathcal{L}^{\tau,\kappa} := \tau I - \nabla \cdot (\kappa \nabla),$$

$$L_n^{\tau,\kappa} := \tau_n + \Delta_n^\kappa,$$

where  $\tau_n := \text{diag}(\tau(x_1), \dots, \tau(x_n))$  and differentiation is defined on the manifold (see e.g. Nicolaescu [2020]). Here  $\Delta_n^\kappa = D - W \in \mathbb{R}^{n \times n}$  and the entries of  $D$  and  $W$  are given by

$$W_{ij} := \frac{2(m+2)}{n\nu_m h_n^{m+2}} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} \sqrt{\kappa(x_i)\kappa(x_j)},$$

$$D_{ii} := \sum_{j=1}^n W_{ij},$$

where  $\nu_m$  is the volume of the  $m$ -dimensional unit ball and  $d_{\mathcal{M}}$  is the geodesic distance on  $\mathcal{M}$ .

**Remark 2.7.1.** *In this appendix we consider the weights  $W_{ij}$  to be defined through the geodesic, rather than the Euclidean distance in (2.11). Since in small neighborhoods both distances agree up to a correction term that is of a higher order than our interest García Trillos et al. [2019a], our results would also hold for the weights in (2.11). We choose to work with the geodesic distance to streamline our presentation.*

As discussed in Section 2.4, the scaling of  $h_n$  will be chosen so that

$$\frac{(\log n)^{c_m}}{n^{1/m}} \ll h_n \ll \frac{1}{n^{1/s}}, \quad (2.36)$$

where  $c_m = 3/4$  if  $m = 2$  and  $c_m = 1/m$  otherwise. We recall that the scaling of  $h_n$  in (2.36) implies that the  $\infty$ -OT distance between  $\gamma_n$  and  $\gamma$  satisfies  $\varepsilon_n = d_\infty(\gamma_n, \gamma) \ll h_n$ . In what follows we assume that we are in a realization where the conclusion of Proposition 2.4.1 holds and we let  $\{T_n\}_{n=1}^\infty$  be a sequence of transport maps satisfying the bound (2.21).

We will use the following inner products and induced norms on continuum and discrete

spaces

$$\begin{aligned}\langle f, g \rangle_{L^2} &:= \int f(x)g(x)d\gamma(x), & \langle f, g \rangle_\tau &:= \int f(x)g(x)\tau(x)d\gamma(x), \\ \langle f, g \rangle_\kappa &:= \int f(x)g(x)\kappa(x)d\gamma(x), & \langle v, w \rangle_2 &:= \frac{1}{n} \sum_{i=1}^n v(x_i)w(x_i), \\ \langle v, w \rangle_{\tau_n} &:= \frac{1}{n} \sum_{i=1}^n v(x_i)w(x_i)\tau(x_i), & \langle v, w \rangle_{\kappa_n} &:= \frac{1}{n} \sum_{i=1}^n v(x_i)w(x_i)\kappa(x_i).\end{aligned}$$

Notice that  $\mathcal{L}^{\tau, \kappa} : \mathcal{D}(\mathcal{L}^{\tau, \kappa}) \subset L^2(\gamma) \rightarrow L^2(\gamma)$ , where  $\mathcal{D}(\mathcal{L}^{\tau, \kappa})$  is the domain of definition, is self-adjoint with respect to the  $\langle \cdot, \cdot \rangle_{L^2}$  inner-product and has a compact resolvent; we will denote by  $\{\lambda^{(k)}\}_{k=1}^\infty$  and  $\{\psi^{(k)}\}_{k=1}^\infty$  its eigenvalues and eigenfunctions and recall that from standard theory

$$\mathcal{D}(\mathcal{L}^{\tau, \kappa}) = \left\{ f \in L^2(\gamma) : \sum_{k=1}^\infty \left[ \lambda^{(k)} \right]^2 \langle f, \psi^{(k)} \rangle_{L^2}^2 < \infty \right\}.$$

We refer to [Nicolaeescu, 2020, Section 10.4.2] for more details. Similarly,  $L_n^{\tau, \kappa}$  is self-adjoint with respect to  $\langle \cdot, \cdot \rangle_2$ . By the minimax principle we can characterize the  $k$ -th smallest eigenvalues of  $\mathcal{L}^{\tau, \kappa}$  and  $L_n^{\tau, \kappa}$  by

$$\begin{aligned}\lambda^{(k)} &= \min_{\mathcal{V} : \dim(\mathcal{V})=k} \max_{f \in \mathcal{V} \setminus \{0\}} \frac{\langle f, \mathcal{L}^{\tau, \kappa} f \rangle_{L^2}}{\langle f, f \rangle_{L^2}}, \\ \lambda_n^{(k)} &= \min_{V : \dim(V)=k} \max_{v \in V \setminus \{0\}} \frac{\langle v, L_n^{\tau, \kappa} v \rangle_2}{\langle v, v \rangle_2}.\end{aligned}$$

We define the continuum and the discrete Dirichlet energies

$$\begin{aligned}D[f] &:= \langle f, \mathcal{L}^{\tau, \kappa} f \rangle_{L^2} = \int \tau(x)f(x)^2 + \int \kappa(x)|\nabla f|^2 = \|f\|_\tau^2 + \|\nabla f\|_\kappa^2 =: D^0[f] + D^1[f], \\ D_{h_n}[v] &:= \langle v, L_n^{\tau, \kappa} v \rangle_2 = \frac{1}{n} \sum_{i=1}^n \tau(x_i)v(x_i)^2 + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n W_{ij} |v(x_i) - v(x_j)|^2 =: D_{h_n}^0[v] + D_{h_n}^1[v].\end{aligned}$$

For the lemmas and theorems below, we shall denote by  $C_{\mathcal{M}}, C_{\mathcal{M},\tau,\kappa}$  etc. constants that depend on the corresponding subscripts.

### 2.7.2 Convergence of Spectrum with Rate

In this subsection we prove Theorem 2.4.6 by establishing a lower and an upper bound on the eigenvalues of  $L_n^{\tau,\kappa}$  in terms of those of  $\mathcal{L}^{\tau,\kappa}$ . With the above definitions of the Dirichlet energies, the eigenvalues have the characterizations

$$\lambda^{(k)} = \min_{\mathcal{V}:\dim(\mathcal{V})=k} \max_{f \in \mathcal{V} \setminus 0} \frac{D[f]}{\|f\|_{L^2}^2}, \quad \lambda_n^{(k)} = \min_{V:\dim(V)=k} \max_{v \in V \setminus 0} \frac{D_{h_n}[v]}{\|v\|_2^2}.$$

In order to compare the Dirichlet energies  $D$  and  $D_{h_n}$ , we need an intermediate quantity defined by

$$E_r[f] := \int_{\mathcal{M}} \int_{B_r(x)} |f(x) - f(y)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x).$$

Notice that  $D_{h_n}^1$  can be seen as a finite sample approximation of  $E_{h_n}$  up to a multiplicative constant. The following lemma, which can be proved with the same argument as [García Trillos et al., 2019a, Lemma 5], connects  $E_{h_n}$  with  $D^1$ .

**Lemma 2.7.2.** *For  $f \in L^2(\gamma)$  and  $r < 2h_n$ ,*

$$E_r[f] \leq (1 + C_{\mathcal{M},\kappa} h_n) \frac{\nu_m r^{m+2}}{m+2} D^1[f].$$

#### 2.7.2.1 Upper Bound

To start with, define the projection map  $P : L^2(\gamma) \rightarrow L^2(\gamma_n)$  by

$$Pf(x_i) := n \int_{U_i} f(x) d\gamma(x),$$

where  $U_i = T_n^{-1}(\{x_i\})$  and  $\{T_n\}_{n=1}^\infty$  is a sequence of transportation maps as in Proposition 2.4.1. The  $U_i$ 's are called transportation cells.

**Lemma 2.7.3** (Discrete Dirichlet Energy Upper Bound). *Let  $f \in H^1(\gamma)$ .*

1.  $\left| \|Pf\|_2 - \|f\|_{L^2} \right| \leq C_{\mathcal{M}} \varepsilon_n \|\nabla f\|_{L^2}$ .
2.  $D_{h_n}[Pf] \leq \left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} \right) \right] D[f]$ .

*Proof.* The first statement is proved in [Burago et al., 2015, Lemma 4.3(1)]. The second statement will be proved by combining upper bounds for  $D_{h_n}^0[Pf]$  and  $D_{h_n}^1[Pf]$ . First, by Hölder's inequality and the fact that  $\gamma(U_i) = \gamma_n(\{x_i\}) = 1/n$ ,

$$\begin{aligned}
D_{h_n}^0[Pf] &= \frac{1}{n} \sum_{i=1}^n \tau(x_i) Pf(x_i)^2 = \frac{1}{n} \sum_{i=1}^n \tau(x_i) n^2 \left| \int_{U_i} f(x) d\gamma(x) \right|^2 \\
&\leq \sum_{i=1}^n \tau(x_i) \int_{U_i} f(x)^2 d\gamma(x) \\
&\leq [1 + \text{Lip}(\tau) \beta \varepsilon_n] \sum_{i=1}^n \int_{U_i} \tau(x) f(x)^2 d\gamma(x) \\
&= [1 + \text{Lip}(\tau) \beta \varepsilon_n] D^0[f].
\end{aligned} \tag{2.37}$$

For the upper bound on  $D_{h_n}^1[Pf]$ , observe that

$$Pf(x_i) - Pf(x_j) = n^2 \int_{U_i} \int_{U_j} f(y) - f(x) d\gamma(y) d\gamma(x),$$

which implies

$$|Pf(x_i) - Pf(x_j)|^2 \leq n^2 \int_{U_i} \int_{U_j} |f(y) - f(x)|^2 d\gamma(y) d\gamma(x).$$

By Lipschitz continuity of  $\kappa$ ,

$$\begin{aligned}
D_{h_n}^1[Pf] &= \frac{m+2}{n^2\nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \sqrt{\kappa(x_i)\kappa(x_j)} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} |Pf(x_i) - Pf(x_j)|^2 \\
&\leq \frac{m+2}{\nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \sqrt{\kappa(x_i)\kappa(x_j)} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} \int_{U_i} \int_{U_j} |f(y) - f(x)|^2 d\gamma(y) d\gamma(x) \\
&\leq [1 + \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} \\
&\quad \cdot \int_{U_i} \int_{U_j} |f(y) - f(x)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x) \\
&\leq [1 + \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \int_{\mathcal{M}} \int_{V(x)} |f(y) - f(x)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x),
\end{aligned}$$

where if  $x \in U_i$  then  $V(x) = \bigcup_{j:j \sim i} U_j$ , with  $j \sim i$  meaning  $d_{\mathcal{M}}(x_i, x_j) < h_n$ . In the second to last step we have used that

$$\begin{aligned}
\left| \sqrt{\kappa(x_i)\kappa(x_j)} - \sqrt{\kappa(x)\kappa(y)} \right| &\leq \left| \sqrt{\kappa(x_i)\kappa(x_j)} - \sqrt{\kappa(x_i)\kappa(y)} \right| + \left| \sqrt{\kappa(x_i)\kappa(y)} - \sqrt{\kappa(x)\kappa(y)} \right| \\
&\leq \sqrt{\alpha} \frac{|\kappa(x_j) - \kappa(y)|}{\sqrt{\kappa(x_j)} + \sqrt{\kappa(y)}} + \sqrt{\alpha} \frac{|\kappa(x_i) - \kappa(x)|}{\sqrt{\kappa(x_i)} + \sqrt{\kappa(x)}} \leq \text{Lip}(\kappa)\alpha\varepsilon_n.
\end{aligned}$$

Notice that  $V(x) \subset B_{h_n+2\varepsilon_n}(x)$  and hence

$$\begin{aligned}
D_{h_n}^1[Pf] &\leq [1 + \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \int_{\mathcal{M}} \int_{B_{h_n+2\varepsilon_n}(x)} |f(y) - f(x)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x) \\
&= [1 + \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} E_{h_n+2\varepsilon_n}[f] \\
&\leq \left[ 1 + \text{Lip}(\kappa)\alpha\varepsilon_n \right] \left[ 1 + C_{\mathcal{M},\kappa}(h_n + 2\varepsilon_n) \right] \left( \frac{h_n + 2\varepsilon_n}{h_n} \right)^{m+2} D^1[f] \\
&\leq \left[ 1 + C_{\mathcal{M},\kappa} \left( h_n + \frac{\varepsilon_n}{h_n} \right) \right] D^1[f], \tag{2.38}
\end{aligned}$$

where we have used Lemma 2.7.2 and the assumption that  $\varepsilon_n \ll h_n$ . The result follows by combining (2.37) and (2.38).  $\square$

**Corollary 2.7.4** (Upper Bound). *Suppose  $k := k_n$  is such that  $\varepsilon_n \sqrt{\lambda^{(k)}} \ll 1$  for  $n$  large.*

*Then*

$$\lambda_n^{(k)} \leq \left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} + \sqrt{\lambda^{(k)}} \varepsilon_n \right) \right] \lambda^{(k)}.$$

*Proof.* Let  $\mathcal{V}$  be the span of eigenfunctions  $f_1, \dots, f_k$  of  $\mathcal{L}^{\tau, \kappa}$  associated with eigenvalues  $\lambda^{(1)}, \dots, \lambda^{(k)}$ . For  $f \in \mathcal{V}$ , we have

$$\|\nabla f\|_{L^2} \leq \sqrt{\alpha} \|\nabla f\|_{\kappa} \leq \sqrt{\alpha} \sqrt{D[f]} \leq \sqrt{\alpha \lambda^{(k)}} \|f\|_{L^2}.$$

Lemma 2.7.3 (1) then implies that

$$\|Pf\|_2 \geq \|f\|_{L^2} - C_{\mathcal{M}} \varepsilon_n \|\nabla f\|_{L^2} \geq \|f\|_{L^2} - C_{\mathcal{M}, \kappa} \varepsilon_n \sqrt{\lambda^{(k)}} \|f\|_{L^2}.$$

Therefore, the assumption that  $\varepsilon_n \sqrt{\lambda^{(k)}} \ll 1$  implies that  $P|_{\mathcal{V}}$  is injective and  $V = P(\mathcal{V})$  has dimension  $k$ . By Lemma 2.7.3 (2) we have

$$\begin{aligned} \lambda_n^{(k)} &\leq \max_{v \in V \setminus 0} \frac{Dh_n[v]}{\|v\|_2^2} = \max_{f \in \mathcal{V} \setminus 0} \frac{Dh_n[Pf]}{\|Pf\|_2^2} \leq \max_{f \in \mathcal{V} \setminus 0} \frac{\left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} \right) \right] D[f]}{\left( 1 - C_{\mathcal{M}, \tau, \kappa} \varepsilon_n \sqrt{\lambda^{(k)}} \right) \|f\|_{L^2}^2} \\ &= \frac{\left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} \right) \right] \lambda^{(k)}}{1 - C_{\mathcal{M}, \tau, \kappa} \varepsilon_n \sqrt{\lambda^{(k)}}} \\ &\leq \left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} + \sqrt{\lambda^{(k)}} \varepsilon_n \right) \right] \lambda^{(k)}. \end{aligned}$$

□

### 2.7.2.2 Lower Bound

Define  $P^* : L^2(\gamma_n) \rightarrow L^2(\gamma)$  by

$$P^*v := \sum_{i=1}^n v(x_i)\mathbf{1}_{U_i},$$

where  $U_i$ 's are the transportation cells as in the definition of  $P$ . We note that  $P^*$  defines a piecewise constant interpolation map; for the subsequent analysis we need to introduce a smoothing operator  $\Lambda$  so that the map  $\mathcal{I} := \Lambda \circ P^*$  satisfies  $\mathcal{I}v \in H^1(\gamma)$ , the Sobolev space of order 1. We now detail the construction of the smoothing operator. Let

$$\psi(t) := \begin{cases} \frac{m+2}{2\nu_m}(1-t^2), & 0 \leq t \leq 1, \\ 0, & t > 1. \end{cases}$$

Consider for  $0 < r < 2h_n$  the kernel

$$k_r(x, y) := r^{-m}\psi\left(\frac{d_{\mathcal{M}}(x, y)}{r}\right)$$

and the associated integral operator

$$\Lambda_r^0 f = \int_{\mathcal{M}} k_r(x, y)f(y)d\gamma(y).$$

Let  $\theta(x) := \Lambda_r^0 \mathbf{1}_{\mathcal{M}} = \int_{\mathcal{M}} k_r(x, y)d\gamma(y)$  and then define

$$\Lambda_r f := \theta^{-1}\Lambda_r^0 f,$$



so that  $\Lambda_r$  preserves constant functions. Finally we define the interpolation operator  $\mathcal{I} : L^2(\gamma_n) \rightarrow L^2(\gamma)$  by

$$\mathcal{I}v := \Lambda_{h_n - 2\varepsilon_n} P^* v.$$

We next present some auxiliary bounds that will be needed later.

**Lemma 2.7.5** (Auxiliary Bounds). *For  $f \in L^2(\gamma)$ , we have*

$$\|\Lambda_r f\|_\tau^2 \leq [1 + \text{Lip}(\tau)r\beta][1 + CmKr^2]^2 \|f\|_\tau^2, \quad (2.39)$$

$$\|\Lambda_r f - f\|_\kappa^2 \leq [1 + \text{Lip}(\kappa)\alpha r] \frac{Cm}{\nu_m r^m} E_r[f], \quad (2.40)$$

$$\|\nabla(\Lambda_r f)\|_\kappa^2 \leq [1 + \text{Lip}(\kappa)\alpha r][1 + Cm^2Kr^2] \frac{m+2}{\nu_m r^{m+2}} E_r[f]. \quad (2.41)$$

*Proof.* The above results are proved in the same way as in [Burago et al., 2015, Lemma 5.3, 5.4, 5.5] with little adjustments and the main differences are the additional factors  $1 + \text{Lip}(\tau)\beta r$  or  $1 + \text{Lip}(\kappa)\alpha r$ . To illustrate the idea, we will prove (2.39) and the generalizations for (2.40) and (2.41) are similar. First by [Burago et al., 2015, Lemma 5.1], we have for each  $x \in \mathcal{M}$ ,

$$(1 + CmKr^2)^{-1} \leq \theta(x) \leq (1 + CmKr^2), \quad (2.42)$$

where recall that  $K$  is an upper bound on the absolute value of the sectional curvature. Then we have

$$\begin{aligned} |\Lambda_r f|^2 &= \theta^{-2} \left| \int_{\mathcal{M}} k_r(x, y) f(y) \right|^2 \leq \theta^{-2} \int_{\mathcal{M}} k_r(x, y) d\gamma(y) \int_{\mathcal{M}} k_r(x, y) |f(y)|^2 d\gamma(y) \\ &= \theta^{-1} \int_{\mathcal{M}} k_r(x, y) |f(y)|^2 d\gamma(y) \\ &\leq [1 + CmKr^2] \int_{\mathcal{M}} k_r(x, y) |f(y)|^2 d\gamma(y). \end{aligned}$$

Noticing that  $k_r(x, y)$  is zero when  $d_{\mathcal{M}}(x, y) > r$  and  $\tau$  is Lipschitz, we have

$$\begin{aligned}
\|\Lambda_r f\|_{\tau}^2 &\leq [1 + CmKr^2] \int_{\mathcal{M}} \int_{\mathcal{M}} k_r(x, y) |f(y)|^2 \tau(x) d\gamma(x) d\gamma(y) \\
&= [1 + CmKr^2] \int_{\mathcal{M}} |f(y)|^2 \left[ \int_{B_r(y)} k_r(x, y) \tau(x) d\gamma(x) \right] d\gamma(y) \\
&\leq [1 + \text{Lip}(\tau)r\beta][1 + CmKr^2] \int_{\mathcal{M}} |f(y)|^2 \left[ \int_{B_r(y)} k_r(x, y) \tau(y) d\gamma(x) \right] d\gamma(y) \\
&= [1 + \text{Lip}(\tau)r\beta][1 + CmKr^2] \int_{B_r(y)} k_r(x, y) d\gamma(x) \int_{\mathcal{M}} |f(y)|^2 \tau(y) d\gamma(y) \\
&= [1 + \text{Lip}(\tau)r\beta][1 + CmKr^2]^2 \|f\|_{\tau}^2.
\end{aligned}$$

□

**Lemma 2.7.6** (Discrete Dirichlet Energy Lower Bound). *For each  $v \in L^2(\gamma_n)$ ,*

1.  $\left| \|\mathcal{I}v\|_{L^2} - \|v\|_2 \right| \leq C_{\kappa} h_n \sqrt{D_{h_n}[v]} .$
2.  $D[\mathcal{I}v] \leq \left[ 1 + C_{\mathcal{M}, \tau, \kappa} \left( h_n + \frac{\varepsilon_n}{h_n} \right) \right] D_{h_n}[v].$

*Proof.* 1. By equation (6.4) in the proof of [Burago et al., 2015, Lemma 6.2(1)],

$$\|\mathcal{I}v - P^*v\|_{L^2} \leq Ch_n \|\delta v\|, \quad (2.43)$$

where

$$\|\delta v\|^2 = \frac{m+2}{\nu_m n^2 h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} |v(x_i) - v(x_j)|^2. \quad (2.44)$$

The result follows by noticing that  $\|\delta v\| \leq \sqrt{\alpha D_{h_n}[v]}$  and

$$\left| \|\mathcal{I}v\|_{L^2} - \|v\|_2 \right| = \left| \|\mathcal{I}v\|_{L^2} - \|P^*v\|_{L^2} \right| \leq \|\mathcal{I}v - P^*v\|_{L^2}.$$

2. The second statement will be proved by combining lower bounds for  $D_{h_n}^0[\mathcal{I}v]$  and  $D_{h_n}^1[\mathcal{I}v]$ . For the lower bound on  $D_{h_n}^0[\mathcal{I}v]$ , we have by (2.39),

$$\begin{aligned} D^0[\mathcal{I}v] &= \|\Lambda_{h_n-2\varepsilon_n} P^*v(x)\|_\tau^2 \\ &\leq \left[1 + \text{Lip}(\tau)\beta(h_n - 2\varepsilon_n)\right] \left[1 + CmK(h_n - 2\varepsilon_n)^2\right]^2 \int \tau(x)|P^*v(x)|^2 d\gamma(x). \end{aligned}$$

We also have

$$\begin{aligned} \int \tau(x)|P^*v(x)|^2 d\gamma(x) &= \sum_{i=1}^n \int_{U_i} \tau(x)|v(x_i)|^2 d\gamma(x) \\ &\leq [1 + \text{Lip}(\tau)\beta\varepsilon_n] \sum_{i=1}^n \int_{U_i} \tau(x_i)v(x_i)^2 d\gamma(x) = [1 + \text{Lip}(\tau)\beta\varepsilon_n] D_{h_n}^0[v]. \end{aligned}$$

Therefore

$$D^0[\mathcal{I}v] \leq (1 + C_{\mathcal{M},\tau}h_n) D_{h_n}^0[v]. \quad (2.45)$$

Next we seek a lower bound for  $D_{h_n}^1[\mathcal{I}v]$ . We have

$$\begin{aligned}
D_{h_n}^1[v] &= \frac{m+2}{n^2 \nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \sqrt{\kappa(x_i) \kappa(x_j)} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} |v(x_i) - v(x_j)|^2 \\
&= \frac{m+2}{n^2 \nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \sqrt{\kappa(x_i) \kappa(x_j)} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} |P^*v(x_i) - P^*v(x_j)|^2 \\
&= \frac{m+2}{\nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \sqrt{\kappa(x_i) \kappa(x_j)} \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} \\
&\quad \cdot \int_{U_i} \int_{U_j} |P^*v(x) - P^*v(y)|^2 d\gamma(y) d\gamma(x) \\
&\geq [1 - \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{d_{\mathcal{M}}(x_i, x_j) < h_n\} \\
&\quad \cdot \int_{U_i} \int_{U_j} |P^*v(x) - P^*v(y)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x) \\
&= [1 - \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \int_{\mathcal{M}} \int_{V(x)} |P^*v(x) - P^*v(y)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x),
\end{aligned}$$

where if  $x \in U_i$  then  $V(x) = \bigcup_{j:j \sim i} U_j$ . Notice that  $V(x) \supset B_{h_n - 2\varepsilon_n}(x)$  and hence,

$$\begin{aligned}
D_{h_n}^1[v] &\geq [1 - \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} \int_{\mathcal{M}} \int_{B_{h_n - 2\varepsilon_n}} |P^*v(x) - P^*v(y)|^2 \sqrt{\kappa(x)\kappa(y)} d\gamma(y) d\gamma(x) \\
&= [1 - \text{Lip}(\kappa)\alpha\varepsilon_n] \frac{m+2}{\nu_m h_n^{m+2}} E_{h_n - 2\varepsilon_n}[P^*v]. \tag{2.46}
\end{aligned}$$

Combining inequality (2.41) with (2.46) gives

$$\begin{aligned}
& D^1[\mathcal{I}v] \\
&= \|\nabla(\mathcal{I}v)\|_{\kappa}^2 \\
&= \|\nabla(\Lambda_{h_n-2\epsilon_n} P^* v)\|_{\kappa}^2 \\
&\leq [1 + \text{Lip}(\kappa)\alpha(h_n - 2\epsilon_n)] \left[1 + Cm^2K(h_n - 2\epsilon_n)^2\right]^2 \frac{m+2}{\nu_m(h_n - \epsilon_n)^{m+2}} E_{h_n-2\epsilon_n}[P^* v] \\
&\leq [1 + \text{Lip}(\kappa)\alpha(h_n - 2\epsilon_n)] \left[1 + Cm^2K(h_n - 2\epsilon_n)^2\right]^2 \left(\frac{h_n}{h_n - 2\epsilon_n}\right)^{m+2} [1 - \text{Lip}(\kappa)\alpha\epsilon_n]^{-1} D_h^1[v] \\
&\leq \left[1 + C_{\mathcal{M},\kappa} \left(h_n + \frac{\epsilon_n}{h_n}\right)\right] D_h^1[v], \tag{2.47}
\end{aligned}$$

where we have used Lemma 2.7.2 and the fact that  $\epsilon_n \ll h_n$ . The second statement follows by combining (2.45) and (2.47).  $\square$

**Corollary 2.7.7** (Lower Bound). *Suppose  $k := k_n$  is such that  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$  for  $n$  large.*

*Then*

$$\lambda_n^{(k)} \geq \left[1 - C_{\mathcal{M},\tau,\kappa} \left(\frac{\epsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right)\right] \lambda^{(k)}.$$

*Proof.* Since we are interested in proving lower bounds for  $\lambda_n^{(k)}$ , the result is trivial if  $\lambda_n^{(k)} \geq \lambda^{(k)}$ . Therefore we shall assume that  $\lambda_n^{(k)} < \lambda^{(k)}$ . Let  $V$  be the span of eigenvectors  $v_1, \dots, v_k$  of  $L_n^{\tau,\kappa}$  associated with eigenvalues  $\lambda_n^{(1)}, \dots, \lambda_n^{(k)}$ . Lemma 2.7.6(1) implies for  $v \in V$

$$\|\mathcal{I}v\|_{L^2} \geq \|v\|_2 - C_{\mathcal{M},\kappa} h_n \sqrt{D_{h_n}[v]} \geq \left[1 - C_{\mathcal{M},\kappa} h_n \sqrt{\lambda_n^{(k)}}\right] \|v\|_2 \geq \left[1 - C_{\mathcal{M},\kappa} h_n \sqrt{\lambda^{(k)}}\right] \|v\|_2.$$

Therefore, the assumption that  $h_n \sqrt{\lambda^{(k)}} \ll 1$  implies that  $\mathcal{I}|_V$  is injective and  $\mathcal{V} = \mathcal{I}(V)$

has dimension  $k$ . Lemma 2.7.6 then gives

$$\begin{aligned}
\lambda^{(k)} &\leq \max_{f \in \mathcal{V} \setminus 0} \frac{D[f]}{\|f\|_{L^2}^2} = \max_{v \in V \setminus 0} \frac{D[\mathcal{I}v]}{\|\mathcal{I}v\|_{L^2}^2} \leq \max_{v \in V \setminus 0} \frac{\left[1 + C_{\mathcal{M}, \tau, \kappa} \left(h_n + \frac{\varepsilon_n}{h_n}\right)\right] D_{h_n}[u]}{\left(1 - C_{\mathcal{M}, \kappa} h_n \sqrt{\lambda^{(k)}}\right)^2 \|u\|_2} \\
&= \max_{v \in V \setminus 0} \frac{\left[1 + C_{\mathcal{M}, \tau, \kappa} \left(h_n + \frac{\varepsilon_n}{h_n}\right)\right]}{\left(1 - C_{\mathcal{M}, \kappa} h_n \sqrt{\lambda^{(k)}}\right)^2} \lambda_n^{(k)} \\
&\leq \left[1 + C_{\mathcal{M}, \tau, \kappa} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right)\right] \lambda_n^{(k)}.
\end{aligned}$$

Therefore

$$\lambda_n^{(k)} \geq \left[1 + C_{\mathcal{M}, \tau, \kappa} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right)\right]^{-1} \lambda^{(k)} \geq \left[1 - C_{\mathcal{M}, \tau, \kappa} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right)\right] \lambda^{(k)}.$$

□

Combining Lemma 2.7.3 and 2.7.6 we have:

**Theorem 2.7.8.** *Suppose  $k := k_n$  is such that  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$  for  $n$  large. Then*

$$\frac{|\lambda_n^{(k)} - \lambda^{(k)}|}{\lambda^{(k)}} \leq C_{\mathcal{M}, \tau, \kappa} \left[\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right],$$

where  $C_{\mathcal{M}, \tau, \kappa}$  is a constant depending on  $\mathcal{M}, \tau, \kappa$ .

### 2.7.3 Convergence of Eigenfunctions with Rate

In this subsection we prove Theorem 2.4.7. Before we proceed, we introduce some additional notations. For any interval  $J$  of  $\mathbb{R}$ , denote  $H_J(\gamma)$  the subspace of  $H^1(\gamma)$  that is spanned by eigenfunctions of  $\mathcal{L}^{\tau, \kappa}$  associated with eigenvalues in  $J$  and  $\mathbb{P}_J(\gamma)$  the orthogonal projection from  $L^2(\gamma)$  onto  $H_J(\gamma)$ . Similarly we use the notation  $H_J(\gamma_n)$  and  $\mathbb{P}_J(\gamma_n)$  for  $L_n^{\tau, \kappa}$ . To ease notation, we will denote  $H_{(-\infty, \lambda)}(\gamma)$  and  $H_{(-\infty, \lambda)}(\gamma_n)$  as  $H_\lambda(\gamma)$  and  $H_\lambda(\gamma_n)$ , respectively.

We shall also denote both projections as  $\mathbb{P}_J$  when no confusion arises.

To start with, we need several auxiliary results.

**Lemma 2.7.9.** *Let  $v \in H_\lambda(\gamma_n)$ .*

1.  $\|P\mathcal{I}v - v\|_2 \leq C_\kappa h_n \sqrt{D_{h_n}[v]}$ .
2.  $D[\mathcal{I}v] \geq \left[1 - C_{\mathcal{M},\tau,\kappa} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda}\right)\right] D_{h_n}[v]$ .

*Proof.* 1. By Burago et al. [2015][Lemma 6.4(2)],

$$\|P\mathcal{I}v - v\|_2 \leq Ch_n \|\delta v\|, \quad (2.48)$$

where  $\|\delta v\|$  is defined in (2.44) The result follows by noticing that  $\|\delta v\| \leq \sqrt{\alpha D_{h_n}[v]}$ .

2. We first bound  $D_{h_n}[P\mathcal{I}v]$  in terms of  $D_{h_n}[v]$ . Denoting  $P\mathcal{I}v$  as  $w$ , we have

$$\begin{aligned} D_{h_n}[w] &= \langle w, L_n^{\tau,\kappa} w \rangle_2 \\ &= \langle w - \mathbb{P}_\lambda w + \mathbb{P}_\lambda w, L_n^{\tau,\kappa} (w - \mathbb{P}_\lambda w + \mathbb{P}_\lambda w) \rangle_2 \\ &= \langle w - \mathbb{P}_\lambda w, L_n^{\tau,\kappa} (w - \mathbb{P}_\lambda w) \rangle_2 + \langle \mathbb{P}_\lambda w, L_n^{\tau,\kappa} \mathbb{P}_\lambda w \rangle_2 + \langle w - \mathbb{P}_\lambda w, L_n^{\tau,\kappa} \mathbb{P}_\lambda w \rangle_2 \\ &= \langle w - \mathbb{P}_\lambda w, L_n^{\tau,\kappa} (w - \mathbb{P}_\lambda w) \rangle_2 + \langle \mathbb{P}_\lambda w, L_n^{\tau,\kappa} \mathbb{P}_\lambda w \rangle_2 \\ &\geq \langle \mathbb{P}_\lambda w, L_n^{\tau,\kappa} \mathbb{P}_\lambda w \rangle_2, \end{aligned}$$

where we have used that  $\mathbb{P}_\lambda w$  and  $L_n^{\tau,\kappa} \mathbb{P}_\lambda w \in H_\lambda(\gamma_n)$  are orthogonal to  $w - \mathbb{P}_\lambda w$ . Since  $L_n^{\tau,\kappa}$  is nonsingular,  $\langle \cdot, L_n^{\tau,\kappa} \cdot \rangle_2$  defines an inner product and the triangle inequality implies

$$\sqrt{D_{h_n}[w]} \geq \sqrt{\langle \mathbb{P}_\lambda w, L_n^{\tau,\kappa} \mathbb{P}_\lambda w \rangle_2} \geq \sqrt{\langle v, L_n^{\tau,\kappa} v \rangle_2} - \sqrt{\langle v - \mathbb{P}_\lambda w, L_n^{\tau,\kappa} (v - \mathbb{P}_\lambda w) \rangle_2}.$$

Now we bound the second term above. Since  $v \in H_\lambda$ , we have  $v = \mathbb{P}_\lambda v$  and

$$\begin{aligned} \langle v - \mathbb{P}_\lambda w, L_n^{\tau, \kappa}(v - \mathbb{P}_\lambda w) \rangle_2 &= \langle \mathbb{P}_\lambda(v - w), L_n^{\tau, \kappa} \mathbb{P}_\lambda(v - w) \rangle_2 \\ &\leq \lambda \|\mathbb{P}_\lambda(v - w)\|_2^2 \leq \lambda \|v - w\|_2^2 \leq C_\kappa \lambda h_n^2 D_{h_n}[v], \end{aligned}$$

where the last step follows from (2.48). Hence

$$\sqrt{D_{h_n}[P\mathcal{I}v]} \geq \left[1 - C_\kappa \sqrt{\lambda h_n}\right] \sqrt{D_{h_n}[v]},$$

and the result follows from Lemma 2.7.3, which says

$$D_{h_n}[P\mathcal{I}v] \leq \left[1 + C_{\mathcal{M}, \tau, \kappa} \left(h_n + \frac{\varepsilon_n}{h_n}\right)\right] D[\mathcal{I}v].$$

□

We fix orthonormal eigenfunctions  $\{v_k\}_{k=1}^n$  and  $\{f_k\}_{k=1}^\infty$  for  $L_n^{\tau, \kappa}$  and  $\mathcal{L}^{\tau, \kappa}$ . The following lemma bounds the projection error when  $J$  is a half-interval.

**Lemma 2.7.10.** *Suppose  $k := k_n$  is such that  $h_n \sqrt{\lambda^{(k)}} \ll 1$ . Then, for any  $a > 0$ ,*

$$\begin{aligned} \|\mathcal{I}v_k - \mathbb{P}_{\lambda^{(k)}+a} \mathcal{I}v_k\|_{L^2}^2 &\leq C_{\mathcal{M}, \tau, \kappa} a^{-1} k \lambda^{(k)} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right), \\ D[\mathcal{I}v_k - \mathbb{P}_{\lambda^{(k)}+a} \mathcal{I}v_k] &\leq C_{\mathcal{M}, \tau, \kappa} a^{-1} (\lambda^{(k)} + a) k \lambda^{(k)} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right). \end{aligned}$$

*Proof.* Let  $V$  be the span of  $v_1, \dots, v_k$  and  $\mathcal{V} = \mathcal{I}(V)$ . Since  $h_n \sqrt{\lambda^{(k)}} \ll 1$ , Theorem 2.4.6 implies  $\lambda_n^{(k)} \leq C \lambda^{(k)}$  and then by Lemma 2.7.6, for any  $v \in V$ ,

$$\begin{aligned} \|\mathcal{I}v\|_{L^2} &\geq \left[1 - C_{\mathcal{M}, \tau, \kappa} h_n \sqrt{\lambda_n^{(k)}}\right] \|v\|_2 \geq \left[1 - C_{\mathcal{M}, \tau, \kappa} h_n \sqrt{\lambda^{(k)}}\right] \|v\|_2 \\ D[\mathcal{I}v] &\leq \left[1 + C_{\mathcal{M}, \tau, \kappa} \left(h_n + \frac{\varepsilon_n}{h_n}\right)\right] D_{h_n}[v]. \end{aligned}$$



The assumption  $h_n \sqrt{\lambda^{(k)}} \ll 1$  also implies that  $\mathcal{I}|_{\mathcal{V}}$  is injective and  $\mathcal{V}$  is  $k$ -dimensional. Let  $\lambda_{\mathcal{V}}^{(1)}, \dots, \lambda_{\mathcal{V}}^{(k)}$  be the eigenvalues of  $A := \mathcal{L}^{\tau, \kappa}|_{\mathcal{V}}$ . The minimax principle implies that, for  $j \leq k$ ,

$$\lambda_{\mathcal{V}}^{(j)} \leq \frac{\left[1 + C_{\mathcal{M}, \tau, \kappa} \left(h_n + \frac{\varepsilon_n}{h_n}\right)\right]}{\left[1 - C_{\mathcal{M}, \tau, \kappa} h_n \sqrt{\lambda^{(k)}}\right]} \lambda_n^{(j)} \leq \lambda_n^{(j)} + C_{\mathcal{M}, \tau, \kappa} \lambda^{(k)} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right). \quad (2.49)$$

Define another operator  $\tilde{\mathcal{L}}^{\tau, \kappa}$  by

$$\tilde{\mathcal{L}}^{\tau, \kappa} f = \mathcal{L}^{\tau, \kappa} \mathbb{P}_{\lambda^{(k)}+a} f + \lambda^{(k)} (f - \mathbb{P}_{\lambda^{(k)}+a} f).$$

Let  $\{f_i\}$ 's be the eigenvectors of  $\mathcal{L}^{\tau, \kappa}$  associated with eigenvalues  $\{\lambda^{(i)}\}$ . We observe that  $\tilde{\mathcal{L}}^{\tau, \kappa}$  is self-adjoint with respect to the  $L^2(\gamma)$  inner product and shares the same eigenvectors with corresponding eigenvalues  $\lambda^{(1)}, \dots, \lambda^{(k)}, \lambda^{(k)}, \dots$ . Let  $\tilde{A} := \tilde{\mathcal{L}}^{\tau, \kappa}|_{\mathcal{V}}$  and  $\tilde{\lambda}_{\mathcal{V}}^{(1)}, \dots, \tilde{\lambda}_{\mathcal{V}}^{(k)}$  be its eigenvalues. Let  $f \in \mathcal{V}$  and  $g = f - \mathbb{P}_{\lambda^{(k)}+a} f$ . Since  $\mathcal{L}^{\tau, \kappa} \mathbb{P}_{\lambda^{(k)}+a} f = \tilde{\mathcal{L}}^{\tau, \kappa} \mathbb{P}_{\lambda^{(k)}+a} f$ , we have by orthogonality

$$\begin{aligned} \langle f, \mathcal{L}^{\tau, \kappa} f \rangle_{L^2} - \langle f, \tilde{\mathcal{L}}^{\tau, \kappa} f \rangle_{L^2} &= \langle g, \mathcal{L}^{\tau, \kappa} g \rangle_{L^2} - \langle g, \tilde{\mathcal{L}}^{\tau, \kappa} g \rangle_{L^2} \\ &= \langle g, \mathcal{L}^{\tau, \kappa} g \rangle_{L^2} - \lambda^{(k)} \|g\|_{L^2}^2 \geq \frac{a}{\lambda^{(k)} + a} \langle g, \mathcal{L}^{\tau, \kappa} g \rangle_{L^2}, \end{aligned} \quad (2.50)$$

where the last inequality follows from the fact that  $\langle g, \mathcal{L}^{\tau, \kappa} g \rangle_{L^2} \geq (\lambda^{(k)} + a) \|g\|_{L^2}^2$ . By the minimax principle, we have  $\tilde{\lambda}_{\mathcal{V}}^{(j)} \geq \lambda^{(j)}$  for  $j \leq k$  and by Theorem 2.4.6 we have

$$\tilde{\lambda}_{\mathcal{V}}^{(j)} \geq \lambda_n^{(j)} - C_{\mathcal{M}, \tau, \kappa} \lambda^{(k)} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right).$$

Together with (2.49), we get

$$\lambda_{\mathcal{V}}^{(j)} - \tilde{\lambda}_{\mathcal{V}}^{(j)} \leq C_{\mathcal{M}, \tau, \kappa} \lambda^{(k)} \left(\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}}\right),$$

and by [Burago et al., 2015, Lemma 7.2], for any  $f \in \mathcal{V}$ ,

$$\begin{aligned} \langle f, \mathcal{L}^{\tau, \kappa} f \rangle_{L^2} - \langle f, \tilde{\mathcal{L}}^{\tau, \kappa} f \rangle_{L^2} &= \langle f, Af \rangle_{L^2} - \langle f, \tilde{A}f \rangle_{L^2} \\ &\leq k \max_{1 \leq j \leq k} \{ \lambda_{\mathcal{V}}^{(j)} - \tilde{\lambda}_{\mathcal{V}}^{(j)} \} \leq C_{\mathcal{M}, \tau, \kappa} k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right), \end{aligned}$$

where we have used the fact that  $A \geq \tilde{A}$  from (2.50). Hence (2.50) implies

$$\begin{aligned} D[g] &= \langle g, \mathcal{L}^{\tau, \kappa} g \rangle_{L^2} \leq C_{\mathcal{M}, \tau, \kappa} a^{-1} (\lambda^{(k)} + a) k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right), \\ \|g\|_{L^2}^2 &\leq C_{\mathcal{M}, \tau, \kappa} a^{-1} k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right). \end{aligned}$$

□

The next lemma bounds the projection error when  $J$  is a finite interval.

**Lemma 2.7.11.** *Suppose  $k := k_n$  is such that  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$ . Let  $a \leq b \leq c \leq \lambda^{(k)}$  be constants so that the interval  $(\lambda^{(k)} + a, \lambda^{(k)} + b)$  does not contain any eigenvalue of  $\mathcal{L}^{\tau, \kappa}$ . Then*

$$\|\mathcal{I}v_k - \mathbb{P}_{(\lambda^{(k)} - c, \lambda^{(k)} + a)} \mathcal{I}v_k\|_{L^2}^2 \leq C_{\mathcal{M}, \tau, \kappa} c^{-1} b^{-1} k \left[ \lambda^{(k)} \right]^2 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right) + c^{-1} a.$$

*Proof.* Let  $f = \mathcal{I}v_k$  and decompose it as

$$f = \mathbb{P}_{(\lambda^{(k)} - c, \lambda^{(k)} + a)} f + \mathbb{P}_{(-\infty, \lambda^{(k)} - c]} f + \mathbb{P}_{(\lambda^{(k)} + a, \infty)} f =: f_0 + f_- + f_+.$$

Orthogonality implies

$$\langle f, \mathcal{L}^{\tau, \kappa} f \rangle_{L^2} = \langle f_0, \mathcal{L}^{\tau, \kappa} f_0 \rangle_{L^2} + \langle f_-, \mathcal{L}^{\tau, \kappa} f_- \rangle_{L^2} + \langle f_+, \mathcal{L}^{\tau, \kappa} f_+ \rangle_{L^2},$$

and we have by assumption that  $f_+ = \mathbb{P}_{[\lambda^{(k)}+b,\infty)}f$ . By Lemma 2.7.10, we have

$$\begin{aligned} \|f_+\|_{L^2}^2 &\leq C_{\mathcal{M},\tau,\kappa} b^{-1} k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right), \\ \langle f_+, \mathcal{L}^{\tau,\kappa} f_+ \rangle_{L^2} &\leq C_{\mathcal{M},\tau,\kappa} b^{-1} (\lambda^{(k)} + b) k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right). \end{aligned} \quad (2.51)$$

By Lemma 2.7.9 (2), we have

$$\begin{aligned} \langle f, \mathcal{L}^{\tau,\kappa} f \rangle_{L^2} &= D[\mathcal{I}v_k] \geq \left[ 1 - C_{\mathcal{M},\tau,\kappa} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right) \right] D_{h_n}[v_k] \\ &= \left[ 1 - C_{\mathcal{M},\tau,\kappa} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right) \right] \lambda_n^{(k)}. \end{aligned}$$

Then

$$\begin{aligned} \langle f_0, \mathcal{L}^{\tau,\kappa} f_0 \rangle_{L^2} + \langle f_-, \mathcal{L}^{\tau,\kappa} f_- \rangle_{L^2} &= \langle f, \mathcal{L}^{\tau,\kappa} f \rangle_{L^2} - \langle f_+, \mathcal{L}^{\tau,\kappa} f_+ \rangle_{L^2} \\ &\geq \lambda_n^{(k)} - C_{\mathcal{M},\tau,\kappa} b^{-1} (\lambda^{(k)} + b) k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right). \end{aligned} \quad (2.52)$$

We also have

$$\begin{aligned} \langle f_0, \mathcal{L}^{\tau,\kappa} f_0 \rangle_{L^2} &\leq (\lambda^{(k)} + a) \|f_0\|_{L^2}, \\ \langle f_-, \mathcal{L}^{\tau,\kappa} f_- \rangle_{L^2} &\leq (\lambda^{(k)} - c) \|f_-\|_{L^2}, \end{aligned}$$

which implies

$$\begin{aligned} \langle f_0, \mathcal{L}^{\tau,\kappa} f_0 \rangle_{L^2} + \langle f_-, \mathcal{L}^{\tau,\kappa} f_- \rangle_{L^2} &\leq \lambda^{(k)} (\|f_0\|_{L^2}^2 + \|f_-\|_{L^2}^2) + a \|f_0\|_{L^2}^2 - c \|f_-\|_{L^2}^2 \\ &\leq \lambda^{(k)} \|f\|_{L^2}^2 + a \|f\|_{L^2}^2 - c \|f_-\|_{L^2}^2. \end{aligned}$$

By Lemma 2.7.6(1), we have

$$\|f\|_{L^2} = \|\mathcal{I}v_k\|_{L^2} \leq \left[1 + C_\kappa h_n \sqrt{\lambda_n^{(k)}}\right] \|v_k\|_2 \leq 1 + C_\kappa h_n \sqrt{\lambda^{(k)}},$$

which gives

$$\langle f_0, \mathcal{L}^{\tau, \kappa} f_0 \rangle_{L^2} + \langle f_-, \mathcal{L}^{\tau, \kappa} f_- \rangle_{L^2} \leq (\lambda^{(k)} + a)(1 + C_\kappa h_n^2 \lambda^{(k)}) - c \|f_-\|_{L^2}^2.$$

Combining with (2.52) we have

$$(\lambda^{(k)} + a)(1 + C_\kappa h_n^2 \lambda^{(k)}) - c \|f_-\|_{L^2}^2 \geq \lambda_n^{(k)} - C_{\mathcal{M}, \tau, \kappa} b^{-1} (\lambda^{(k)} + b) k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right),$$

and then

$$\begin{aligned} \|f_-\|_{L^2}^2 &\leq C_{\mathcal{M}, \tau, \kappa} c^{-1} b^{-1} (\lambda^{(k)} + b) k \lambda^{(k)} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right) + c^{-1} a + c^{-1} \left| \lambda_n^{(k)} - \lambda^{(k)} \right| \\ &\leq C_{\mathcal{M}, \tau, \kappa} c^{-1} b^{-1} k \left[ \lambda^{(k)} \right]^2 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(k)}} \right) + c^{-1} a, \end{aligned} \quad (2.53)$$

where the assumption  $b \leq \lambda^{(k)}$  is used in the last step. The result then follows by combining (2.51) and (2.53) and noticing that  $\|\mathcal{I}v_k - \mathbb{P}_{(\lambda^{(k)} - c, \lambda^{(k)} + a)} \mathcal{I}v_k\|_{L^2}^2 = \|f_+\|_{L^2}^2 + \|f_-\|_{L^2}^2$ .  $\square$

Now we are ready to prove Theorem 2.4.7.

**Theorem 2.7.12** (Eigenfunction Approximation). *Let  $\lambda$  be an eigenvalue of  $\mathcal{L}^{\tau, \kappa}$  with multiplicity  $\ell$ , i.e.,*

$$\lambda^{(k_n-1)} < \lambda^{(k_n)} = \lambda = \dots = \lambda^{(k_n+\ell-1)} < \lambda^{(k_n+\ell)}.$$

*Suppose that  $h_n \sqrt{\lambda^{(k_n)}} \ll 1$  and  $\varepsilon_n \ll h_n$  for  $n$  large. Let  $\psi_n^{(k_n)}, \dots, \psi_n^{(k_n+\ell-1)}$  be orthonormal eigenvectors of  $L_n^{\tau, \kappa}$  associated with eigenvalues  $\lambda_n^{(k_n)}, \dots, \lambda_n^{(k_n+\ell-1)}$ . Then there exists*

orthonormal eigenfunctions  $\psi^{(k_n)}, \dots, \psi^{(k_n+\ell-1)}$  of  $\mathcal{L}^{\tau, \kappa}$  so that for  $j = k_n, \dots, k_n + \ell - 1$

$$\begin{aligned} \|\psi_n^{(j)} \circ T_n - \psi^{(j)}\|_{L^2}^2 &\leq C_{\mathcal{M}, \tau, \kappa} j^3 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right), \\ \|\psi_n^{(j)} \circ \mathcal{T}_n - \psi^{(j)}\|_{L^2}^2 &\leq C_{\mathcal{M}, \tau, \kappa} (\log n)^{mc_m} j^3 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right). \end{aligned}$$

*Proof.* For each  $j = k, \dots, k + \ell - 1$ , let  $a = (\varepsilon_n/h_n + h_n \sqrt{\lambda^{(j)}})$  and  $b = c = \frac{\delta_\lambda}{2}$ , where

$$\delta_\lambda = \min\{\lambda^{(k_n)} - \lambda^{(k_n-1)}, \lambda^{(k_n+\ell)} - \lambda^{(k_n+\ell-1)}\}$$

so that the assumptions of Lemma 2.7.11 are satisfied. Indeed,  $a \leq b \leq c \leq \lambda^{(k)}$  and the interval  $(\lambda^{(j)} + a, \lambda^{(j)} + b)$  does not contain any eigenvalue of  $\mathcal{L}^{\tau, \kappa}$  and  $\mathbb{P}_{(\lambda^{(j)}-c, \lambda^{(j)}+a)} = \mathbb{P}_{\{\lambda^{(j)}\}}$ . Hence we obtain

$$\|\mathcal{I}\psi_n^{(j)} - \tilde{\psi}^{(j)}\|_{L^2}^2 \leq C_{\mathcal{M}, \tau, \kappa} \delta_\lambda^{-2} j \left[ \lambda^{(j)} \right]^2 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right),$$

where  $\tilde{\psi}^{(j)} = \mathbb{P}_{\{\lambda\}} \mathcal{I}\psi_n^{(j)}$  is a  $\lambda$ -eigenfunction of  $\mathcal{L}^{\tau, \kappa}$ . Lemma 2.7.6(1) implies that  $\mathcal{I}$  is almost an isometry on the span of  $\psi_n^{(k)}, \dots, \psi_n^{(k+\ell-1)}$  and by the polarization identity we get that the  $\mathcal{I}\psi_n^{(j)}$ 's are almost orthonormal up to  $C_{\mathcal{M}, \tau, \kappa} h_n \sqrt{\lambda^{(j)}}$ . This implies the  $\tilde{\psi}^{(j)}$ 's are almost orthogonal up to  $C_{\mathcal{M}, \tau, \kappa} \delta_\lambda^{-2} j [\lambda^{(j)}]^2 (\varepsilon_n/h_n + h_n \sqrt{\lambda^{(j)}})$ . Hence letting  $\{\psi^{(j)}\}_{j=k}^{k+\ell-1}$  be the Gram-Schmidt orthogonalization of  $\{\tilde{\psi}^{(j)}\}_{j=k}^{k+\ell-1}$ , we get

$$\|\mathcal{I}\psi_n^{(j)} - \psi^{(j)}\|_{L^2}^2 \leq C_{\mathcal{M}, \tau, \kappa} \delta_\lambda^{-2} j \left[ \lambda^{(j)} \right]^2 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right).$$

Using (2.43) that  $\|\mathcal{I}v - P^*v\|_{L^2}^2 \leq Ch_n^2 D_{h_n}[v]$  gives

$$\|P^*\psi_n^{(j)} - \psi^{(j)}\|_{L^2}^2 \leq C_{\mathcal{M}, \tau, \kappa} \delta_\lambda^{-2} j \left[ \lambda^{(j)} \right]^2 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right).$$

By Weyl's law that  $\lambda^{(j)} \asymp j^{\frac{2}{m}}$  and hence  $\delta_\lambda \asymp j^{\frac{2}{m}-1} \asymp j^{-1}\lambda^{(j)}$ , we conclude that

$$\|P^*\psi_n^{(j)} - \psi^{(j)}\|_{L^2}^2 \leq C_{\mathcal{M},\tau,\kappa} j^3 \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}} \right),$$

which is the first assertion of the theorem by noticing that  $P^*\psi_n^{(j)} = \psi_n^{(j)} \circ T_n$ . Now by Lemma 17 and the proof of [García Trillos et al., 2019a, Theorem 6], we have

$$\begin{aligned} \|\psi_n^{(j)} \circ T_n - \psi^{(j)}\|_{L^2} &\leq C_{\mathcal{M}} \left[ \lambda_j^{\frac{m+1}{4}} \varepsilon_n + (\log n)^{\frac{m c_m}{2}} \|P^*\psi_n^{(j)} - \psi^{(j)}\|_{L^2} \right] \\ &\leq C_{\mathcal{M},\tau,\kappa} \left[ j^{\frac{m+1}{2m}} \varepsilon_n + (\log n)^{\frac{m c_m}{2}} j^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}}} \right] \\ &\leq C_{\mathcal{M},\tau,\kappa} (\log n)^{\frac{m c_m}{2}} j^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(j)}}}, \end{aligned}$$

where we have used that  $\varepsilon_n \ll h_n$  in the last step. □

### 2.7.4 Convergence of Graph Matérn Field

Now we are ready to prove Theorem 2.4.2. Theorem 2.4.4 can be proved in the same fashion using the second assertion of Theorem 2.7.12.

**Theorem 2.7.13.** *Suppose  $\tau$  is Lipschitz,  $\kappa \in C^1(\mathcal{M})$  and both are bounded below by positive constants. Let  $s > m$  and*

$$\frac{(\log n)^{c_m}}{n^{1/m}} \ll h_n \ll \frac{1}{n^{1/2s}},$$

where  $c_m = 3/4$  if  $m = 2$  and  $c_m = 1/m$  otherwise. Then, with probability one,

$$\mathbb{E} \|u_n \circ T_n - u\|_{L^2} \xrightarrow{n \rightarrow \infty} 0.$$

If further  $s > (5m + 1)/2$  and

$$h_n \asymp \sqrt{\frac{(\log n)^{c_m}}{n^{1/m}}}. \quad (2.54)$$

Then, with probability one,

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} = O\left(\sqrt{h_n}\right) = O\left(\frac{(\log n)^{c_m/4}}{n^{1/4m}}\right).$$

*Proof.* Suppose we are in a realization where the conclusion of Proposition 2.4.1 holds. Suppose  $k_n$  is chosen so that  $n^{m/2s} \ll k_n \ll h_n^{-m}$ . Note that this is possible given the scaling of  $h_n$ . Then by Theorem 2.7.12 we can fix orthonormal eigenfunctions  $\{\psi^{(i)}\}_{i=1}^\infty$  of  $\mathcal{L}^{\tau, \kappa}$  and  $\{\psi_i^{(n)}\}_{i=1}^n$  of  $L_n^{\tau, \kappa}$  for each  $n$  so that

$$\|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} \leq C_{\mathcal{M}, \tau, \kappa} i^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}}}, \quad (2.55)$$

for  $i = 1, \dots, k_n$ . Recall  $u_n$  and  $u$  have the following representations

$$\begin{aligned} u_n &:= \tau_n^{\frac{s}{2} - \frac{m}{4}} \kappa_n^{\frac{m}{2}} \sum_{i=1}^n \left[\lambda_n^{(i)}\right]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)}, \\ u &= \tau^{\frac{s}{2} - \frac{m}{4}} \kappa^{\frac{m}{2}} \sum_{i=1}^\infty \left[\lambda^{(i)}\right]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)}. \end{aligned}$$

Since  $\tau_n$  is the restriction of  $\tau$  to  $\mathcal{M}_n$ , we have  $\tau_n \circ T_n = \tau \circ T_n$  and similarly for  $\kappa_n$ . Therefore  $u_n \circ T_n$  has the expression

$$u_n \circ T_n = [\tau \circ T_n]^{\frac{s}{2} - \frac{m}{4}} [\kappa \circ T_n]^{\frac{m}{2}} \sum_{i=1}^n \left[\lambda_n^{(i)}\right]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)} \circ T_n.$$

To bound the expected  $L^2$  distance between  $u_n \circ T_n$  and  $u$ , we introduce four intermediate

functions.

$$\begin{aligned}
u_n^{k_n} &:= [\tau \circ T_n]^{\frac{s}{2} - \frac{m}{4}} [\kappa \circ T_n]^{\frac{m}{2}} \sum_{i=1}^{k_n} [\lambda_n^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)} \circ T_n, \\
\tilde{u}_n^{k_n} &:= [\tau \circ T_n]^{\frac{s}{2} - \frac{m}{4}} [\kappa \circ T_n]^{\frac{m}{2}} \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi_n^{(i)} \circ T_n, \\
\tilde{u}^{k_n} &:= [\tau \circ T_n]^{\frac{s}{2} - \frac{m}{4}} [\kappa \circ T_n]^{\frac{m}{2}} \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)} \\
u^{k_n} &:= \tau^{\frac{s}{2} - \frac{m}{4}} \kappa^{\frac{m}{2}} \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}} \xi^{(i)} \psi^{(i)}.
\end{aligned}$$

It then suffices to bound the difference between any two consecutive functions. By Theorem 2.4.6 and Weyl's law we have that  $\lambda_n^{(k_n)} \gtrsim \lambda^{(k_n)} \gtrsim k_n^{2/m}$ , which gives

$$\mathbb{E} \|u_n \circ T_n - u_n^{k_n}\|_{L^2} \leq \beta^{\frac{s}{2} - \frac{m}{4}} \alpha^{\frac{m}{2}} \left( \sum_{i=k_n+1}^n [\lambda_n^{(i)}]^{-s} \right)^{\frac{1}{2}} \lesssim \left( n [\lambda_n^{(k_n)}]^{-s} \right)^{\frac{1}{2}} \lesssim \sqrt{n} k_n^{-\frac{s}{m}}. \tag{2.56}$$

Similarly,

$$\mathbb{E} \|u^{k_n} - u\|_{L^2} \lesssim \left( \sum_{i=k_n+1}^{\infty} [\lambda^{(i)}]^{-s} \right)^{\frac{1}{2}} \lesssim \left( \sum_{i=k_n+1}^{\infty} i^{-\frac{2s}{m}} \right)^{\frac{1}{2}} \lesssim \left( \int_{k_n}^{\infty} x^{-\frac{2s}{m}} \right)^{\frac{1}{2}} \lesssim k_n^{\frac{1}{2} - \frac{s}{m}}. \tag{2.57}$$

Both (2.56) and (2.57) converges to 0 by the choice of  $k_n$ .

Next, since both  $\lambda_n^{(i)}$  and  $\lambda^{(i)}$  are bounded below by  $\min \tau > 0$ , by Lipschitz continuity of  $x^{-s/2}$  away from 0 we have, for  $i = 1, \dots, k_n$ ,

$$\left| [\lambda_n^{(i)}]^{-\frac{s}{2}} - [\lambda^{(i)}]^{-\frac{s}{2}} \right| \lesssim \left( [\lambda_n^{(i)}] \wedge [\lambda^{(i)}] \right)^{-\frac{s}{2} - 1} |\lambda_n^{(i)} - \lambda^{(i)}| \lesssim [\lambda^{(i)}]^{-\frac{s}{2}} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}} \right)$$



Hence

$$\begin{aligned}
\mathbb{E}\|u_n^{k_n} - \tilde{u}_n^{k_n}\|_{L^2} &\lesssim \left( \sum_{i=1}^{k_n} \left( [\lambda_n^{(i)}]^{-\frac{s}{2}} - [\lambda^{(i)}]^{-\frac{s}{2}} \right)^2 \right)^{\frac{1}{2}} \\
&\lesssim \left( \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-s} \left( \frac{\varepsilon_n}{h_n} + h_n \sqrt{\lambda^{(i)}} \right)^2 \right)^{\frac{1}{2}} \\
&\lesssim \left( \frac{\varepsilon_n}{h_n} + h_n k_n^{\frac{1}{m}} \right) \left( \sum_{i=1}^{k_n} i^{-\frac{2s}{m}} \right)^{\frac{1}{2}}.
\end{aligned} \tag{2.58}$$

The last expression goes to 0 since  $s > m$ ,  $\varepsilon_n \ll h_n$  and  $k_n \ll h_n^{-m}$ .

Now by Lipschitz continuity of  $\tau$  and  $\kappa$ , and Proposition 2.4.1 that  $d(x, T_n(x)) \leq \varepsilon_n$ , we have for all  $x \in \mathcal{M}$

$$\left| \tau(T_n(x))^{\frac{s}{2} - \frac{m}{2}} \kappa(T_n(x))^{\frac{m}{2}} - \tau(x)^{\frac{s}{2} - \frac{m}{2}} \kappa(x)^{\frac{m}{2}} \right| \lesssim \varepsilon_n.$$

Therefore

$$\mathbb{E}\|\tilde{u}_n^{k_n} - u_n^{k_n}\|_{L^2} \lesssim \left\| (\tau \circ T_n)^{\frac{s}{2} - \frac{m}{4}} (\kappa \circ T_n)^{\frac{m}{2}} - \tau^{\frac{s}{2} - \frac{m}{4}} \kappa^{\frac{m}{2}} \right\|_{\infty} \left( \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-s} \right) \lesssim \varepsilon_n, \tag{2.59}$$

which converges to zero.

Finally, for fixed  $\ell \in \mathbb{N}$ , we have by using the fact that  $\|\psi_n^{(i)} \circ T_n\|_{L^2} = \|\psi_n^{(i)}\|_2 = 1$

$$\begin{aligned}
\mathbb{E}\|\tilde{u}_n^{k_n} - \tilde{u}_n^{k_n}\|_{L^2} &\lesssim \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}} \|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} \\
&\lesssim \sum_{i=1}^{\ell} [\lambda^{(i)}]^{-\frac{s}{2}} \|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} + \sum_{i=\ell+1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}}.
\end{aligned} \tag{2.60}$$

By (2.55) we have  $\|\psi_n^{(i)} \circ T_n - \psi^{(i)}\|_{L^2} \xrightarrow{n \rightarrow \infty} 0$  for  $i = 1, \dots, \ell$  since  $\ell$  is fixed. Therefore

we have

$$\limsup_{n \rightarrow \infty} \mathbb{E} \|\tilde{u}_n^{k_n} - \tilde{u}^{k_n}\|_{L^2} \lesssim \sum_{i=\ell+1}^{\infty} [\lambda^{(i)}]^{-\frac{s}{2}}.$$

Since  $\ell$  is arbitrary, the last expression goes to 0 as  $\ell \rightarrow \infty$  under the assumption  $s > m$ . Hence by combining all the pieces we get  $\mathbb{E} \|u_n \circ T_n - u\|_{L^2} \xrightarrow{n \rightarrow \infty} 0$ .

Now in order to obtain rates of convergence, we need the additional assumption that  $s > \frac{5}{2}m + \frac{1}{2}$  and to refine in particular the estimates for  $\mathbb{E} \|u_n^{k_n} - \tilde{u}_n^{k_n}\|_{L^2}$  and  $\mathbb{E} \|\tilde{u}_n^{k_n} - \tilde{u}^{k_n}\|_{L^2}$ . By (2.58), we have

$$\mathbb{E} \|u_n^{k_n} - \tilde{u}_n^{k_n}\|_{L^2} \lesssim \left( \frac{\varepsilon_n}{h_n} + h_n \right) \left( \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-s+1} \right)^{\frac{1}{2}} \lesssim \left( \frac{\varepsilon_n}{h_n} + h_n \right), \quad (2.61)$$

where the last step follows by the assumption  $s > \frac{5}{2}m + \frac{1}{2}$ . By (2.55), we can further bound (2.60) by

$$\begin{aligned} \mathbb{E} \|\tilde{u}_n^{k_n} - \tilde{u}^{k_n}\|_{L^2} &\lesssim \sum_{i=1}^{k_n} [\lambda^{(i)}]^{-\frac{s}{2}} i^{\frac{3}{2}} \sqrt{\frac{\varepsilon_n}{h_n} + h_n} \sqrt{\lambda^{(i)}} \\ &\lesssim \sqrt{\frac{\varepsilon_n}{h_n} + h_n} \sum_{i=1}^{k_n} i^{\frac{3}{2}} [\lambda^{(i)}]^{-\frac{s}{2} + \frac{1}{4}} \\ &\lesssim \sqrt{\frac{\varepsilon_n}{h_n} + h_n} \sum_{i=1}^{k_n} i^{\frac{3}{2} - \frac{s}{m} + \frac{1}{2m}} \lesssim \sqrt{\frac{\varepsilon_n}{h_n} + h_n}, \end{aligned} \quad (2.62)$$

where the last step follows from that  $s > \frac{5}{2}m + \frac{1}{2}$ . Now combining (2.56), (2.57), (2.59), (2.61), (2.62), we see that the error is dominated by

$$\mathbb{E} \|u_n \circ T_n - u\|_{L^2} \lesssim \sqrt{n k_n^{-\frac{s}{m}}} + \sqrt{\frac{\varepsilon_n}{h_n} + h_n}.$$

Therefore by setting

$$h_n \asymp (\log n)^{\frac{cm}{2}} n^{-\frac{1}{2m}}, \quad k_n \asymp n^{\frac{2m+1}{4s}}$$

we have

$$\mathbb{E}\|u_n \circ T_n - u\|_{L^2} \lesssim (\log n)^{\frac{cm}{4}} n^{-\frac{1}{4m}}.$$

□

We end this subsection with a remark that Theorem 2.7.13 can be stated in terms of the  $TL^2$  metric proposed in García Trillos and Slepčev [2016b]. Let  $\mathcal{P}(\mathcal{M})$  be the space of Borel probability measures on  $\mathcal{M}$ . Define the  $TL^2$  space as

$$TL^2 := \left\{ (\mu, f) : \mu \in \mathcal{P}(\mathcal{M}), f \in L^2(\mu) \right\},$$

endowed with the metric

$$d_{TL^2}\left((\mu_1, f_1), (\mu_2, f_2)\right) := \inf_{\omega \in \mathcal{C}(\mu_1, \mu_2)} \left[ \int_{\mathcal{M}} \int_{\mathcal{M}} \left( d_{\mathcal{M}}(x, y)^2 + |f_1(x) - f_2(y)|^2 \right) d\omega(x, y) \right]^{\frac{1}{2}},$$

where  $\mathcal{C}$  is the set of couplings between  $\mu_1$  and  $\mu_2$  and  $d_{\mathcal{M}}$  is the geodesic distance on  $\mathcal{M}$ . The  $d_{TL^2}$  metric is a natural generalization of  $L^2$  convergence of functions and weak convergence of measures [García Trillos and Slepčev, 2016b], which allows comparison of functions defined over the point cloud with functions defined on  $\mathcal{M}$ . It bypasses the need to consider a specific transport map and thus may be of independent interest.

The assertions of Theorem 2.7.13 continue to hold if  $\mathbb{E}\|u_n \circ T_n - u\|_{L^2}$  is replaced by  $\mathbb{E}\left[ d_{TL^2}\left((\gamma_n, u_n), (\gamma, u)\right) \right]$ . The proof follows immediately from the definition since the transport maps  $T_n$  induce a coupling defined as  $\omega_{T_n} := (I \times T_n)_\# \gamma$ , the push-forward of  $\gamma$

under  $I \times T_n : \mathcal{M} \rightarrow \mathcal{M} \times \mathcal{M}_n$ , where  $(I \times T_n)(x) = (x, T_n(x))$ . Hence we see that

$$d_{TL^2}((\gamma_n, u_n), (\gamma, u)) \leq \left[ \int_{\mathcal{M}} \left( d_{\mathcal{M}}(x, T_n(x))^2 + |u_n(T_n(x)) - u(x)|^2 \right) d\gamma(x) \right]^{\frac{1}{2}}$$

$$\lesssim \varepsilon_n + \|u_n \circ T_n - u\|_{L^2}.$$

# CHAPTER 3

## KERNEL METHODS FOR BAYESIAN ELLIPTIC INVERSE PROBLEMS ON MANIFOLDS

### 3.1 Introduction

Partial Differential Equations (PDEs) on manifolds are used to model a variety of physical and biological phenomena including pattern formation on biological surfaces, phase separation in bio-membranes, tumor growth, and surfactants on fluid interfaces [Eilks and Elliott, 2008, Elliott and Stinner, 2010, Chaplain et al., 2003, Xu et al., 2006]. In this chapter we focus on the inversion of *elliptic* PDEs for two main reasons. First, elliptic PDEs are ubiquitous in applications and they are used, for instance, as simplified models for groundwaterflow and oil reservoir simulation. The need to specify uncertain input parameters of these models leads naturally to the inverse problem of determining the permeability from the pressure under a Darcy model of flow in a porous medium [McLaughlin and Townley, 1996, Lorentzen et al., 2012, Iglesias and Dawson, 2007, Ping and Zhang, 2014]. Second, elliptic models are widely used to test algorithms for forward propagation of uncertainty [Frauenfelder et al., 2005, Cohen and Schwab, 2011, Babuska et al., 2004] and Bayesian inversion [Stuart, 2010, Cotter et al., 2010b, García Trillos and Sanz-Alonso, 2017]. Despite the applied importance of elliptic inversion, the manifold setting that we consider is largely unexplored and may allow for more realistic modelling in applications. For example, the variables of interest in the groundwaterflow problem may not be confined to a *flat* two-dimensional domain and knowledge of the underlying flow surface may be limited to a point cloud of landmark locations.

The aim of this chapter is to study the formulation and implementation of Bayesian inverse problems to learn input parameters of PDEs defined on manifolds. Specifically, we study the inverse problem of recovering the diffusion coefficient of a second-order, divergence-

form elliptic equation, given noisy measurements of the solution. While our interest lies in solving the inverse problem, most of our efforts are devoted to studying the approximation of the *forward map* (the operator that takes the input parameter to the solution of the PDE). Several techniques to approximate the forward map have been proposed in the extensive literature on numerical methods for PDEs on manifolds. For example, finite element methods [Dziuk and Elliott, 2013, Camacho and Demlow, 2015, Bonito et al., 2016], level-set methods [Bertalmío et al., 2001, Mévoli et al., 2004], closest point methods [Ruuth and Merriman, 2008], or mesh-free radial basis function methods [Piret, 2012]. The implementation detail of each of the existing methods is different, but a unifying theme is the need to have a representation of the manifold in order to approximate the differential operator. Unfortunately, these approaches are difficult to implement when one only has access to an unstructured point cloud of manifold samples and meshing is challenging, or when the dimension of the ambient space is large but the manifold dimension is moderate.

In this chapter, we avoid the problems associated with the representation of the manifold by directly approximating the differential operator in the forward map with an appropriate kernel integral operator. With a consistent kernel approximation to the differential operator on the manifold, the numerical implementation can be performed naturally by discretizing the corresponding integral operator on a *point cloud* of manifold samples without further knowledge of the underlying manifold or its Riemannian metric. Building on this construction, we propose a fully discrete, mesh-free approach to the numerical solution of Bayesian inverse problems on point clouds. The idea of facilitating the discretization of PDEs on manifolds by an integral equation approximation can also be found in the recent papers [Li and Shi, 2016, Li et al., 2017, Gilani and Harlim, 2019], all of which build on manifold learning techniques and analyses. Our perspective in this chapter and in Gilani and Harlim [2019], Berry and Sauer [2016] is in contrast to the one taken in Belkin and Niyogi [2004, 2008], Coifman et al. [2005], Coifman and Lafon [2006], Berry and Harlim [2016]. Rather than iden-

tifying the limiting continuum operator of different normalizations of graph-Laplacians, our interest is to define a suitable kernel to approximate a given anisotropic diffusion operator on the underlying manifold.

We adopt a Bayesian approach to the inverse problem [Kaipo and Somersalo, 2006, Calvetti and Somersalo, 2007, Stuart, 2010, Sanz-Alonso et al., 2019], where we set a *prior* distribution on the unknown PDE input parameter, and condition on observed data to find its *posterior* distribution. The Bayesian approach is largely motivated by the following advantages. First, the posterior covariance and posterior confidence intervals may be used to quantify the uncertainty in the parameter reconstruction. Second, the Bayesian formulation leads to a well-posed inverse problem [Marzouk and Xiu, 2009, Stuart, 2010, Sanz-Alonso et al., 2019] by which a small perturbation on the data, the prior distribution, or the forward map leads to a small perturbation in the posterior solution. Our main theoretical result is an example of the well-posedness of the Bayesian formulation: we deduce a total variation error bound between the true posterior distribution and its kernel-based approximation from a new error bound between the forward map and its kernel approximation. The new forward map error bound, with a dependence on the diffusion coefficient, builds on existing results on the point-wise convergence kernel approximations to elliptic operators [Coifman and Lafon, 2006, Berry and Harlim, 2016].

The advantages of the Bayesian approach outlined above come with a cost: the need to specify a prior distribution on the unknown. The choice of prior is crucial as it determines the support of the posterior, but unfortunately this choice is often only guided by ad-hoc and computational considerations. In this chapter we consider a two-parameter family of log-Gaussian field priors on manifolds, defining the covariance through the Laplace Beltrami operator on the manifold [Lindgren et al., 2011, Dunlop et al., 2017]. The two prior parameters allow the specification of the smoothness and lengthscale of prior (and hence posterior) draws, and the lengthscale may be learned from data using a hierarchical approach as de-

tailed in our numerical experiments. In addition to their flexibility, a further advantage of our choice of priors is that they allow the infusion of geometric information from the manifold on the reconstructed input by expressing it as a random combination of the first eigenfunctions of the Laplacian. Moreover, in the absence of a full representation of the manifold, these priors can be consistently discretized using a graph-Laplacian [Von Luxburg, 2007, García Trillos and Sanz-Alonso, 2018]. We refer to Bertozzi et al. [2018], Gao et al. [2019], García Trillos and Sanz-Alonso [2018], García Trillos et al. [2020b] for recent applications and references on Gaussian processes on manifolds. From a computational viewpoint, the use of log-Gaussian priors and the existence of a continuum limit facilitate the design of MCMC algorithms that scale well with the size of the point cloud, as shown in a linear inverse problem in García Trillos et al. [2020b]. In this regard, a simple but powerful idea is to use a proposal kernel that satisfies detailed balance with respect to the prior [Cotter et al., 2010a].

**Outline and Main Contributions** We close this introduction with an outline of the rest of the chapter, summarizing the main contributions of each section.

- In Section 3.2 we give a brief introduction to the Bayesian formulation of inverse problems, and formulate the problem on a manifold. The main novel contributions of this section are: i) to introduce a kernel-based approximation to the forward map and an associated approximation to the posterior in the continuum (Section 3.2.2); and ii) to employ the kernel approximations to formulate a Bayesian solution to the inverse problems on point clouds (Section 3.2.3). These kernel and point cloud approximations are inspired by manifold learning and data analysis techniques.
- Section 3.3 contains the main theoretical contributions of this chapter. Theorem 3.3.1 gives an error bound between the true and kernel-based forward maps, and Theorem 3.3.6 establishes a bound on the total variation distance between the posterior and



its kernel-based approximation. The main novelty of Theorem 3.3.1 is to generalize the analysis in Coifman and Lafon [2006] to account for anisotropic diffusion and, more importantly, to explicitly track the dependence of the diffusion coefficient in the error bounds. Understanding this dependency is necessary in order to guarantee the closeness of the true and approximate posterior distributions shown in Theorem 3.3.6.

- In Section 3.4 we discuss the practical implementation of the methods, provide guidelines for the choice of tuning parameters and for the posterior sampling, and conduct three numerical experiments of increasing difficulty to illustrate the applicability of our approach. We also consider in Section 3.4.5 a hierarchical formulation of the inverse problem, where the prior lengthscale is learned from the data, when in absence of such knowledge.
- We close in Section 3.5 with conclusions and open directions for research that stem from our work.

**Notation and Setting** Throughout this chapter  $\mathcal{M}$  will denote an  $m$ -dimensional smooth Riemannian manifold embedded in  $\mathbb{R}^d$ . We will denote by  $\mathcal{C}^k := \mathcal{C}^k(\mathcal{M})$  the space of  $k$ -times differentiable functions on  $\mathcal{M}$  and by  $\mathcal{C}^{k,\alpha} := \mathcal{C}^{k,\alpha}(\mathcal{M})$  the space of  $k$ -times differentiable functions whose  $k$ -th partial derivatives are Hölder continuous with exponent  $\alpha$ . We will assume that the manifold  $\mathcal{M}$  is compact and has no boundary, thus avoiding the technicalities necessary to deal with boundary conditions. The theoretical and computational investigation of point cloud approximation to PDEs supplemented with boundary conditions is a topic of current research Li and Shi [2016], Li et al. [2017], Gilani and Harlim [2019]. Due to the lack of boundary conditions, the elliptic problem that we consider is unique up to a constant. We will enforce uniqueness by working on the space  $L_0^2 := L_0^2(\mathcal{M})$  of mean-zero square integrable functions on  $\mathcal{M}$ .

## 3.2 Bayesian Inverse Problems on Manifolds and Point Clouds

We start in Section 3.2.1 by recalling the Bayesian formulation of elliptic inverse problems in a given manifold. We then introduce in Section 3.2.2 a kernel-based approximation to the forward map and a corresponding approximation to the posterior, both of which will be analyzed in Section 3.3. Finally in Section 3.2.3 we introduce a point cloud approximation of the kernel forward map leading to a formulation of the elliptic problem on point clouds without reference to the underlying manifold. We will investigate numerically the implementation of elliptic Bayesian inverse problems on point clouds in Section 3.4.

### 3.2.1 Bayesian Elliptic Inverse Problems on Manifolds

We consider the elliptic equation

$$\mathcal{L}^\kappa u := -\operatorname{div}(\kappa \nabla u) = f, \quad x \in \mathcal{M}, \quad (3.1)$$

where  $\kappa$  is a function on  $\mathcal{M}$ . Here and throughout, the differential operators are defined with respect to the Riemannian metric inherited by  $\mathcal{M}$  from  $\mathbb{R}^d$ . We are interested in the inverse problem of determining the diffusion coefficient  $\kappa$  from noisy measurements of  $u$  of the form

$$y = \mathcal{D}(u) + \eta, \quad (3.2)$$

where the *observation map*  $\mathcal{D} : L^2 \rightarrow \mathbb{R}^J$  will be assumed to be known. Examples of observation maps will be discussed in Section 3.2.1.2. We adopt a Bayesian perspective to the inverse problem, described succinctly in what follows; we refer to Kaipio and Somersalo [2006], Stuart [2010], Sanz-Alonso et al. [2019] for a more detailed account. In short, the Bayesian formulation of inverse problems involves specifying a *prior* distribution  $\pi$  for the unknown PDE input  $\kappa$  and a distribution for the observation noise  $\eta$ . Once these distributions

have been specified, the solution to the inverse problem is the *posterior* distribution of the variable  $\kappa$  conditioned on the observed data  $y$ . For simplicity of exposition, we will assume throughout that the observation noise is centered and Gaussian,  $\eta \sim \mathcal{N}(0, \Gamma)$  for given positive definite  $\Gamma \in \mathbb{R}^{J \times J}$ .

Writing  $\kappa = e^\theta$ , we take a Gaussian prior for  $\theta$  supported on a Banach space  $\mathcal{B}$ . A specific form of prior, widely used in applications, will be described in Section 3.2.1.1. Our assumptions on  $\kappa$  and  $f$  in Section 3.3 will guarantee the existence of a unique solution to equation (3.1) in the space  $L_0^2$  of mean-zero square integrable functions on  $\mathcal{M}$ . This, in turn, allows us to define a *forward map*  $\mathcal{F} : \theta \in \mathcal{B} \mapsto u \in L_0^2$ . Provided that the map  $\mathcal{G} := \mathcal{D} \circ \mathcal{F} : \mathcal{B} \rightarrow \mathbb{R}^J$  is measurable and that the prior is supported on  $\mathcal{B}$ , the posterior  $\pi^y$  can be written as a change of measure with respect to the prior

$$\frac{d\pi^y}{d\pi}(\theta) \propto \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right), \quad (3.3)$$

with  $|\cdot|_\Gamma^2 := \langle \cdot, \Gamma^{-1} \cdot \rangle$ . Equation (3.3) shows that the posterior distribution  $\pi^y$  is defined by reweighting the prior, favoring unknowns  $\theta$  that produce a good match with the data  $y$  through a likelihood function (the right-hand term), implied by equation (3.2) and the assumed Gaussian distribution of the noise  $\eta$ .

For our theoretical results in Section 3.3 we will take  $\mathcal{B} = \mathcal{C}^4$  and assume that  $f \in \mathcal{C}^{3,\alpha}$ , for  $0 < \alpha < 1$ . These assumptions guarantee [Gilbarg and Trudinger, 2015] that almost surely with respect to the prior, the diffusion coefficient  $\kappa$  is uniformly elliptic, and the unique solution of equation (3.1) in  $L_0^2$  lives in  $\mathcal{C}^{5,\alpha}$ , allowing us to establish a stability result for an approximation of the forward map. We believe, however, that these strong regularity conditions can be relaxed.

### 3.2.1.1 Matérn-type Prior

Here we describe a choice of prior that is widely used in applications in the geophysical sciences and spatial statistics [Stein, 1999a]; in Section 3.2.3.1 we will introduce a point cloud approximation to this prior used in our numerical experiments in Section 3.4. Since equation (3.3) implies that the support of the prior determines the support of the posterior, it should both capture the geometry of the manifold  $\mathcal{M}$  and have enough expressivity to include a wide class of functions. This motivates to choose the prior from a flexible two-parameter family of Gaussian measures on  $L^2$ . Precisely, we will consider priors of the form

$$\pi = \mathcal{N}(0, C_{\tau,s}), \quad C_{\tau,s} = c(\tau)(\tau I + \Delta_{\mathcal{M}})^{-s}, \quad (3.4)$$

where  $\Delta_{\mathcal{M}} := -\operatorname{div}(\nabla \cdot)$  is the Laplace-Beltrami operator on  $\mathcal{M}$ ,  $\tau > 0, s > \frac{m}{2}$  are two free parameters, whose intuitive interpretation will be given below, and  $c(\tau)$  is a normalizing constant. Let  $\{(\lambda_i, \phi_i)\}_{i=1}^{\infty}$  be eigenvalue-eigenvector pairs for  $\Delta_{\mathcal{M}}$  with  $\lambda_i$ 's increasing. Then by the Karhunen-Loève expansion, random samples of  $\pi$  admit a series expansion

$$v = c(\tau)^{1/2} \sum_{i=1}^{\infty} (\tau + \lambda_i)^{-s/2} \xi_i \phi_i, \quad (3.5)$$

where  $\xi_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ . The eigenfunctions of the Laplacian contain geometric information on the underlying manifold, and therefore constitute a natural basis for functions on the manifold. By Weyl's law,  $\lambda_i \asymp i^{2/m}$  and so the requirement  $s > \frac{m}{2}$  is to ensure that samples from  $\pi$  belong to  $L^2$  almost surely. Moreover, the parameter  $s$  controls the rate of decay of the coefficients and hence characterizes the regularity of the samples. The role of  $\tau$  is more delicate. If we write the coefficients as  $v_i := (\tau + \lambda_i)^{-s/2} = \tau^{-s/2} (1 + \frac{\lambda_i}{\tau})^{-s/2}$ , then we can see that the  $v_i$ 's will be small for  $\lambda_i$ 's that are much larger than  $\tau$ . In particular, the only significant  $v_i$ 's are those where the corresponding  $\lambda_i$ 's are on the same order of  $\tau$  and hence

$\tau$  determines the significant basis functions in the expansion (3.5). Since the eigenfunctions  $\{\phi_i\}_{i=1}^\infty$  represent increasing frequencies,  $\tau$  can be interpreted as a lengthscale parameter. It can be seen from (3.5) that  $\tau$  also affects the amplitude of the samples and this motivates to choose the normalizing constant so that  $v$  has a fixed variance, which we set to be 1:

$$c(\tau) = \frac{1}{\sum_{i=1}^\infty (\tau + \lambda_i)^{-s}}. \quad (3.6)$$

Such priors are widely used when  $\mathcal{M}$  is a domain in a Euclidean space and are related to the Whittle-Matérn distributions [Dunlop et al., 2017]. In Lindgren et al. [2011] the authors also considered their extension to manifolds. It can be shown that for  $s$  large enough, samples from  $\pi$  belong to  $\mathcal{C}^k$  almost surely. And since the embedding of  $\mathcal{C}^k$  into  $L^2$  is continuous, the restriction of  $\pi$  to  $\mathcal{C}^k$  is again a Gaussian measure [Bogachev, 1998]. Hence for our purpose, we choose  $s$  so that  $\pi$  is a Gaussian measure on  $\mathcal{C}^4$ . In particular we will need later in Section 3.3 the result from Fernique’s theorem [Fernique, 1970] that there exists  $\alpha > 0$  such that

$$\int_{\mathcal{B}} \exp\left(\alpha \|\theta\|_{\mathcal{C}^4}^2\right) d\pi(\theta) < \infty.$$

**Remark 3.2.1.** *Choosing a prior with parameter  $\tau$  that is far from the true lengthscale of the unknown parameter would lead to poor Bayesian inversion. This can be problematic if such prior knowledge is not available, but may be at least partially alleviated by considering a hierarchical formulation specifying a joint prior on both  $\tau$  and  $\theta$ , so that the lengthscale is learned from data simultaneously with the unknown  $\theta$ ; implementation details of the hierarchical formulation will be given in Section 3.4.5.*

### 3.2.1.2 Observation Maps

Here we give two examples of observation maps that we shall consider. For theoretical considerations, we assume that the observation map is of the form  $\mathcal{D}(u) = (\ell_1(u), \dots, \ell_J(u))^T$ ,

where each  $\ell_j$  is a bounded linear functional on  $L^2$ . A widely used example is the smoothed observation at a point  $x_j$ :  $\ell_j(u) = \int K(x_j, x)u(x)dV(x)$ , where  $K$  is a kernel such as the Gaussian kernel [Brennan et al., 2020]; this type of observations arise in practice when the data is gathered from a collection of spatially distributed sensors located in the vicinity of landmark points  $x_j$ . Equally common is the pointwise evaluation [García Trillos and Sanz-Alonso, 2017, Dashti and Stuart, 2017]:  $\ell_j = u(x_j)$ . Notice that pointwise evaluation is not a bounded linear functional on  $L^2$  but can be approximated by smoothed observation arbitrarily well for continuous  $u$ 's. We remark that the boundedness assumption of  $\ell_j$  is only a technical one and for the numerical experiments in Section 3.4 we will consider only pointwise evaluations.

### 3.2.2 Kernel Approximation of the Forward and Inverse Problem

In this subsection we introduce a kernel approximation  $\mathcal{L}_\varepsilon^k$  to the operator  $\mathcal{L}^k$ . Instead of directly discretizing the differential operators on  $\mathcal{M}$ , the new kernel operator is defined by an integral that can be discretized by Monte-Carlo integration as described in the next subsection. Our kernel approximation is inspired by the following construction and result found in Coifman and Lafon [2006].

Let

$$G_\varepsilon u(x) := \varepsilon^{-\frac{m}{2}} \int_{\mathcal{M}} h\left(\frac{|x - \tilde{x}|^2}{\varepsilon}\right) u(\tilde{x}) dV(\tilde{x}), \quad h(z) := \frac{1}{\sqrt{4\pi}} e^{-\frac{z}{4}},$$

where  $dV$  denotes the volume form inherited by  $\mathcal{M}$  from the ambient space  $\mathbb{R}^d$ . Then Lemma 8 in Coifman and Lafon [2006] shows that, for  $u$  sufficiently smooth,

$$G_\varepsilon u(x) = u(x) + \varepsilon(\omega u(x) - \Delta_{\mathcal{M}} u(x)) + O(\varepsilon^2), \quad x \in \mathcal{M}. \quad (3.7)$$

Here,  $\Delta_{\mathcal{M}} := -\operatorname{div}(\nabla \cdot)$ , and  $\omega$  is a function that depends only on the embedding of  $\mathcal{M}$ .

Now, note that by direct calculation

$$\mathcal{L}^\kappa u := -\operatorname{div}(\kappa \nabla u) = \sqrt{\kappa} [\Delta_{\mathcal{M}}(u\sqrt{\kappa}) - u\Delta_{\mathcal{M}}\sqrt{\kappa}], \quad (3.8)$$

and that the expansion (3.7) for  $\sqrt{\kappa}$  and  $u\sqrt{\kappa}$  yields

$$\begin{aligned} uG_\varepsilon\sqrt{\kappa} &= u\sqrt{\kappa} + \varepsilon(\omega u\sqrt{\kappa} - u\Delta_{\mathcal{M}}\sqrt{\kappa}) + O(\varepsilon^2), \\ G_\varepsilon(u\sqrt{\kappa}) &= u\sqrt{\kappa} + \varepsilon(\omega u\sqrt{\kappa} - \Delta_{\mathcal{M}}(u\sqrt{\kappa})) + O(\varepsilon^2). \end{aligned}$$

Subtracting both equations and using (3.8) gives that

$$uG_\varepsilon\sqrt{\kappa} - G_\varepsilon(u\sqrt{\kappa}) = \varepsilon [\Delta_{\mathcal{M}}(u\sqrt{\kappa}) - u\Delta_{\mathcal{M}}\sqrt{\kappa}] + O(\varepsilon^2) = \frac{\varepsilon}{\sqrt{\kappa}} \mathcal{L}^\kappa u + O(\varepsilon^2).$$

This equation motivates the following definition of the integral operator  $\mathcal{L}_\varepsilon^\kappa$

$$\begin{aligned} \mathcal{L}_\varepsilon^\kappa u(x) &:= \frac{\sqrt{\kappa(x)}}{\varepsilon} [u(x)G_\varepsilon\sqrt{\kappa}(x) - G_\varepsilon(u(x)\sqrt{\kappa}(x))] \\ &= \frac{1}{\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int_{\mathcal{M}} \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) \sqrt{\kappa(x)\kappa(\tilde{x})} [u(x) - u(\tilde{x})] dV(\tilde{x}), \end{aligned}$$

which satisfies

$$\mathcal{L}_\varepsilon^\kappa u(x) = \mathcal{L}^\kappa u(x) + \mathcal{O}(\varepsilon), \quad x \in \mathcal{M}.$$

We will make rigorous this formal derivation in Section 3.3.

We next consider the following analogue to equation (3.1), defined by replacing the differential operator  $\mathcal{L}^\kappa$  with the kernel approximation  $\mathcal{L}_\varepsilon^\kappa$ :

$$\mathcal{L}_\varepsilon^\kappa u_\varepsilon = f, \quad x \in \mathcal{M}. \quad (3.9)$$

Lemma 3.3.2 below guarantees the existence of a unique weak solution  $u_\varepsilon \in L_0^2$  to equation

(3.9) provided that  $f \in L^2$  and that the original PDE is uniformly elliptic. In other words, the solution  $u_\varepsilon$  satisfies

$$\int_{\mathcal{M}} \mathcal{L}_\varepsilon^\kappa u_\varepsilon v = \int_{\mathcal{M}} f v, \quad \forall v \in L_0^2. \quad (3.10)$$

We define  $\mathcal{F}_\varepsilon$  as the map that associates  $\theta = \log(\kappa)$  to the solution  $u_\varepsilon$  to (3.9). Denoting  $\mathcal{G}_\varepsilon = \mathcal{D} \circ \mathcal{F}_\varepsilon$ , the approximate posterior  $\pi_\varepsilon^y$  has the following form

$$\frac{d\pi_\varepsilon^y}{d\pi}(\theta) \propto \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right). \quad (3.11)$$

In Section 3.3 we will establish a bound on the total variation distance between the posterior distribution  $\pi^y$  defined in equation (3.3) and its approximation  $\pi_\varepsilon^y$ . We note, however, that the sample-based discretization of the kernel operator  $\mathcal{L}_\varepsilon^\kappa$ —that we will introduce in the next subsection—will involve another layer of approximation not accounted for by the theory in Section 3.3, but necessary in practice.

**Remark 3.2.2.** *As will be seen in Section 3.3, a weak solution to equation (3.9) is sufficient for all the results to hold. We remark that one can show, using Fredholm alternative, the existence of a unique mean zero strong solution with the additional condition that  $f$  has mean zero.*

### 3.2.3 Kernel-Based Elliptic Inverse Problem on a Point Cloud

In this subsection we assume that we are given a point cloud  $X = \{x_1, \dots, x_n\}$ , sampled independently according to an unknown density  $q$  on  $\mathcal{M}$ , but that  $\mathcal{M}$  is otherwise unknown. In applications,  $x_i$  may represent landmarks on the underlying manifold, that may correspond to sensor locations. We consider the inverse problem of determining the value of the unknown input parameter  $\kappa$  at the points  $x_i \in \mathcal{M}$  given the observed data  $y$ . Again we will follow a Bayesian perspective, defining a suitable prior  $\pi_n$  over functions on the point cloud, as well



as a sample-based approximation to the composition map  $\mathcal{G}_\varepsilon = \mathcal{D} \circ \mathcal{F}_\varepsilon$ . We discuss the priors in Section 3.2.3.1 and the approximation to  $\mathcal{G}_\varepsilon$  in Section 3.2.3.2.

### 3.2.3.1 Prior on Point Cloud Functions

We now present the choice of priors that we will use for our numerical experiments in Section 3.4. These will be defined in analogy to (3.4), replacing  $\Delta_{\mathcal{M}}$  by a graph Laplacian. More explicitly, given  $n$  points  $x_1, \dots, x_n$ , we set the prior to be

$$\pi_n = \mathcal{N}(0, C_{\tau,s}^n), \quad C_{\tau,s}^n = c_n(\tau)(\tau I + \Delta_n)^{-s}, \quad (3.12)$$

where  $\Delta_n \in \mathbb{R}^{n \times n}$  is a graph Laplacian constructed with  $x_1, \dots, x_n$  and  $c_n(\tau)$  is a normalizing constant. We refer to Von Luxburg [2007] for a detailed account of graph Laplacians and to Sanz-Alonso and Yang [2022a] for extensive theoretical and computational motivation for our choice of priors. Note that draws from  $\pi_n$  are functions defined intrinsically in the point cloud  $\mathcal{M}_n$  rather than on the (unknown) manifold  $\mathcal{M}$ . The two parameters  $\tau$  and  $s$  play the same role as discussed above in equation (3.5). Again samples from  $\pi_n$  can be expressed by Karhunen-Loève expansion,

$$v_n = c_n(\tau)^{1/2} \sum_{i=1}^n (\tau + \lambda_i^{(n)})^{-s} \xi_i \phi_i^{(n)},$$

where  $\{(\lambda_i^{(n)}, \phi_i^{(n)})\}_{i=1}^n$  are the eigenvalue-eigenvector pairs for  $\Delta_n$  and  $\xi_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ . Similarly as in equation (3.6), we normalize the draws so that the variance per node is 1:

$$c_n(\tau) = \frac{n}{\sum_{i=1}^n (\tau + \lambda_i^{(n)})^{-s}}.$$

For practical considerations, we advocate to set  $\Delta_n$  as the self-tuning graph Laplacian proposed in Zelnik-Manor and Perona [2005]. To illustrate the idea, let  $X = \{x_1, \dots, x_n\}$

be the given point cloud. Then the symmetric graph Laplacian is constructed as the matrix

$$\Delta_n = I - A^{-1/2}SA^{-1/2}, \quad (3.13)$$

where  $S \in \mathbb{R}^{n \times n}$  is a similarity matrix and  $A$  is a diagonal matrix with entries  $A_{ii} = \sum_{j=1}^n S_{ij}$ . We set the entries of the similarity matrix  $S$  to be

$$S_{ij} = \exp\left(-\frac{|x_i - x_j|^2}{2d(i)d(j)}\right),$$

where  $d(i)$  is the distance from  $x_i$  to its  $k$ -th nearest neighbor, and  $k$  is a tunable parameter. The idea is similar to the standard Gaussian similarities except that the local bandwidth parameter is allowed to change adaptively based on the density of points  $x_i$ 's. Moreover, the bandwidth parameter is specified through  $k$ , a positive integer, which can be easily tuned empirically.

### 3.2.3.2 Posterior on Point Cloud Functions

In this subsection we discuss how to discretize the posterior by constructing a point cloud approximation of  $\mathcal{G}_\varepsilon$ . We first approximate  $\mathcal{L}_\varepsilon^\kappa$  by discretizing the integral

$$\mathcal{I}u(x) := \int_{\mathcal{M}} \exp\left(-\frac{|x - \tilde{x}|^2}{4\varepsilon}\right) \sqrt{\kappa(\tilde{x})}u(\tilde{x})dV(\tilde{x})$$

by a Monte-Carlo sum with a reweighting which employs a kernel density estimation. Precisely, we have

$$\mathcal{I}u(x_i) \approx \frac{1}{n} \sum_{j=1}^n \exp\left(-\frac{|x_i - x_j|^2}{4\varepsilon}\right) \sqrt{\kappa(x_j)}u(x_j)q_\varepsilon(x_j)^{-1}, \quad (3.14)$$

where the approximate density, applying (3.7) that  $G_\varepsilon q \approx q$  up to an error of order  $\varepsilon$ , is given by

$$q_\varepsilon(x_j) = \frac{1}{\sqrt{4\pi n\varepsilon^{\frac{m}{2}}}} \sum_{k=1}^n \exp\left(-\frac{|x_j - x_k|^2}{4\varepsilon}\right).$$

The approximation in (3.14) can be interpreted as combining a kernel density estimation [Wasserman, 2006] with importance sampling [Agapiou et al., 2017, Sanz-Alonso, 2018]. In Section 3.4 we will use this observation, where the point clouds come from uniform grids. Then  $\mathcal{L}_\varepsilon^\kappa u$  evaluated at the point cloud is approximated by

$$\begin{aligned} \mathcal{L}_\varepsilon^\kappa u(x_i) &\approx \frac{1}{\sqrt{4\pi n\varepsilon^{\frac{m}{2}+1}}} \sum_{j=1}^n \exp\left(-\frac{|x_i - x_j|^2}{4\varepsilon}\right) \sqrt{\kappa(x_i)\kappa(x_j)} q_\varepsilon(x_j)^{-1} [u(x_i) - u(x_j)] \\ &:= L_{\varepsilon,n}^\kappa u(x_i). \end{aligned} \tag{3.15}$$

More concisely, we can write  $L_{\varepsilon,n}^\kappa$  in matrix form in a series of steps. Define  $P$  to be the kernel matrix with entries  $P_{ij} = \exp(-|x_i - x_j|^2/4\varepsilon)$ . Let  $Q$  be a vector with entries  $Q_i = \sum_{j=1}^n P_{ij}$  and define  $W$  to be the matrix with entries  $W_{ij} = P_{ij} \sqrt{\kappa(x_i)\kappa(x_j)} Q_j^{-1}$ . Then we have

$$L_{\varepsilon,n}^\kappa = \frac{1}{\varepsilon} (D - W), \tag{3.16}$$

where  $D$  is a diagonal matrix with entry  $D_{ii} = \sum_{j=1}^n W_{ij}$ . Notice that the above construction resembles that of the unnormalized graph Laplacian. Indeed, if  $\kappa \equiv 1$ , then (3.16) is exactly the unnormalized graph Laplacian up to a factor of the density [Coifman and Lafon, 2006].

Given the above discretization, we consider the following analogue to equation (3.9), by replacing  $\mathcal{L}_\varepsilon^\kappa$  with  $L_{\varepsilon,n}^\kappa$  and restricting  $f$  to the point cloud:

$$L_{\varepsilon,n}^\kappa u_n = f_n, \tag{3.17}$$

where  $f_n$  is the  $n$ -dimensional vector with entries  $f(x_i)$ , or an approximation thereof when  $f$  is not smooth. One can see from (3.16) that  $L_{\varepsilon,n}^{\kappa}$  is self-adjoint and positive semi-definite under the weighted inner product  $\langle u, v \rangle_q := \frac{1}{n} \sum_{j=1}^n u(x_j)v(x_j)q_{\varepsilon}(x_j)^{-1}$ . Hence  $L_{\varepsilon,n}^{\kappa}$  admits a nonnegative spectrum  $\{\lambda_i\}_{i=1}^n$  with  $\lambda_1 = 0$  and an orthonormal basis of eigenfunctions  $\{v_i\}_{i=1}^n$  with respect to  $\langle \cdot, \cdot \rangle_q$ , with  $v_1 \equiv 1$ . We then set the solution to be

$$u_n = \sum_{i=2}^n \frac{f_n^i}{\lambda_i} v_i, \quad (3.18)$$

where the  $f_n = \sum_{i=1}^n f_n^i v_i$ . Notice that the mean zero condition of  $u$  translates into  $\langle u, 1 \rangle_q = 0$ , taking into account the density. By the orthogonality of the  $v_i$ 's, we see that the solution  $u_n$  in (3.18) satisfies  $\langle u_n, 1 \rangle_q = 0$  and moreover,  $\{v_2, \dots, v_n\}$  forms a basis for  $\ell_0^2 = \{v : \langle v, 1 \rangle_q = 0\}$ , which is the discrete analogue of  $L_0^2$ . One can also check that  $u_n$  satisfies  $\langle L_{\varepsilon,n}^{\kappa} u_n, v \rangle_q = \langle f, v \rangle_q$  for all  $v \in \ell_0^2$ , is consistent with equation (3.10). We remark that if in addition  $\langle f, 1 \rangle_q = 0$ , then  $u_n$  given by equation (3.18) is a strong solution of equation (3.17), in analogy to Remark 3.2.2.

Hence we can now define the discrete forward map  $F_{\varepsilon,n} : \mathbb{R}^n \mapsto \mathbb{R}^n$  as the map that associates  $\theta_n = \log(\kappa_n) := (\log(\kappa(x_1)), \dots, \log(\kappa(x_n)))$  to the solution  $u_n$ . Approximating the pointwise observation map is straightforward. We may also approximate the smoothed observation map introduced in Section 3.2.1.2 by Monte-Carlo as follows:

$$\ell_j^{(n)}(u_n) = \frac{1}{n} \sum_{k=1}^n K(x_j, x_k) u_n(x_k) q_{\varepsilon}(x_k)^{-1}.$$

In either case, denoting  $D_n(u_n) = (\ell_1^{(n)}(u_n), \dots, \ell_J^{(n)}(u_n))^T$  and  $G_{\varepsilon,n} = D_n \circ F_{\varepsilon,n}$ , the graph posterior has the following form

$$\frac{d\pi_{\varepsilon,n}^y}{d\pi_n}(\theta_n) \propto \exp\left(-\frac{1}{2} |y - G_{\varepsilon,n}(\theta_n)|_{\Gamma}^2\right). \quad (3.19)$$

A full analysis of the convergence of the sample-based posteriors  $\pi_{\varepsilon,n}^y$  to the ground-truth posterior  $\pi^y$  is beyond the scope of this chapter. For a *linear* regression problem, the convergence of such graph-based posteriors has been established in García Trillos et al. [2020b] and García Trillos and Sanz-Alonso [2018] using spectral graph theory and variational techniques.

### 3.3 Analysis of Kernel Approximation to the Forward and Inverse Problem

In this section we study the error incurred by replacing the differential operator in the forward map by its kernel approximation, and the effect of such error in the posterior solution to the Bayesian inverse problem. The approximation of the forward map is analyzed in Section 3.3.1 and the approximation of the posterior in Section 3.3.2.

#### 3.3.1 Forward Map Approximation

The main result of this subsection is the following theorem which bounds the difference between the solution to the PDE (3.1) and the solution to the kernel-based equation (3.9).

**Theorem 3.3.1** (Forward map approximation). *Suppose that  $f \in \mathcal{C}^{3,\alpha}$  and  $\kappa \in \mathcal{C}^4$ , with  $\kappa$  bounded below by  $\kappa_{\min} > 0$ . Let  $u$  solve  $\mathcal{L}^\kappa u = f$ , and  $u_\varepsilon$  solve  $\mathcal{L}_\varepsilon^\kappa u_\varepsilon = f$  weakly in  $L_0^2$ . Then for  $\frac{1}{4} < \beta < \frac{1}{2}$  and  $\varepsilon$  small enough depending on  $\beta$ ,*

$$\|u - u_\varepsilon\|_{L^2} \leq CA(\kappa)\|f\|_{H^3}\varepsilon^{4\beta-1},$$

where  $C$  is a constant depending only on  $\mathcal{M}$  and

$$A(\kappa) = \sqrt{\kappa_{\min}^{-5} + \kappa_{\min}^{-6}(\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3}) + \kappa_{\min}^{-7}(\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^2 + \kappa_{\min}^{-8}(\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^3 \|\sqrt{\kappa}\|_{\mathcal{C}^4}^2}.$$

The novelty is to generalize previous analysis in Coifman and Lafon [2006], Gilani and Harlim [2019] to the case of anisotropic diffusions and, more importantly, to keep track of the dependence  $A(\kappa)$  of the error bound on the diffusion coefficient  $\kappa$ . As we will show in Section 3.3.2, understanding this dependence is a crucial ingredient in establishing an approximation result for the inverse problem.

The proof of Theorem 3.3.1 follows the classical numerical analysis argument of combining stability and consistency, coupled with an  $H^4$  norm estimate for solutions to PDE (3.1). Lemma 3.3.2 below establishes the stability of solutions to the kernel-based equation (3.9), Lemma 3.3.3 shows consistency, and Lemma 3.3.5 shows an  $H^4$  norm bound on solutions to (3.1). The proof of Theorem 3.3.1 will be given at the end of this subsection by combining these three lemmas. To streamline the presentation we postpone the proofs of the lemmas to the Appendix.

**Lemma 3.3.2** (Stability). *The equation  $\mathcal{L}_\varepsilon^\kappa u_\varepsilon = f$ , with  $f \in L^2$  and  $\kappa \in L^2$  satisfying  $\kappa(x) \geq \kappa_{\min}$  for a.e.  $x \in \mathcal{M}$  has a unique weak solution  $u_\varepsilon \in L^2_0$ . Moreover, there is  $C > 0$  independent of  $\varepsilon$  and  $\kappa$  such that*

$$\|u_\varepsilon\|_{L^2} \leq C \kappa_{\min}^{-1} \|f\|_{L^2}. \quad (3.20)$$

The next lemma makes rigorous the argument in Section 3.2.2 and characterizes the error between  $\mathcal{L}^\kappa$  and  $\mathcal{L}_\varepsilon^\kappa$  by accounting for its dependence on  $\kappa$ .

**Lemma 3.3.3** (Consistency). *Let  $u \in \mathcal{C}^4$  and  $\kappa \in \mathcal{C}^4$ . Then, for  $\frac{1}{4} < \beta < \frac{1}{2}$  and  $\varepsilon$  sufficiently small depending on  $\beta$ , we have*

$$\|(\mathcal{L}_\varepsilon^\kappa - \mathcal{L}^\kappa)u\|_{L^2} \leq C \|u\|_{H^4} \|\sqrt{\kappa}\|_{\mathcal{C}^4}^2 \varepsilon^{4\beta-1}.$$

**Remark 3.3.4.** *In the proof of Lemma 3.3.3, found in the Appendix, we cannot set  $\beta = \frac{1}{2}$ . However we can choose  $\beta$  arbitrarily close to  $\frac{1}{2}$  so that the rate is essentially  $O(\varepsilon)$ . We*

remark that the proof of Lemma 3.3.3 suggests that the  $\mathcal{C}^3$  assumption in Coifman and Lafon [2006] may not be sufficient.

The last lemma bounds the  $H^4$  norm of the solution to equation (3.1) in terms of the diffusion coefficient  $\kappa$ .

**Lemma 3.3.5** ( $H^4$ -norm bound). *Suppose that  $\kappa \in C^4$  and  $f \in C^{3,\alpha}$  with  $0 < \alpha < 1$ . Let  $u \in \mathcal{C}^5$  be the zero-mean solution to the equation  $\mathcal{L}^\kappa u = f$ . Then*

$$\|u\|_{H^4}^2 \leq C \|f\|_{H^3}^2 \left[ \kappa_{\min}^{-5} + \kappa_{\min}^{-6} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3}) + \kappa_{\min}^{-7} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^2 + \kappa_{\min}^{-8} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^3 \right],$$

where  $C$  is a constant that depends only on  $\mathcal{M}$ .

We are now ready to prove Theorem 3.3.1.

*Proof of Theorem 3.3.1.* Recall that we need to show that

$$\|u - u_\varepsilon\|_{L^2} \leq CA(\kappa) \|f\|_{H^3} \varepsilon^{4\beta-1},$$

where  $u$  solves  $\mathcal{L}^\kappa u = f$  and  $u_\varepsilon$  solves  $\mathcal{L}^\kappa u_\varepsilon = f$  weakly, and  $A(\kappa)$  is defined in the statement of Theorem 3.3.1. Notice that in the weak sense

$$\mathcal{L}_\varepsilon^\kappa(u - u_\varepsilon) = \mathcal{L}_\varepsilon^\kappa u - f = \mathcal{L}_\varepsilon^\kappa u - \mathcal{L}^\kappa u.$$

Hence using Lemma 3.3.2 for the first inequality, and Lemma 3.3.3 for the second one noting that  $f \in \mathcal{C}^{3,\alpha}$  implies that  $u \in \mathcal{C}^5$  Gilbarg and Trudinger [2015], we have

$$\|u - u_\varepsilon\|_{L^2} \leq C \kappa_{\min}^{-1} \|(\mathcal{L}_\varepsilon^\kappa - \mathcal{L}^\kappa)u\|_{L^2} \leq C \kappa_{\min}^{-1} \|u\|_{H^4} \|\sqrt{\kappa}\|_{\mathcal{C}^4}^2 \varepsilon^{4\beta-1}.$$

The result follows by combining this inequality with the bound on  $\|u\|_{H^4}$  derived in Lemma 3.3.5. □

### 3.3.2 Posterior Approximation

In this subsection we characterize the total variation distance between the two posteriors:

$$\begin{aligned} \frac{d\pi^y}{d\pi}(\theta) &= \frac{1}{Z} \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_{\Gamma}^2\right), & Z &:= \int \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_{\Gamma}^2\right) d\pi(\theta), \\ \frac{d\pi_{\varepsilon}^y}{d\pi}(\theta) &= \frac{1}{Z_{\varepsilon}} \exp\left(-\frac{1}{2}|y - \mathcal{G}_{\varepsilon}(\theta)|_{\Gamma}^2\right), & Z_{\varepsilon} &:= \int \exp\left(-\frac{1}{2}|y - \mathcal{G}_{\varepsilon}(\theta)|_{\Gamma}^2\right) d\pi(\theta), \end{aligned}$$

where  $Z$  and  $Z_{\varepsilon}$  are normalizing constants and recall that  $\mathcal{G}(\theta) = (\ell_1(u), \dots, \ell_J(u))^T$  and  $\mathcal{G}_{\varepsilon}(\theta) = (\ell_1(u_{\varepsilon}), \dots, \ell_J(u_{\varepsilon}))^T$  where  $\ell_j$ 's are bounded linear functionals on  $L^2$ .

The main result is Theorem 3.3.6 below. Its proof relies on Theorem 3.3.1 and a standard argument [Stuart, 2010, García Trillos and Sanz-Alonso, 2017, Sanz-Alonso et al., 2019] for the analysis of approximations of Bayesian inverse problems. In particular, the proof makes use of the integrability of the function  $A(\kappa)$  defined in Theorem 3.3.1 with respect to the prior  $\pi$ , guaranteed by Fernique's theorem [Fernique, 1970].

**Theorem 3.3.6** (Posterior approximation). *Let  $\pi$  be a Gaussian measure on  $\mathcal{C}^4$ , and suppose that  $f \in C^{3,\alpha}$  for  $0 < \alpha < 1$ . Then for any  $\frac{1}{4} < \beta < \frac{1}{2}$  and  $\varepsilon$  sufficiently small depending on  $\beta$ ,*

$$d_{\text{TV}}(\pi^y, \pi_{\varepsilon}^y) \leq C\varepsilon^{4\beta-1},$$

where  $C$  is constant depending only on  $\mathcal{M}$ .

The proof of Theorem 3.3.6 relies on the following lemma, whose proof can be found in the Appendix by making use of the integrability of  $A(\kappa)$  with respect to the prior.

**Lemma 3.3.7.** *For  $\frac{1}{4} < \beta < \frac{1}{2}$  and  $\varepsilon$  small enough depending on  $\beta$ , we have*

$$\int \left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_{\varepsilon}(\theta)|_{\Gamma}^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_{\Gamma}^2\right) \right| d\pi(\theta) \leq C\varepsilon^{4\beta-1},$$



where  $C$  is independent of  $\varepsilon$ .

*Proof of Theorem 3.3.6.* We have

$$\begin{aligned}
d_{\text{TV}}(\pi^y, \pi_\varepsilon^y) &= \int \left| \frac{1}{Z_\varepsilon} \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) - \frac{1}{Z} \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right) \right| d\pi(\theta) \\
&\leq \int \left| \frac{1}{Z_\varepsilon} - \frac{1}{Z} \right| \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) d\pi(\theta) \\
&\quad + \int \frac{1}{Z} \left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right) \right| d\pi(\theta) \\
&= \frac{|Z - Z_\varepsilon|}{Z} + \int \frac{1}{Z} \left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right) \right| d\pi(\theta) \\
&\leq \frac{2}{Z} \int \left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right) \right| d\pi(\theta).
\end{aligned}$$

Then using Lemma 3.3.7 it follows that

$$d_{\text{TV}}(\pi^y, \pi_\varepsilon^y) \leq C\varepsilon^{4\beta-1},$$

where  $C$  is independent of  $\varepsilon$ . □

**Remark 3.3.8.** Similarly as for Lemma 3.3.3 our proof fails when  $\beta = \frac{1}{2}$ . However one can choose  $\beta$  arbitrarily close to  $\frac{1}{2}$  so that the rate in 3.3.6 is essentially  $O(\varepsilon)$ .

### 3.4 Numerical Experiments

In this section we investigate numerically the point cloud formulation of the inverse problem introduced in Section 3.2.3. We start in Section 3.4.1 by considering some aspects of the implementation. Then in Sections 3.4.2, 3.4.3, and 3.4.4 we give three numerical examples, where the underlying manifold is chosen to be an ellipse, the torus, and a cow-shaped manifold. In Section 3.4.5, we study a hierarchical approach where the prior lengthscale parameter is learned from data.

### 3.4.1 Implementation

When it comes to practical applications, care must be taken when one chooses the parameters. Central in our kernel method is the parameter  $\varepsilon$ . While Theorem 3.3.1 characterizes the error in approximating  $\mathcal{L}^\kappa$  with  $\mathcal{L}_\varepsilon^\kappa$  and suggests the consistency of the estimator as  $\varepsilon$  goes to 0, in practice one cannot take  $\varepsilon$  too small as we explain now. One can indeed establish the consistency of  $L_{\varepsilon,n}^\kappa$  with  $\mathcal{L}_\varepsilon^\kappa$ , using the same discrete estimation technique as in Berry and Harlim [2016], Berry and Sauer [2016], Gilani and Harlim [2019]. We should point out that while the resulting discrete error bound in Berry and Harlim [2016], Berry and Sauer [2016], Gilani and Harlim [2019] does not show any dependence on  $\kappa$ , which is needed for proving the convergence of the discrete posterior density estimate in (3.19) to (3.11), this result is sufficient for understanding the consistency of  $L_{\varepsilon,n}^\kappa$  with  $\mathcal{L}_\varepsilon^\kappa$ . Specifically, for point cloud with distribution characterized by density  $q(x)$ , defined with respect to the volume form inherited by  $\mathcal{M} \subset \mathbb{R}^d$ , the discrete estimate for fixed-bandwidth Gaussian kernel (e.g., see Corollary 1 of Berry and Harlim [2016] with  $\alpha = 1/2, \beta = 0$  in their setup) states that the sampling error for obtaining an order- $\varepsilon^2$  of the density  $q$  with  $q_\varepsilon$  is of order  $\mathcal{O}(q(x_i)^{1/2}n^{-1/2}\varepsilon^{-(2+m/4)})$  and the error between  $L_{\varepsilon,n}^\kappa$  and  $\mathcal{L}_\varepsilon^\kappa$  is of order  $n^{-1/2}\varepsilon^{-(1/2+m/4)}$ . The fact that  $\varepsilon$  appears in the denominator of these estimates suggests that one cannot take  $\varepsilon$  too small in practice and it also implies that  $\varepsilon$  should be adequately scaled with the size of the data,  $n$ . Since a direct use of these estimates requires knowing the constants that depend on the volume of  $\mathcal{M}$  that are difficult to estimate, we instead adopt an automated empirical method for choosing  $\varepsilon$ . Precisely, we follow Coifman et al. [2008] and plot

$$T(\varepsilon) := \sum_{i,j} \exp\left(-\frac{|x_i - x_j|^2}{4\varepsilon}\right)$$

as a function of  $\varepsilon$  and choose  $\varepsilon$  to be in the region where  $\log(T(\varepsilon))$  is approximately linear.

In the following three subsections we demonstrate the local kernel method for solving

inverse problems through three numerical examples. In the first two examples, the embeddings are known and we set the model analytically, i.e., we first choose the ground truth  $\kappa^\dagger$  and  $u^\dagger$ , and then compute the corresponding  $f$  as

$$f = \operatorname{div}(e^\theta \nabla u) = \frac{1}{\sqrt{\det g}} \partial_i \left( e^\theta g^{ij} \sqrt{\det g} \partial_j u \right), \quad (3.21)$$

where  $g$  is the Riemannian metric on  $\mathcal{M}$ . The third example will be an artificial surface where the embedding is unknown. We will then generate the truth using our kernel method. We will use the pCN algorithm to sample from the posterior [Cotter et al., 2010b, Bertozzi et al., 2018, García Trillos et al., 2020b]. This is a Metropolis-Hastings algorithm with proposal mechanism to move from  $u_n$  to  $u_n^*$  given by

$$u_n^* \sim \left(1 - c^2\right)^{1/2} u_n + c \xi^{(n)}, \quad \xi \sim \pi_n = \mathcal{N}(0, C_{\tau, s}^n), \quad (3.22)$$

where  $c \in (0, 1)$  is a tuning parameter. Note that if  $u_n \sim \pi_n$  then  $u_n^* \sim \pi_n$  showing that the prior is invariant for this kernel. Moreover, it is not difficult to see that detailed balance holds, and as a consequence the Metropolis-Hastings accept/reject mechanism involves only evaluation of the likelihood function. The advantage of pCN in our setting over a standard random walk or Langevin algorithm is that the rate of convergence of pCN does not deteriorate with  $n$ ; this has been established rigorously for a linear inverse problem in García Trillos et al. [2020b].

**Remark 3.4.1.** *At each iteration of the MCMC algorithm, the forward map involves an eigenvalue decomposition of a different matrix for different  $\theta$ 's as shown in Section 3.2.3.2. Hence large  $n$ 's are not favored for computational purposes and this can be an issue for high dimensional  $\mathcal{M}$ 's where the number of points grow as  $n^m$  if one discretizes each dimension by  $n$ .*

### 3.4.2 One-Dimensional Elliptic Problem on an Unknown Ellipse

In this subsection we take  $\mathcal{M}$  to be an ellipse with semi-major and semi-minor axis of length  $a = 3$  and  $1$ , embedded through

$$\iota(\omega) = (\cos \omega, a \sin \omega)^T, \quad \omega \in [0, 2\pi], \quad (3.23)$$

and the Riemannian metric is

$$g_{11}(\omega) = \sin^2 \omega + a^2 \cos^2 \omega.$$

The truth is set to be

$$\kappa^\dagger = 2 + \cos \omega, \quad u^\dagger = \cos \omega,$$

and right-hand side  $f$  in equation (3.1) is defined through equation (3.21). One can check that both  $u^\dagger$  and  $f$  have mean zero, i.e.,  $\int_0^{2\pi} u^\dagger(\omega) \sqrt{g_{11}(\omega)} d\omega = \int_0^{2\pi} f(\omega) \sqrt{g_{11}(\omega)} d\omega = 0$ . We generate the point cloud  $\{x_1, \dots, x_n\}$  according to (3.23) from a uniform grid of  $\omega$  of size  $n = 400$ . The observations are given as noisy pointwise evaluations at subsets of the point cloud:

$$\ell_j = u(x_j) + \eta_j, \quad j = 1, \dots, J,$$

where  $\eta_j \sim \mathcal{N}(0, \sigma^2)$  are assumed to be independent. We will take  $J = 100, 200, 400$  respectively with noise size varying as  $\sigma = 0.01, 0.005, 0.1$ . As discussed in Section 3.2.3.1, we construct the prior with a self-tuning graph Laplacian, using  $k = 2$  neighbors. We empirically discover that such choice of  $k$  gives the best spectral approximation towards the Laplace-Beltrami operator on the ellipse, which has spectrum  $\{i^2\}_{i=0}^\infty$  with eigenfunctions

$\{\sin(i\omega), \cos(i\omega)\}_{i=0}^{\infty}$ . We also tune empirically the parameters in (3.12) as  $\tau = 0.05$  and  $s = 4$ .

In Figure 3.1, we plot the posterior means as functions of  $\omega \in [0, 2\pi]$  and the 95% credible intervals for different  $\sigma$  and  $J$ 's. While the point cloud Bayesian solution is only defined at the discrete point cloud, to ease the visualization we represent the outcome as continuous functions defined on  $\omega \in [0, 2\pi]$ . We see that the truth is mostly captured in the Bayesian confidence intervals. To quantify the error of reconstruction, we compute the relative  $\ell_2$  distances between the posterior mean  $\bar{\kappa}$  and the truth  $\kappa^\dagger$ . Moreover, we compute the solution  $\bar{u}$  of (3.17) with  $\bar{\kappa}$  and its relative  $\ell_2$  distance to the true solution  $u^\dagger$ . As shown in Table 3.1, the reconstruction error for  $u^\dagger$  is much smaller than the relative noise level defined as  $\sqrt{n}\sigma/\|u^\dagger\|_2$ .

$\sigma$	0.01			0.05			0.1		
$J$	100	200	400	100	200	400	100	200	400
$\frac{\ \bar{\kappa}-\kappa^\dagger\ _2}{\ \kappa^\dagger\ _2}$	0.60%	0.80%	0.62%	2.85%	1.96%	2.18%	5.46%	3.90%	3.45%
$\frac{\ \bar{u}-u^\dagger\ _2}{\ u^\dagger\ _2}$	0.26%	0.23%	0.23%	1.08%	0.83%	0.90%	1.70%	1.37%	1.70%
$\frac{\sqrt{n}\sigma}{\ u^\dagger\ _2}$	1.41%			7.07%			14.14%		

Table 3.1: Relative error of  $\bar{\kappa}$  and  $\bar{u}$  for different noise level,  $\sigma$ 's and number of observations,  $J$ , where  $\bar{\kappa}$  and  $\bar{u}$  are the posteriors means for  $\kappa$  and  $u$  respectively. In the last row, the relative noise level for each  $\sigma$  is reported for diagnostic purposes. Particularly, note that the reconstruction error for  $u^\dagger$  is much smaller than the relative noise level.

**Remark 3.4.2.** *Since our prior is on  $\theta = \log(\kappa)$ , we are actually approximating  $\log(2+\cos\omega)$  with trigonometric functions and hence the truth  $\kappa^\dagger$  is not simply the combination of the first two eigenfunctions in the prior. In other words, although the truth  $\kappa^\dagger$  is in the support of the prior, the fact that its coordinates in the eigenbasis do not decay like that in the expansion (3.5) makes it difficult to reconstruct.*

Regarding Remark 3.4.2, we consider another prior with self-tuning graph constructed with  $k = 0.2n$  points. This new graph Laplacian gives a worse spectral approximation to

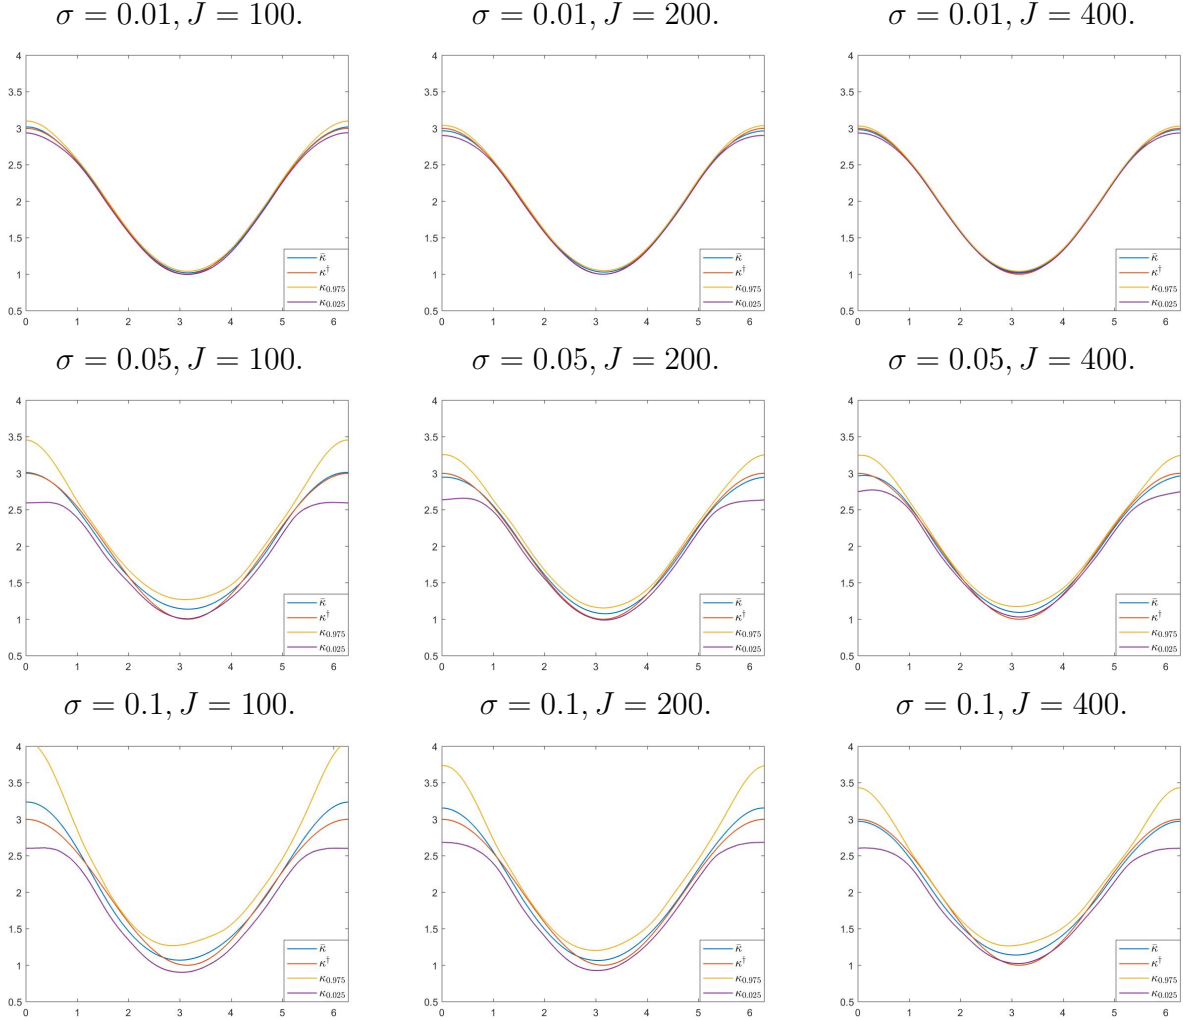


Figure 3.1: Posterior means and 95% credible intervals for different  $\sigma$ 's and  $J$ 's. Here  $\bar{\kappa}$ ,  $\kappa_{0.025}$  and  $\kappa_{0.975}$  represent the posterior mean, 2.5% and 97.5% posterior quantiles respectively.

the Laplace Beltrami operator in the underlying manifold, as its spectrum saturates instead of growing at the appropriate rate. In other words, the basis functions associated with the high frequencies will be given more weight in the expansion (3.5). This can be beneficial in practical applications since it effectively enlarges the support of the prior. Below in Figure 3.2, we solve the inverse problem using this new prior in the case  $\sigma = 0.1$ . The parameters are tuned empirically:  $\tau = 0.75$ ,  $s = 8$ . It can be seen that although the reconstructions are rougher than those in Figure 3.1, they capture better the shape of  $\kappa^\dagger$ , with the help of the higher frequencies. Essentially, larger  $\tau$  (corresponds to more nontrivial modes in the

representation in (3.5)) gives less bias but larger variance, which is consistent with the theory of nonparameteric statistical estimation (e.g, Section 1.7 of Tsybakov [2008]).




Figure 3.2: Posterior means and 95% credible intervals for  $\sigma = 0.1$  and different  $J$ 's. Here  $\kappa_{0.025}$  and  $\kappa_{0.975}$  represent the 2.5% and 97.5% posterior quantiles respectively.

**Remark 3.4.3.** *We note that the inexact reconstruction is partially due to the ill-posedness of the elliptic inverse problem [Mueller and Siltanen, 2012]. As can be seen in Table 3.1, the reconstruction error for  $\bar{\kappa}$  is much larger than that for  $\bar{u}$ : a wide range of  $\kappa$ 's around  $\kappa^\dagger$  give solutions  $u$  which are “close enough” to  $u^\dagger$  (within a range of order  $\sigma$ ) that the algorithm cannot distinguish. When  $\sigma$  is large, such tolerance is larger and the inverse problem becomes more difficult. This issue, together with Remark 3.4.2, explains why one cannot expect exact recovery of  $\kappa^\dagger$  as seen in Figure 3.1.*

### 3.4.3 Two-Dimensional Elliptic Problem on an Unknown Torus

In this subsection we take  $\mathcal{M}$  to be  $\mathbb{T}^2$  embedded in  $\mathbb{R}^3$  through

$$\iota(\omega_1, \omega_2) = ((2 + \cos \omega_1) \cos \omega_2, (2 + \cos \omega_1) \sin \omega_2, \sin \omega_1)^T, \quad \omega_1, \omega_2 \in [0, 2\pi], \quad (3.24)$$

and the Riemannian metric is

$$g(\omega_1, \omega_2) = \begin{bmatrix} 1 & 0 \\ 0 & (2 + \cos \omega_1)^2 \end{bmatrix}.$$

The truth is set to be

$$\kappa^\dagger(\omega_1, \omega_2) = 2 + \sin \omega_1 \sin \omega_2, \quad u^\dagger = \sin \omega_1 \sin \omega_2,$$

and  $f$  is again specified through (3.21). One can check that both  $u^\dagger$  and  $f$  have mean zero, i.e.,  $\int u^\dagger \sqrt{\det g} = \int f \sqrt{\det g} = 0$ . For computational reasons as in Remark 3.4.1, we generate the point cloud according to (3.24) from a  $20 \times 20$  uniform grid on  $[0, 2\pi] \times [0, 2\pi]$  and the observations are given as noisy pointwise evaluations at all points. The graph Laplacian for the prior is constructed with  $k = 16$  neighbors and again we empirically tune the parameters:  $\tau = 0.08$  and  $s = 6$ . Unlike the ellipse case, we cannot get almost exact spectral approximation given that we are only discretizing each dimension by 20 points. The eigenfunctions of the graph Laplacian are wiggly in this case, for which reason we need a large  $s$  to ensure sufficient decay of the spectrum to obtain relatively smooth reconstructions. In Figure 3.3, we plot the posterior means and the standard deviations as functions on  $[0, 2\pi] \times [0, 2\pi]$ ; we note that the uncertainty is large when the function  $\sin \omega_1 \sin \omega_2$  crosses 0. Table 3.2 quantifies the reconstruction error as usual and the reconstruction error for  $u^\dagger$  is again much smaller than the noise level, which are 2%, 10%, and 20% respectively. However, the reconstruction error for  $u^\dagger$  decreases with  $\sigma$  decreasing while the error for  $\kappa^\dagger$  does the opposite, a manifestation of the ill-posedness.

$\sigma$	0.01	0.05	0.1
$\frac{\ \bar{\kappa} - \kappa^\dagger\ _2}{\ \kappa^\dagger\ _2}$	8.56%	8.34%	6.94%
$\frac{\ \bar{u} - u^\dagger\ _2}{\ u^\dagger\ _2}$	1.8%	2.8%	4.8%
$\frac{\sqrt{n}\sigma}{\ u^\dagger\ _2}$	2%	10%	20%

Table 3.2: Relative error of  $\kappa^\dagger$  and  $\bar{u}$  for different noise level,  $\sigma$ 's. In the last row, the relative noise level is reported for diagnostic purpose. Particularly, note that the reconstruction error for  $u^\dagger$  is much smaller than the relative noise level.

**Remark 3.4.4.** For this example we solved equation (3.17) by taking the pseudoinverse  $\hat{u}_n =$



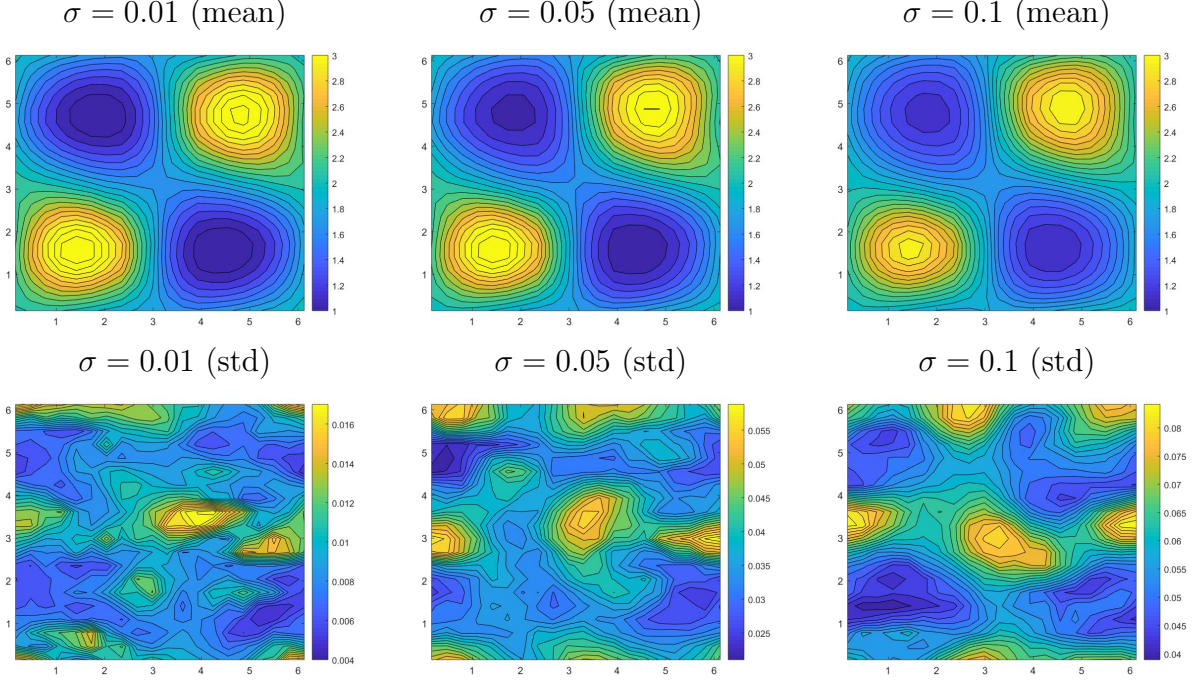


Figure 3.3: Posterior means (first row) and standard deviations (second row) of posteriors for different  $\sigma$ 's. horizontally. We have extended by interpolation the point cloud solution in order to ease visualization.

$(L_{\varepsilon,n}^{\kappa})^{\dagger} f_n$  instead, as this is numerically more stable than taking the eigenvalue decomposition of the asymmetric matrix  $L_{\varepsilon,n}^{\kappa}$ . Moreover,  $\hat{u}_n$  is consistent with  $u^{\dagger}$  for this specific problem as explained below. We have that  $\hat{u}_n$  solves the following problem:

$$\hat{u}_n = \min \{ \|u\|_2 : u \in \operatorname{argmin} \|L_{\varepsilon,n}^{\kappa} u - f_n\|_2 \}. \quad (3.25)$$

The fact that  $f$  has zero mean implies  $\langle f_n, 1 \rangle_q = 0$  in the large  $n$  limit. Then equation (3.17) has a strong solution as mentioned in Section 3.2.3.2 and so the characterization (3.25) implies that  $\hat{u}_n$  is also a strong solution, with  $\sum_{i=1}^n \hat{u}_n^i = 0$ . Notice that the truth  $u^{\dagger}$  also satisfies  $\int u^{\dagger} = 0$  and hence makes  $\hat{u}_n$  consistent.

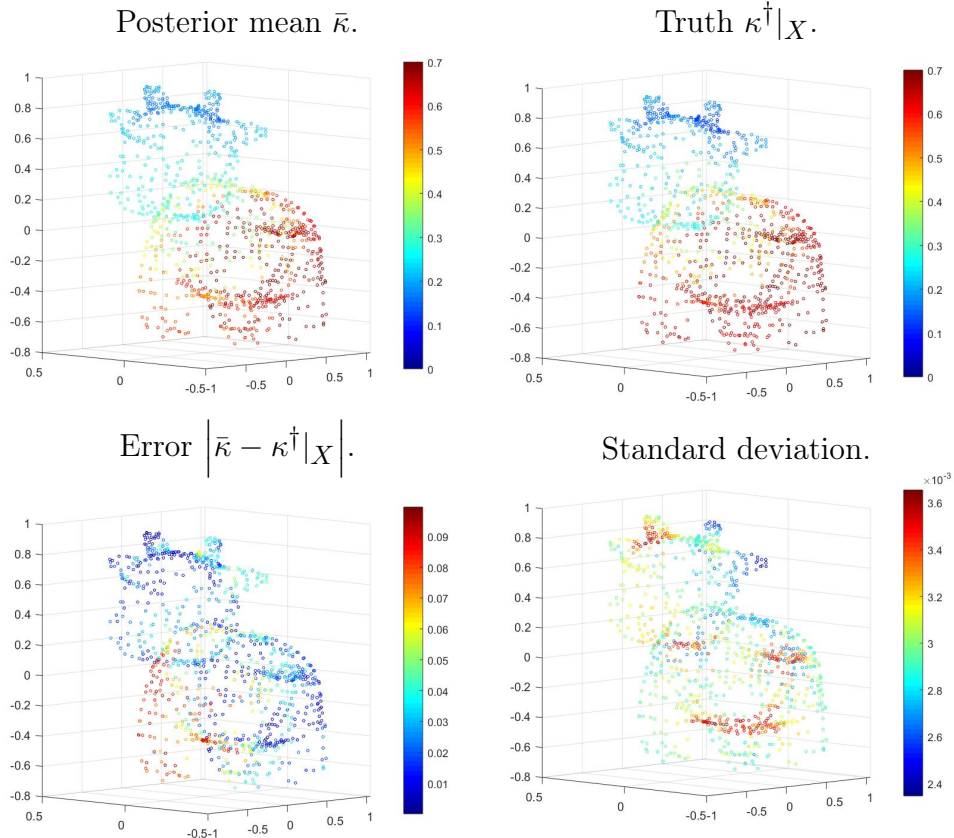


Figure 3.4: Reconstruction for the cow-shaped manifold.

#### 3.4.4 Two-Dimensional Elliptic Problem on an Unknown Artificial Surface

In this subsection we consider an artificial dataset from Keenan Crane’s 3D repository [Crane]. The dataset is made of 2930 points sampled from a cow-shaped surface homeomorphic to the two-dimensional sphere. The purpose of this subsection is to demonstrate that our kernel method can be applied to more complicated manifolds. To avoid an inverse crime [Kaipo and Somersalo, 2006], we generate the truth using all 2930 points but solve the inverse problem on a subset of size  $n = 1000$ .

To be more precise, we generate  $\kappa^\dagger$  from the Gaussian measure  $\mathcal{N}(0, (\tau I + \Delta_{2930}))^{-s}$ , where  $\Delta_{2930}$  is the self-tuning graph Laplacian constructed with  $k = 100$  neighbors and  $\tau = 0.7$ ,  $s = 6$ . We then set  $u^\dagger$  to be  $10(\phi_2 - c)$ , where  $\phi_2$  is the second eigenvector of  $\Delta_{2930}$  and  $c$  is a constant chosen below. The factor 10 in the definition of  $u$  is only to match the

order of magnitude with  $\kappa^\dagger$ . We then set  $f = L_{\varepsilon, 2930}^{\kappa^\dagger} u^\dagger$ . This would serve as our ground truth, and now we solve the inverse problem on a random subset  $X$  of  $n = 1000$  points.

On the point cloud  $X$ , the truth becomes  $\kappa^\dagger|_X$  and  $u^\dagger|_X$ . As mentioned above, the solution  $u^\dagger|_X$  needs to have zero mean to be consistent with our theory. Since we do not have access to the Riemannian metric as in the previous two examples, we instead require  $\langle u^\dagger|_X, 1 \rangle_q = 0$ , which gives the choice of  $c$  above. Again we consider noisy pointwise observations at all 1000 points. The inputs of the problem are now  $f|_X$  and noisy  $u^\dagger|_X$ . The noise level is 10%, which gives  $\sigma = 0.0186$ . The prior that we use has the same parameter  $\tau$  as used to obtain our synthetic truth, but is defined using a graph Laplacian on  $X$ . Namely, we consider

$$\pi_n = \mathcal{N}(0, (0.7I + \Delta_{1000})^{-6}),$$

where  $\Delta_{1000}$  is constructed with  $k = 80$  neighbors. Figure 3.4 shows the plots of the posterior mean, the truth, the error, and the standard deviation. The reconstruction errors are  $\|\bar{\kappa} - \kappa^\dagger|_X\|_2 / \|\kappa^\dagger|_X\|_2 = 9.73\%$  and  $\|\bar{u} - u^\dagger|_X\|_2 / \|u^\dagger|_X\|_2 = 16.24\%$ . We remark that the large relative error for  $u^\dagger$  is partly due to the fact that the point cloud of size 1000 does not approximate the original one well and in particular

$$L_{\varepsilon, 1000}^{\kappa^\dagger|_X} u^\dagger|_X \neq f|_X. \quad (3.26)$$

When we solve for  $\hat{u}$  in equation (3.26), i.e., solving

$$L_{\varepsilon, 1000}^{\kappa^\dagger|_X} \hat{u} = f|_X,$$

we get  $\|\hat{u} - u^\dagger|_X\|_2 / \|u^\dagger|_X\|_2 = 17.01\%$  and this is the best one can hope for in terms of reconstructing  $u^\dagger|_X$ . Hence we see that the above relative error has already reached the

limit of the method.

### 3.4.5 Hierarchical Bayesian Formulation

As mentioned in Section 3.2.1.1, the choice of  $\tau$  is crucial and would require some prior knowledge of the lengthscale of the function to be reconstructed. In this section, we demonstrate how one can learn the parameter  $\tau$  together with  $\kappa$  through a hierarchical Bayesian approach proposed in Dunlop et al. [2017]. We emphasize that the hierarchical approach may not be able to find the precise lengthscale of the parameter to be reconstructed, and hence should only be applied when little prior knowledge on the lengthscale is available. We will only focus on the point-cloud inverse problem as in Section 3.2.3.

We remark that our choice of priors in (3.4) and (3.12) differ from the ones used in Dunlop et al. [2017] by the scaling constants. In the continuum space, the family of measures defined by equation (3.4) are mutually singular, which leads to technical difficulties when designing hierarchical methods. However, in the point cloud setting, the family of measures as in (3.12) are simply multivariate normal and are equivalent. The formulation in Dunlop et al. [2017] then carries over.

The idea is to consider a joint prior on  $(\theta_n, \tau)$  that takes the form

$$\Pi(\theta_n, \tau) = \pi_0(\tau)\pi(\theta_n|\tau) := \pi_0(\tau)\pi_\tau(\theta_n),$$

where  $\pi_0$  is a distribution on  $\mathbb{R}^+$  and the conditional distribution  $\pi_\tau$  is taken as in (3.12).

Recall that  $\pi_\tau$  has the form

$$\pi_\tau = \mathcal{N}(0, C_{\tau,s}^n), \quad C_{\tau,s}^n = c_n(\tau)(\tau I + \Delta_n)^{-s},$$

where  $c_n(\tau)$  normalizes the draws so that  $u \sim \pi_n$  satisfies  $\mathbb{E}|u_n|^2 = n$ . Now we can define the forward map  $F_n : \mathbb{R}^n \times \mathbb{R}^+ \rightarrow \mathbb{R}^n$  that associates a pair  $(\theta_n, \tau)$  with the unique mean-zero

solution  $u_n$  of equation (3.17). We notice that  $F_n$  is essentially the same as before except the additional  $\tau$  that does not play a role in the definition. Denoting  $G_n = D_n \circ F_n$ , the joint posterior  $\Pi^y$  can be written as a change of measure with respect to the prior,

$$\frac{d\Pi^y}{d\Pi}(\theta_n, \tau) \propto \exp\left(-\frac{1}{2}|y - G_n(\theta_n, \tau)|_\Gamma^2\right). \quad (3.27)$$

### 3.4.5.1 Sampling

To sample from the joint posterior  $\theta_n, \tau|y$ , we will use a Metropolis within Gibbs sampling scheme by updating of  $\theta_n|\tau, y$  and  $\tau|\theta_n, y$  alternately. Sampling of  $\theta_n|\tau, y$  reduces to the non-hierarchical setting, where we use the pCN algorithm. Sampling of  $\tau|\theta_n, y$  is more delicate since one needs first to make sense of this conditional distribution. Instead of making the presentation too involved, we will only present the algorithm and refer to Dunlop et al. [2017] for more details. The idea is to use symmetric random-walk Metropolis-Hastings, with acceptance probability to accept the proposal  $\tau$ , given the current chain value  $\gamma$  as,

$$a(\tau, \gamma) = \exp\left(-\frac{1}{2}[H(\tau) - H(\gamma)]\right) \frac{\pi_0(\tau)}{\pi_0(\gamma)} \wedge 1, \quad (3.28)$$

where,

$$H(\tau) = \sum_{i=1}^n \log \lambda_i(\tau) + \frac{\langle \theta, \phi_i \rangle^2}{\lambda_i(\tau)}.$$

Here,  $\{(\lambda_i(\tau), \phi_i)\}_{i=1}^n$  are the eigenvalue-eigenfunction pairs of  $C_{\tau, s}^n$ .

From (3.28), we see that the algorithm favors lengthscales  $\tau$ 's that give small values of  $H(\tau)$ . As  $\tau$  approaches 0, the first eigenvalue of  $(\tau I + \Delta_n)^{-s}$  goes to infinity and so the normalizing constant  $c_n(\tau)$  approaches 0. However, the eigenvalues of  $(\tau I + \Delta_n)^{-s}$ , except the first one, do not change much since  $\tau$  is now a much smaller quantity. Hence when multiplied by  $c_n(\tau)$ , they converge to 0. In other words  $\lambda_i(\tau)$  approaches 0 as  $\tau$  goes to 0 for

$i \geq 2$ , which then implies that  $\sum_{i=1}^n \log \lambda_i(\tau) \rightarrow -\infty$ . So the first term in  $H$  is minimized at  $\tau = 0$ , while the second term  $\langle \theta, \phi_i \rangle^2 / \lambda_i(\tau)$  favors large  $\tau$  by a similar argument as above. Then the minimum of  $H$  is attained by balancing the two sums, using the coefficients  $\langle \theta, \phi_i \rangle$ . Hence the algorithm will give a  $\tau$  that is consistent with these coefficients, reflecting the lengthscale of  $\theta$ .

### 3.4.5.2 Numerical Experiments

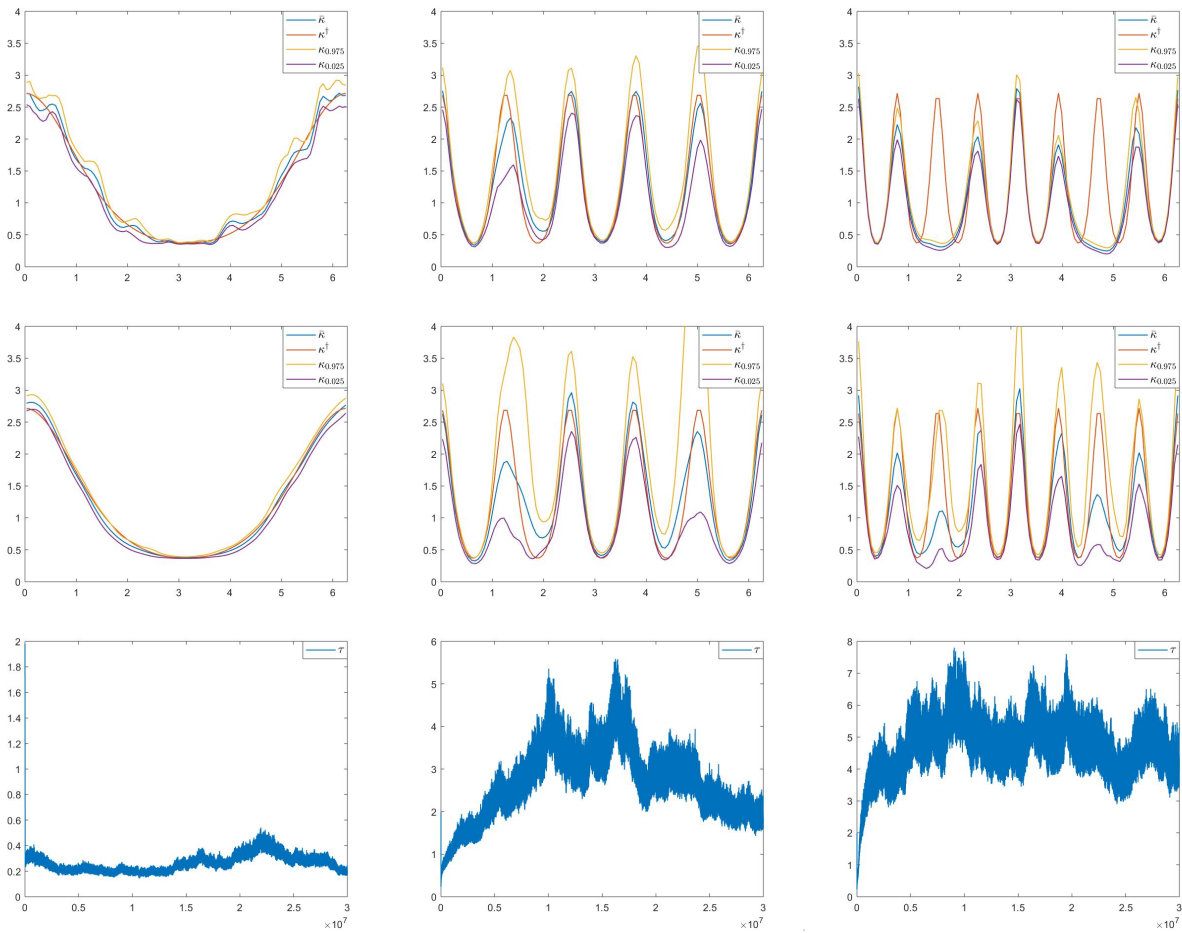


Figure 3.5: Posterior means and 95% credible intervals for different truths. Figure are arranged so that the first two rows correspond to non-hierarchical and hierarchical respectively, and the third row shows the sample paths for  $\tau$ . The three columns represent the truths  $e^{\cos(\omega)}$ ,  $e^{\cos(5\omega)}$ ,  $e^{\cos(8\omega)}$  respectively.

To demonstrate the hierarchical approach, we focus on the ellipse case with three truths

of different lengthscales:  $\kappa^\dagger = e^{\cos(\omega)}, e^{\cos(5\omega)}, e^{\cos(8\omega)}$  with a fixed  $f = \frac{1}{5} \sin \omega$ . We fix  $f$  so that we are solving the same inverse problem, with different underlying truth  $\kappa^\dagger$ 's. In this case we no longer have analytic solutions for  $u$  and we use the MATLAB PDE toolbox. The point cloud is generated as in Section 3.4.2 with  $n = 100$ , with pointwise observations with noise level  $\sigma = 0.01$  at all points. We take  $\pi_0 = \mathcal{N}(2, 1)$  and  $s = 4$ . In the hierarchical setting, it takes longer for the chains to mix, where we run the chain for  $3 \times 10^7$  iterations and use the last  $5 \times 10^6$  samples for computations. We compare the hierarchical approach with the non-hierarchical one, where we use a same  $\tau$  to reconstruct the three  $\kappa^\dagger$ 's. Figure 3.5 below shows the corresponding reconstructions.

The first row corresponds to the non-hierarchical approach where we fix  $\tau = 2$ , which is finely tuned to match the lengthscale of  $e^{\cos(5\omega)}$ , for all three problems. The reconstruction is then acceptable for  $e^{\cos(5\omega)}$  but is poorer for the other two. For  $e^{\cos(\omega)}$ , the reconstruction still fits the shape of the truth, but since the prior now has a lengthscale much larger than that of  $e^{\cos(\omega)}$ , the reconstruction is oscillatory. On the other hand, the reconstruction of  $e^{\cos(8\omega)}$  fails to capture the shape of the truth. This is because the prior now has a smaller lengthscale than that of  $e^{\cos(8\omega)}$ , so that frequencies high as  $\cos(8\omega)$  barely belongs to the prior. The second and third row corresponds to the reconstructions of the hierarchical approach and the corresponding sample paths for  $\tau$ . We see that the credible intervals capture most of the truths and the reconstructions are much better for  $e^{\cos(\omega)}$  and  $e^{\cos(8\omega)}$  than with the non-hierarchical approach. Table 3.3 quantifies the reconstruction error. We notice that the hierarchical approach performs worse than the non-hierarchical one for  $e^{\cos(5\omega)}$ ; this is because we have chosen  $\tau = 2$  agrees with the lengthscale of the true diffusion coefficient. This fact suggests that the hierarchical approach only improves the performance when little prior knowledge on the lengthscale is known. From the sample paths for  $\tau$ 's, we see that the chains have large variance and do not concentrate on a particular value. This is due to the ill-posedness of the inverse problem where  $\kappa$ 's of different lengthscales give equally good

reconstruction of  $u$ , and hence the algorithm cannot distinguish between them. So in general the algorithm may not give the precise lengthscale but a possible range of it.

$\kappa^\dagger$	$e^{\cos(\omega)}$		$e^{\cos(5\omega)}$		$e^{\cos(8\omega)}$	
Method	NH	H	NH	H	NH	H
$\frac{\ \bar{\kappa} - \kappa^\dagger\ _2}{\ \kappa^\dagger\ _2}$	9.07%	3.69%	11.17%	17.41%	39.38%	28.82%
$\frac{\ \bar{u} - u^\dagger\ _2}{\ u^\dagger\ _2}$	0.84%	0.72%	0.73%	0.69%	0.87%	0.83%
$\frac{\sqrt{n\sigma}}{\ u^\dagger\ _2}$	1.87%		1.26%		1.27%	

Table 3.3: Relative error of  $\bar{\kappa}$  and  $\bar{u}$  for different truths. Here “N” and “H” stand for non-hierarchical and hierarchical respectively. In the last row, the relative noise level for each  $\sigma$  is reported for diagnostic purposes.

**Remark 3.4.5.** Notice that in the above the noise level has been set to be small. When the noise level  $\sigma$  is large, the performance of the hierarchical approach may be worse, as shown in Figure 3.6. The reason is that the algorithm sees only the noisy data, which is the truth  $u^\dagger$  perturbed by noise. In other words, the lengthscale of the data is corrupted by the noise, which has lengthscale converging to 0 ( $\tau \rightarrow \infty$ ) in the large  $n$  and  $J$  limit if the noise is independent, i.e.,  $\Gamma = I$ . As shown in Figure 3.6, the chain for  $\tau$  oscillates in a wide range of values, suggesting that the data contains little information on this parameter.




Figure 3.6: Reconstruction of  $e^{\cos(\omega)}$  and sample path for  $\tau$  when  $\sigma = 0.1$ .



### 3.5 Conclusions and Future Work

- This chapter introduced kernel-based methods for the solution of inverse problems on manifolds. We have shown through rigorous analysis that the forward map can be replaced by a kernel approximation while keeping small the total variation distance to the true posterior. Through numerical experiments we have shown that a point cloud discretization to the kernel approximation may allow to implement the inverse problem on point clouds, without reference to the underlying manifold.
- An important question of theoretical interest when solving the inverse problem on the point cloud is how to choose optimally the kernel bandwidth  $\varepsilon$  in terms of the number  $n$  of manifold samples. We conjecture that the convergence of the graph posteriors to the ground truth posterior, and guidelines on the choice of kernel bandwidth, may be established by generalizing the spectral graph theory results in Burago et al. [2015], García Trillos et al. [2019a] to anisotropic diffusion, and using the variational techniques introduced in García Trillos and Sanz-Alonso [2018], García Trillos et al. [2020b]. The analysis of these questions will be the subject of future work.
- We streamlined the presentation by working on a closed manifold with no boundary. We expect that the numerical and theoretical results may be extended to Neumann, Robin, and Dirichlet boundary conditions using the results and ideas in Gilani and Harlim [2019], Li and Shi [2016], Shi and Sun [2016], Thiede et al. [2019].
- The practical success of the Bayesian approach is heavily dependent on the choice of prior. Here, we have used Matern-type priors that are flexible models widely used in spatial statistics and the geophysical sciences [Stein, 1999a]. While the hierarchical approach to the inverse problem [Dunlop et al., 2017, Geoga et al., 2019] is effective for learning the prior lengthscale from data in certain regimes as we have numeri-

cally shown, a more robust algorithm is needed and this merits an extensive further investigation.

- A topic of further research will be the extension of the kernel-based approximation to PDEs and inverse problems to other PDEs and ODEs beyond the elliptic model considered here.

## 3.6 Appendix

### 3.6.1 Proof of Lemma 3.3.2

The proof proceeds by a standard Lax-Milgram argument. Throughout,  $C > 0$  denotes a constant independent of  $\varepsilon$  and  $\kappa$  that may change from line to line. Consider the bilinear and linear functionals

$$\begin{aligned} B : L_0^2 \times L_0^2 &\rightarrow \mathbb{R}, & F : L_0^2 &\rightarrow \mathbb{R}, \\ (u, v) &\mapsto \langle u, \mathcal{L}_\varepsilon^\kappa v \rangle, & v &\mapsto \langle v, f \rangle. \end{aligned}$$

Clearly,  $B$  and  $F$  are bounded. To show that  $B$  is coercive, note that by [Shi and Sun, 2016, Theorem 7.2] there exists  $C > 0$  such that, for all  $\varepsilon > 0$   $v \in L_0^2(\mathcal{M})$ ,

$$\langle v, \mathcal{L}_\varepsilon v \rangle \geq C \|v\|_{L^2}, \tag{3.29}$$

where

$$\mathcal{L}_\varepsilon v := \frac{1}{\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) [v(x) - v(\tilde{x})] dV(\tilde{x}).$$

It follows that, for  $v \in L_0^2$ ,

$$\begin{aligned}
\langle v, \mathcal{L}_\varepsilon^\kappa v \rangle &= \frac{1}{\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int \int \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) \sqrt{\kappa(x)\kappa(\tilde{x})} v(x) [v(x) - v(\tilde{x})] dV(\tilde{x}) dV(x) \\
&= \frac{1}{\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int \int \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) \sqrt{\kappa(x)\kappa(\tilde{x})} v(\tilde{x}) [v(\tilde{x}) - v(x)] dV(\tilde{x}) dV(x) \\
&= \frac{1}{2\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int \int \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) \sqrt{\kappa(x)\kappa(\tilde{x})} |v(x) - v(\tilde{x})|^2 dV(\tilde{x}) dV(x) \\
&\geq \kappa_{\min} \frac{1}{2\sqrt{4\pi\varepsilon^{\frac{m}{2}+1}}} \int \int \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) |v(x) - v(\tilde{x})|^2 dV(\tilde{x}) dV(x) \\
&= \kappa_{\min} \langle v, \mathcal{L}_\varepsilon v \rangle \geq C \kappa_{\min} \|v\|_{L^2}^2,
\end{aligned}$$

establishing the coercivity of  $B$ . The existence and uniqueness of a weak solution, as well as the bound (3.20) follow from the Lax Milgram theorem.

### 3.6.2 Proof of Lemma 3.3.3

Our proof follows the same argument as [Coifman and Lafon, 2006, Lemma 8] but keeps track of the coefficients of the higher order terms. Let

$$G_\varepsilon u(x) = \varepsilon^{-\frac{m}{2}} \int h\left(\frac{|x-\tilde{x}|^2}{\varepsilon}\right) u(\tilde{x}) dV(\tilde{x}), \quad h(z) = \frac{1}{\sqrt{4\pi}} e^{-\frac{z}{4}}.$$

Let  $0 < \beta < \frac{1}{2}$ . We can localize the integration near  $x$  due to the exponential decay of  $e^{-x^2}$ :

$$\begin{aligned}
& \left| \varepsilon^{-\frac{m}{2}} \int_{\tilde{x} \in \mathcal{M}: |\tilde{x}-x| > \varepsilon^\beta} \exp\left(-\frac{|x-\tilde{x}|^2}{4\varepsilon}\right) u(\tilde{x}) dV(\tilde{x}) \right| \\
& \leq \varepsilon^{-\frac{m}{4}} \|u\|_{L^2} \sqrt{\varepsilon^{-\frac{m}{2}} \int_{\tilde{x} \in \mathcal{M}: |\tilde{x}-x| > \varepsilon^\beta} \exp\left(-\frac{|x-\tilde{x}|^2}{2\varepsilon}\right) dV(\tilde{x})} \\
& \leq \varepsilon^{-\frac{m}{4}} \|u\|_{L^2} \sqrt{\int_{x+\sqrt{\varepsilon}\tilde{x} \in \mathcal{M}: |\tilde{x}| > \varepsilon^{\beta-1/2}} \exp\left(-\frac{1}{2}|\tilde{x}|^2\right) dV(\tilde{x})} \\
& \leq C\varepsilon^{-\frac{m}{4}} \|u\|_{L^2} \sqrt{\mathbb{P}\{\mathcal{N}(0,1) > \varepsilon^{\beta-1/2}\}} \\
& \leq C\varepsilon^{-\frac{m}{4}} \|u\|_{L^2} \exp\left(-c\varepsilon^{2\beta-1}\right) \leq C\|u\|_{L^2}\varepsilon^2,
\end{aligned}$$

where in the last inequality since  $2\beta < 1$ ,  $\exp(-c\varepsilon^{2\beta-1})$  decays faster than any polynomial in  $\varepsilon$  and in particular for  $\varepsilon$  small enough it decays faster than  $\varepsilon^{2+\frac{m}{4}}$ . Therefore,

$$G_\varepsilon u(x) = \varepsilon^{-\frac{m}{2}} \int_{\tilde{x} \in \mathcal{M}: |\tilde{x}-x| < \varepsilon^\beta} h\left(\frac{|x-\tilde{x}|^2}{\varepsilon}\right) u(\tilde{x}) dV(\tilde{x}) + O(\|u\|_{L^2}\varepsilon^2).$$

Now we Taylor expand  $u$  near  $x$ . Let  $(s_1, \dots, s_m)$  be the geodesic coordinates at  $x$  and  $u(\tilde{x}) = u(\tilde{x}(s_1, \dots, s_m)) = \tilde{u}(s_1, \dots, s_m) = \tilde{u}(s)$ . Then

$$\begin{aligned}
u(\tilde{x}) - u(x) &= \tilde{u}(s) - \tilde{u}(0) = \sum_{i=1}^m s_i \frac{\partial \tilde{u}}{\partial s_i}(0) + \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m s_i s_j \frac{\partial^2 \tilde{u}}{\partial s_i \partial s_j}(0) \\
&+ \frac{1}{6} \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m s_i s_j s_k \frac{\partial^3 \tilde{u}}{\partial s_i \partial s_j \partial s_k}(0) + \delta(s),
\end{aligned}$$

where

$$\delta(s) = \frac{1}{24} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m \sum_{\ell=1}^m s_i s_j s_k s_\ell \frac{\partial^4 \tilde{u}}{\partial s_i \partial s_j \partial s_k \partial s_\ell}(t_1 t_2 t_3 t_4) dt_1 dt_2 dt_3 dt_4.$$

Then expanding the 4-th order term in  $G_\varepsilon u$ , we have

$$\begin{aligned}
|T_u(x)| &:= \left| \varepsilon^{-\frac{m}{2}} \int_{|s| < \varepsilon^\beta} h\left(\frac{|s|^2}{\varepsilon}\right) \delta(s) ds \right| \\
&\leq \frac{m^4 \varepsilon^{4\beta}}{24} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_{|s| < \varepsilon^\beta} \varepsilon^{-\frac{m}{2}} h\left(\frac{|s|^2}{\varepsilon}\right) \|\nabla^4 \tilde{u}(t_1 t_2 t_3 t_4 s)\| ds dt_1 dt_2 dt_3 dt_4 \\
&= \frac{m^4 \varepsilon^{4\beta}}{24} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} (t_1 t_2 t_3 t_4)^{-d} \varepsilon^{-\frac{d}{2}} h\left(\frac{|r|^2}{t_1^2 t_2^2 t_3^2 t_4^2 \varepsilon}\right) \|\nabla^4 \tilde{u}(r)\| dr dt_1 dt_2 dt_3 dt_4 \\
&:= \frac{m^4 \varepsilon^{4\beta}}{24} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) \|\nabla^4 \tilde{u}(r)\| dr d\tau,
\end{aligned}$$

where we have used the notation  $\tau := (t_1, t_2, t_3, t_4)$  and  $d\tau := dt_1 dt_2 dt_3 dt_4$ . By interchanging the order of integration and noticing that  $\nabla^4 \tilde{u}(r) = \nabla^4 u(x + \xi)$  where  $\xi$ 's are directional vectors that are independent of  $x$ , we have

$$\begin{aligned}
\|T_u\|_{L^2}^2 &\leq \left(\frac{m^4 \varepsilon^{4\beta}}{24}\right)^2 \int_{\mathcal{M}} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) \|\nabla^4 \tilde{u}(r)\| dr \right]^2 d\tau dV(x) \\
&\leq \left(\frac{m^4 \varepsilon^{4\beta}}{24}\right)^2 \int_{\mathcal{M}} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) dr \right] \\
&\quad \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) \|\nabla^4 \tilde{u}(r)\|^2 dr \right] d\tau dV(x) \\
&= \left(\frac{m^4 \varepsilon^{4\beta}}{24}\right)^2 \int_0^1 \int_0^1 \int_0^1 \int_0^1 \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) dr \right] \\
&\quad \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) \int_{\mathcal{M}} \|\nabla^4 u(x + \xi)\|^2 dV(x) dr \right] d\tau \\
&= \left(\frac{m^4 \varepsilon^{4\beta}}{24}\right)^2 \|\nabla^4 u\|_{L^2}^2 \int_0^1 \int_0^1 \int_0^1 \int_0^1 \left[ \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) dr \right]^2 d\tau.
\end{aligned}$$

By a change of variable  $r = t_1 t_2 t_3 t_4 \sqrt{\varepsilon} s$ , we notice that

$$\int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} K(r, \tau) dr = \int_{|r| < t_1 t_2 t_3 t_4 \varepsilon^\beta} (t_1 t_2 t_3 t_4)^{-d} \varepsilon^{-\frac{d}{2}} h\left(\frac{|r|^2}{t_1^2 t_2^2 t_3^2 t_4^2 \varepsilon}\right) dr = \int_{|z| < \varepsilon^{\beta-1/2}} h(|z|^2),$$

which can be bounded by a constant independent of  $\varepsilon$  by the Gaussianity of  $h(|z|^2)$ . Hence

$$\|T_u\|_{L^2} \leq C \|\nabla^4 u\|_{L^2} \varepsilon^{4\beta}, \quad (3.30)$$

where  $C$  is a constant that does not depend on  $u$  or  $\varepsilon$ . Now following the same argument as in Coifman and Lafon [2006] and keeping track of the derivatives of  $u$ , we have

$$G_\varepsilon u(x) = u(x) + \varepsilon [\omega(x)u(x) + \Delta u(x)] + R_u(x),$$

where

$$R_u(x) = T_u(x) + O\left(\|u\|_{L^2} + \|\nabla u(x)\| + \|\nabla^2 u(x)\| + \|\nabla^3 u(x)\|\right) \varepsilon^2.$$

Applying  $G_\varepsilon$  to  $u\sqrt{\kappa}$ , we have

$$G_\varepsilon(u\sqrt{\kappa}) = u\sqrt{\kappa} + \varepsilon [\omega u\sqrt{\kappa} + \Delta(u\sqrt{\kappa})] + R_{u\sqrt{\kappa}}. \quad (3.31)$$

By expanding the derivatives of  $u\sqrt{\kappa}$  and bounding them by the  $\infty$ -norms of  $\kappa$  and its derivatives, we have

$$R_{u\sqrt{\kappa}}(x) = T_{u\sqrt{\kappa}}(x) + O\left(\|\sqrt{\kappa}\|_{C^4} \left[\|u\|_{L^2} + |u(x)| + \|\nabla u(x)\| + \|\nabla^2 u(x)\| + \|\nabla^3 u(x)\|\right] \varepsilon^2\right),$$

where

$$\|T_{u\sqrt{\kappa}}\|_{L^2} \leq C \|\sqrt{\kappa}\|_{C^4} \|\nabla^4 u\|_{L^2} \varepsilon^{4\beta}.$$

By setting  $u = 1$ , we have

$$G_\varepsilon \sqrt{\kappa} = \sqrt{\kappa} + \varepsilon [\omega \sqrt{\kappa} + \Delta \sqrt{\kappa}] + O\left(\|\sqrt{\kappa}\|_{\mathcal{C}^4} \varepsilon^2\right). \quad (3.32)$$

By combining (3.31) and (3.32), we have

$$\mathcal{L}_\varepsilon^\kappa u = \frac{\sqrt{\kappa}}{\varepsilon} [u G_\varepsilon \sqrt{\kappa} - G_\varepsilon(u \sqrt{\kappa})] = \mathcal{L}^\kappa u + O\left(\sqrt{\kappa} u \|\sqrt{\kappa}\|_{\mathcal{C}^4} \varepsilon\right) + \frac{\sqrt{\kappa} R_u \sqrt{\kappa}}{\varepsilon}.$$

Hence it follows that

$$\|(\mathcal{L}_\varepsilon^\kappa - \mathcal{L}^\kappa)u\|_{L^2} \leq C \|u\|_{H^4} \|\sqrt{\kappa}\|_{\mathcal{C}^4}^2 \varepsilon^{4\beta-1}.$$

### 3.6.3 Proof of Lemma 3.3.5

First we multiply the equation by  $u$  and integrate over  $\mathcal{M}$ . Integrating by parts, we get

$$\int f u = - \int \operatorname{div}(\kappa \nabla u) u = \int \kappa |\nabla u|^2 \geq \kappa_{\min} \|\nabla u\|_{L^2}^2.$$

By Hölder and Poincaré inequalities, there is a constant  $C$  that depends only on  $\mathcal{M}$  so that

$$\|\nabla u\|_{L^2} \leq C \kappa_{\min}^{-1} \|f\|_{L^2} \leq C \kappa_{\min}^{-1} \|f\|_{H^3}, \quad (3.33)$$

$$\|u\|_{L^2} \leq C \kappa_{\min}^{-1} \|f\|_{L^2} \leq C \kappa_{\min}^{-1} \|f\|_{H^3}. \quad (3.34)$$

Now differentiating the equation with respect to  $x_k$  and testing against  $u_{x_k}$ , we get

$$\begin{aligned} \int f_{x_k} u_{x_k} &= - \int \operatorname{div} \left( \frac{\partial}{\partial x_k} (\kappa \nabla u) \right) u_{x_k} = \int \frac{\partial}{\partial x_k} (\kappa \nabla u) \cdot \nabla u_{x_k} \\ &= \int \kappa_{x_k} \nabla u \cdot \nabla u_{x_k} + \kappa \nabla u_{x_k} \cdot \nabla u_{x_k}. \end{aligned}$$

Using Young's inequality that  $|ab| \leq \varepsilon a^2 + \frac{1}{4\varepsilon} b^2$  we get

$$\begin{aligned} \int f_{x_k} u_{x_k} &\geq -\|\kappa_{x_k}\|_\infty \left( \varepsilon \|\nabla u_{x_k}\|_{L^2}^2 + \frac{1}{4\varepsilon} \|\nabla u\|_{L^2}^2 \right) + \kappa_{\min} \|\nabla u_{x_k}\|_{L^2}^2 \\ &= (\kappa_{\min} - \varepsilon \|\kappa_{x_k}\|_\infty) \|\nabla u_{x_k}\|_{L^2}^2 - \frac{\|\kappa_{x_k}\|_\infty}{4\varepsilon} \|\nabla u\|_{L^2}^2. \end{aligned}$$

Choosing  $\varepsilon = \frac{\kappa_{\min}}{2(\|\kappa_{x_k}\|_\infty + 1)}$  and rearranging terms, we get

$$\begin{aligned} \|\nabla u_{x_k}\|_{L^2}^2 &\leq \kappa_{\min}^{-2} \|\kappa_{x_k}\|_\infty (\|\kappa_{x_k}\|_\infty + 1) \|\nabla u\|_{L^2}^2 + 2\kappa_{\min}^{-1} \int f_{x_k} u_{x_k} \\ &\leq \kappa_{\min}^{-2} \|\kappa_{x_k}\|_\infty (\|\kappa_{x_k}\|_\infty + 1) \|\nabla u\|_{L^2}^2 + \kappa_{\min}^{-1} \|f_{x_k}\|_{L^2}^2 + \kappa_{\min}^{-1} \|u_{x_k}\|_{L^2}^2. \end{aligned}$$

Then we have

$$\begin{aligned} \|\nabla^2 u\|_{L^2}^2 &= \sum_{k=1}^m \|\nabla u_{x_k}\|_{L^2}^2 \leq m\kappa_{\min}^{-2} (\|\nabla \kappa\|_\infty^2 + \|\nabla \kappa\|_\infty) \|\nabla u\|_{L^2}^2 + \kappa_{\min}^{-1} \|\nabla f\|_{L^2}^2 + \kappa_{\min}^{-1} \|\nabla u\|_{L^2}^2 \\ &\leq C \|f\|_{H^3}^2 \left[ \kappa_{\min}^{-3} + \kappa_{\min}^{-4} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3}) \right], \end{aligned} \tag{3.35}$$

where we have used (3.33) and  $C$  only depends on  $\mathcal{M}$ . Moreover, we have used the fact that  $\kappa_{\min} \geq e^{-\|u\|_\infty}$ , which implies  $\kappa_{\min}^{-n_1} \leq \kappa_{\min}^{-n_2}$  if  $n_1 \leq n_2$ . Now we bound the norm of the third derivatives by further differentiating the equation with respect to  $x_j$  and integrating



against  $u_{x_k x_j}$ . We have again by Cauchy's inequality

$$\begin{aligned}
\int f_{x_k x_j} u_{x_k x_j} &= - \int \operatorname{div} \left( \frac{\partial^2}{\partial x_j \partial x_k} (\kappa \nabla u) \right) u_{x_k x_j} \\
&= \int \frac{\partial^2}{\partial x_j \partial x_k} (\kappa \nabla u) \cdot \nabla u_{x_k x_j} \\
&= \int \left[ \kappa_{x_k x_j} \nabla u + \kappa_{x_k} \nabla u_{x_j} + \kappa_{x_j} \nabla u_{x_k} + \kappa \nabla u_{x_k x_j} \right] \cdot \nabla u_{x_k x_j} \\
&\geq \kappa_{\min} \|\nabla u_{x_k x_j}\|_{L^2}^2 - \|\kappa_{x_k x_j}\|_{\infty} \left( \varepsilon_1 \|\nabla u_{x_k x_j}\|_{L^2}^2 + \frac{1}{4\varepsilon_1} \|\nabla u\|_{L^2}^2 \right) \\
&\quad - \|\kappa_{x_k}\|_{\infty} \left( \varepsilon_2 \|\nabla u_{x_k x_j}\|_{L^2}^2 + \frac{1}{4\varepsilon_2} \|\nabla u_{x_j}\|_{L^2}^2 \right) - \|\kappa_{x_j}\|_{\infty} \left( \varepsilon_3 \|\nabla u_{x_k x_j}\|_{L^2}^2 + \frac{1}{4\varepsilon_3} \|\nabla u_{x_k}\|_{L^2}^2 \right) \\
&= \|\nabla u_{x_k x_j}\|_{L^2}^2 \left( \kappa_{\min} - \varepsilon_1 \|\kappa_{x_k x_j}\|_{\infty} - \varepsilon_2 \|\kappa_{x_k}\|_{\infty} - \varepsilon_3 \|\kappa_{x_j}\|_{\infty} \right) \\
&\quad - \frac{1}{4\varepsilon_1} \|\kappa_{x_k x_j}\|_{\infty} \|\nabla u\|_{L^2}^2 - \frac{1}{4\varepsilon_2} \|\kappa_{x_k}\|_{\infty} \|\nabla u_{x_j}\|_{L^2}^2 - \frac{1}{4\varepsilon_3} \|\kappa_{x_j}\|_{\infty} \|\nabla u_{x_k}\|_{L^2}^2.
\end{aligned}$$

Now choosing

$$\varepsilon_1 = \frac{\kappa_{\min}}{4(\|\kappa_{x_k x_j}\|_{\infty} + 1)}, \quad \varepsilon_2 = \frac{\kappa_{\min}}{4(\|\kappa_{x_k}\|_{\infty} + 1)}, \quad \varepsilon_3 = \frac{\kappa_{\min}}{4(\|\kappa_{x_j}\|_{\infty} + 1)}$$

and rearranging terms, we get

$$\begin{aligned}
\|\nabla u_{x_k x_j}\|_{L^2}^2 &\leq 2\kappa_{\min}^{-1} \left( \|f_{x_k x_j}\|_{L^2}^2 + \|u_{x_k x_j}\|_{L^2}^2 \right) + 4\kappa_{\min}^{-2} \|\kappa_{x_k x_j}\|_{\infty} (\|\kappa_{x_k x_j}\|_{\infty} + 1) \|\nabla u\|_{L^2}^2 \\
&\quad + 4\kappa_{\min}^{-2} \|\kappa_{x_k}\|_{\infty} (\|\kappa_{x_k}\|_{\infty} + 1) \|\nabla u_{x_j}\|_{L^2}^2 + 4\kappa_{\min}^{-2} \|\kappa_{x_j}\|_{\infty} (\|\kappa_{x_j}\|_{\infty} + 1) \|\nabla u_{x_k}\|_{L^2}^2.
\end{aligned}$$

Then we have

$$\begin{aligned}
\|\nabla^3 u\|_{L^2}^2 &= \sum_{j=1}^m \sum_{k=1}^m \|\nabla u_{x_k x_j}\|_{L^2}^2 & (3.36) \\
&\leq 2\kappa_{\min}^{-1} \left( \|\nabla^2 f\|_{L^2}^2 + \|\nabla^2 u\|_{L^2}^2 \right) + 4m^2 \kappa_{\min}^{-2} (\|\nabla^2 \kappa\|_{\infty}^2 + \|\nabla^2 \kappa\|_{\infty}) \|\nabla u\|_{L^2}^2 \\
&\quad + 8m \kappa_{\min}^{-2} (\|\nabla \kappa\|_{\infty}^2 + \|\nabla \kappa\|_{\infty}) \|\nabla^2 u\|_{L^2}^2 \\
&\leq C \|f\|_{H^3}^2 \left[ \kappa_{\min}^{-4} + \kappa_{\min}^{-5} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3}) + \kappa_{\min}^{-6} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^2 \right], & (3.37)
\end{aligned}$$

where again we only keep the highest order in  $\kappa_{\min}$ . Differentiating further and applying a similar argument, gives that

$$\begin{aligned}
\|\nabla^4 u\|_{L^2}^2 &\leq 4\kappa_{\min}^{-1} (\|\nabla^3 f\|_{L^2}^2 + \|\nabla^3 u\|_{L^2}^2) + 16m^3 \kappa_{\min}^{-2} \|\nabla u\|_{L^2}^2 (\|\nabla^3 \kappa\|_{\infty}^2 + \|\nabla^3 \kappa\|_{\infty}) \\
&\quad + 48m^2 \kappa_{\min}^{-2} \|\nabla^2 u\|_{L^2}^2 (\|\nabla^2 \kappa\|_{\infty}^2 + \|\nabla^2 \kappa\|_{\infty}) + 48m \kappa_{\min}^{-2} \|\nabla^3 u\|_{L^2}^2 (\|\nabla \kappa\|_{\infty}^2 + \|\nabla \kappa\|_{\infty}) \\
&\leq C \|f\|_{H^3}^2 \left[ \kappa_{\min}^{-5} + \kappa_{\min}^{-6} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3}) + \kappa_{\min}^{-7} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^2 + \kappa_{\min}^{-8} (\|\kappa\|_{\mathcal{C}^3}^2 + \|\kappa\|_{\mathcal{C}^3})^3 \right]. & (3.38)
\end{aligned}$$

The desired result follows by combining equations (3.34), (3.33), (3.35), (3.37) and (3.38).

### 3.6.4 Proof of Lemma 3.3.7

By Lipschitz continuity of  $e^{-x}$  when  $x > 0$ , we have,

$$\begin{aligned}
&\left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_{\varepsilon}(\theta)|_{\Gamma}^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_{\Gamma}^2\right) \right| \\
&\leq \left| \frac{1}{2}|y - \mathcal{G}_{\varepsilon}(\theta)|_{\Gamma}^2 - \frac{1}{2}|y - \mathcal{G}(\theta)|_{\Gamma}^2 \right| \\
&= \frac{1}{2} \left| (\mathcal{G}(\theta)^T + \mathcal{G}_{\varepsilon}(\theta)^T) \Gamma^{-1} (\mathcal{G}(\theta) - \mathcal{G}_{\varepsilon}(\theta)) + 2y^T \Gamma^{-1} (\mathcal{G}(\theta) - \mathcal{G}_{\varepsilon}(\theta)) \right| \\
&\leq \frac{1}{2} \|\Gamma^{-1}\|_2 \left( |\mathcal{G}(\theta)| + |\mathcal{G}_{\varepsilon}(\theta)| \right) |\mathcal{G}(\theta) - \mathcal{G}_{\varepsilon}(\theta)| + \|\Gamma^{-1}\|_2 |y| |\mathcal{G}(\theta) - \mathcal{G}_{\varepsilon}(\theta)|,
\end{aligned}$$

where  $\|\Gamma^{-1}\|_2$  is the operator 2-norm of  $\Gamma^{-1}$ . By Theorem 3.3.1, Lemma 3.3.2 and (3.34),

$$\begin{aligned} |\mathcal{G}(\theta) - \mathcal{G}_\varepsilon(\theta)| &\leq \sqrt{\sum \|\ell_j\|^2} \|u - u_\varepsilon\|_{L^2} \leq C \sqrt{\sum \|\ell_j\|^2} A(\theta) \|f\|_{H^3} \varepsilon^{4\beta-1} \\ |\mathcal{G}(\theta)| + |\mathcal{G}_\varepsilon(\theta)| &\leq \sqrt{\sum \|\ell_j\|^2} (\|u\|_{L^2} + \|u_\varepsilon\|_{L^2}) \leq C \sqrt{\sum \|\ell_j\|^2} e^{\|\theta\|_\infty} \|f\|_{L^2}, \end{aligned}$$

where  $u$  and  $u_\varepsilon$  are the zero-mean solutions associated with  $\theta$ , and  $\|\ell_j\|$  is the operator norm of  $\ell_j$ . Here we have written  $A$  as a function of  $\theta$  and use the fact that  $\kappa_{\min} \geq e^{-\|\theta\|_\infty}$ . So

$$\begin{aligned} &\int \left| \exp\left(-\frac{1}{2}|y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2\right) - \exp\left(-\frac{1}{2}|y - \mathcal{G}(\theta)|_\Gamma^2\right) \right| d\pi(\theta) \\ &\leq C \|\Gamma^{-1}\|_2 \varepsilon^{4\beta-1} \left[ \sum \|\ell_j\|^2 \|f\|_{H^3}^2 \int e^{\|\theta\|_\infty} A(\theta) d\pi(\theta) + \sqrt{\sum \|\ell_j\|^2} \|y\| \|f\|_{H^3} \int A(\theta) d\pi(\theta) \right]. \end{aligned}$$

It now suffices to show  $\int (e^{\|\theta\|_\infty} \vee 1) A(\theta) d\pi(\theta) < \infty$ . Since  $\kappa = e^\theta$ , we have

$$\|\kappa\|_{\mathcal{C}^4} \leq C e^{\|\theta\|_\infty} \left( \|\theta\|_{\mathcal{C}^4} + \|\theta\|_{\mathcal{C}^4}^2 + \|\theta\|_{\mathcal{C}^4}^3 + \|\theta\|_{\mathcal{C}^4}^4 \right),$$

where  $C$  is a constant depending on the dimension  $m$  and a similar relation is true for  $\|\sqrt{\kappa}\|_{\mathcal{C}^4}$ . Keeping only the highest order term in  $e^{\|\theta\|_\infty}$ , we have

$$\begin{aligned} A(\theta) &\leq C \sqrt{P_1(\|\theta\|_{\mathcal{C}^4}) e^{14\|\theta\|_\infty} P_2(\|\theta\|_{\mathcal{C}^4})} e^{\|\theta\|_\infty} \\ &\leq C \sqrt{P_1(\|\theta\|_{\mathcal{C}^4}) e^{14\|\theta\|_{\mathcal{C}^4}} P_2(\|\theta\|_{\mathcal{C}^4})} e^{\|\theta\|_{\mathcal{C}^4}}, \end{aligned}$$

where  $P_1$  and  $P_2$  are polynomials. Since  $\pi$  is a Gaussian measure on  $\mathcal{C}^4$ , by Fernique's theorem [Fernique, 1970],

$$\int (e^{\|\theta\|_\infty} \vee 1) A(\theta) d\pi(\theta) < \infty.$$

It follows that

$$\int \left| \exp \left( -\frac{1}{2} |y - \mathcal{G}_\varepsilon(\theta)|_\Gamma^2 \right) - \exp \left( -\frac{1}{2} |y - \mathcal{G}(\theta)|_\Gamma^2 \right) \right| d\pi(\theta) \leq C \varepsilon^{4\beta-1},$$

where  $C$  depends on  $\Gamma, y, f$  and the  $\ell_j$ 's, but is independent of  $\varepsilon$ .

CHAPTER 4

UNLABELED DATA HELP IN GRAPH-BASED  
SEMI-SUPERVISED LEARNING: A BAYESIAN  
NONPARAMETRICS PERSPECTIVE

4.1 Introduction

Semi-supervised learning (SSL) refers to machine learning techniques that combine during training a small amount of labeled data with a large amount of unlabeled data. SSL has received increasing attention over the past decades because in many recent applications labeled data are expensive to collect while unlabeled data are abundant. Examples include analysis of body-worn videos and video surveillance [Qiao et al., 2019], text categorization and translation [Shi et al., 2010], image classification [Zhu, 2005] and protein structure prediction [Weston et al., 2005]. The aim of this chapter is to investigate whether unlabeled data can enhance learning performance. The answer to such question necessarily depends on the model assumptions and the methodology employed. Our main contribution is to show that under a standard manifold assumption, unlabeled data are helpful when using graph-based methods in a Bayesian setting. We do so by establishing that the optimal posterior contraction rate is achieved (up to a logarithmic factor) provided that the size of the unlabeled dataset grows sufficiently fast with the size of the labeled dataset.

The SSL problem of interest can be informally described as follows. Given labeled data  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$  and unlabeled data  $\{X_{n+1}, \dots, X_{N_n}\}$ , the goal is to predict  $Y$  from  $X$ . More precisely, we are interested in the SSL problem of inferring the regression function  $f_0(x) := \mathbb{E}(Y|X = x)$  at the given features  $\mathcal{X}_{N_n} := \{X_1, \dots, X_{N_n}\}$  in either of these two settings:

1. Regression:  $Y = f_0(X) + \eta$ , where  $\eta \sim \mathcal{N}(0, \sigma^2)$  with  $\sigma^2$  known.

2. Binary classification:  $\mathbb{P}(Y = 1|X) = f_0(X)$ .

We analyze graph-based methods applied to this inference task under the *manifold assumption* that  $X$  takes values on a hidden manifold  $\mathcal{M}$ . This assumption is natural when the features  $\mathcal{X}_{N_n}$  live in a high-dimensional ambient space but admit a low-dimensional representation [Bickel and Li, 2007, Niyogi, 2013, García Trillos et al., 2019b, 2020b]. For instance, Hein and Audibert [2005] and Costa and Hero [2006] demonstrate the low intrinsic dimension of standard datasets for image classification. Graph-based methods are well-suited under the manifold assumption as they allow to uncover the geometry of the hidden manifold and promote smoothness of the inferred  $f_0$  along it. Indeed, graph-based methods are among the most powerful classification techniques for SSL problems where similar features are expected to belong to the same class [Belkin et al., 2004, Belkin and Niyogi, 2004, Zhu, 2005]. The central idea behind traditional graph-based methods is to infer  $f_0$  by minimizing an objective function comprising at least two terms: (i) a regularization term constructed with a graph-Laplacian of the features  $\mathcal{X}_{N_n}$ , which leverages the ability of unsupervised graph-based techniques to extract geometric information from  $\mathcal{M}$ ; and (ii) a data-fidelity term that incorporates the labeled data. We adopt an analogous Bayesian perspective where: (i) the prior distribution  $\Pi_n(\cdot | \mathcal{X}_{N_n})$  will be defined using a graph-Laplacian of the features  $\mathcal{X}_{N_n}$  (hence the notation “given  $\mathcal{X}_{N_n}$ ”) to extract geometric information from  $\mathcal{M}$ ; and (ii) the likelihood function promotes matching the labeled data. Assuming the data are independent so that the likelihood factorizes, the posterior takes the form

$$\Pi_n(B | \mathcal{X}_{N_n}, \mathcal{Y}_n) \propto \int_B \prod_{i=1}^n L_{Y_i|X_i}(f) d\Pi_n(f | \mathcal{X}_{N_n}), \quad B \in \mathcal{B}, \quad (4.1)$$

where  $\mathcal{Y}_n := \{Y_1, \dots, Y_n\}$  and  $\mathcal{B}$  is the Borel  $\sigma$ -algebra on  $\mathbb{R}^{N_n}$  (here we are identifying

functions over  $\mathcal{X}_{N_n}$  with  $\mathbb{R}^{N_n}$ ). The conditional likelihood of  $Y_i|X_i$  is given by

$$L_{Y_i|X_i}(f) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{|Y_i - f(X_i)|^2}{2\sigma^2}\right) \quad (4.2)$$

in the regression setting and

$$L_{Y_i|X_i}(f) = f(X_i)^{Y_i} [1 - f(X_i)]^{1-Y_i} \quad (4.3)$$

in the classification setting.

We shall analyze our Bayesian approach in a frequentist perspective by assuming the data is generated from a fixed truth  $f_0$  and studying the contraction rate of the posterior around  $f_0$ , defined as the sequence of real numbers  $\varepsilon_n$  such that

$$\mathbb{E}_X \mathbb{E}_{f_0} \Pi_n (f : d_n(f, f_0) \geq M_n \varepsilon_n \mid \mathcal{X}_{N_n}, \mathcal{Y}_n) \xrightarrow{n \rightarrow \infty} 0,$$

for any sequence  $M_n \rightarrow \infty$  and some suitable semi-metric  $d_n$ . Here the double expectation  $\mathbb{E}_X \mathbb{E}_{f_0}$  is taken first with respect to the conditional data distribution of  $\mathcal{Y}_n|X_1, \dots, X_n$  specified by  $f_0$  and then with respect to the marginal of  $\mathcal{X}_{N_n}$ . The idea of posterior contraction rate was formally introduced in Ghosal et al. [2000] and Shen and Wasserman [2001] and has since then become a popular criterion for analyzing Bayesian methods. In particular, contraction of the posterior with rate  $\varepsilon_n$  gives a point estimator defined as

$$\hat{f}_n = \arg \max_g \left[ \Pi_n(f : d_n(f, g) \leq M_n \varepsilon_n \mid \mathcal{X}_{N_n}, \mathcal{Y}_n) \right]$$

that converges to  $f_0$  with the same rate  $\varepsilon_n$ , which together with the minimax theory for statistical estimation quantifies the performance of our Bayesian SSL procedure.

Our main result, Theorem 4.2.1, shows that the posterior (4.1) contracts around  $f_0$  at the minimax optimal rate (up to logarithmic factors), provided that the prior distribution

$\Pi_n(\cdot | \mathcal{X}_{N_n})$  is carefully designed and that  $N_n$  grows at a certain polynomial rate with  $n$ . More broadly, our theory suggests that graph-based methods require abundant unlabeled data in order to effectively extract geometric information to regularize SSL problems.

#### 4.1.1 Related Work

The question of whether unlabeled data enhance SSL performance has been widely studied. Positive and negative conclusions have been reached depending on the assumptions made on the relationship between the target function and the marginal data distribution, and also on the methodology employed [Cozman and Cohen, 2006, Bickel and Li, 2007, Liang et al., 2007, Wasserman and Lafferty, 2008, Niyogi, 2013, Singh et al., 2008]. Two standard model assumptions are the *clustering assumption* [Seeger, 2000a], which states that the target function is smooth over high density regions, and the *manifold assumption* that we have introduced. The latter has been extensively used in the statistical learning literature, and specifically in SSL, e.g. Bickel and Li [2007], Wasserman and Lafferty [2008], Castillo et al. [2014], Yang and Dunson [2016], García Trillos et al. [2019b, 2020b].

Several methodologies for SSL have been developed based on generative modeling, support vector machines, semi-definite programming, graph-based methods, etc. (see the overview in Zhu [2005] and Chapelle et al. [2006]). Our focus is on graph-based approaches that combine label information with geometric information extracted from the unlabeled data employing graphical unsupervised techniques [Von Luxburg, 2007]. This heuristic has motivated the use of graph-based regularizations in a wide number of applications, but a rigorous analysis of the mechanisms by which unlabeled data enhance the performance of graph-based SSL methods is only starting to emerge. The recent papers Bertozzi et al. [2021] and Hoffmann et al. [2020] studied posterior *consistency* (the asymptotics of the posterior probability  $\Pi_n(f : d_n(f, f_0) \geq \varepsilon | \text{data})$  for a fixed small number  $\varepsilon$ ) for a fixed sample size in the small noise limit, whereas we consider the large  $n$  limit and further establish posterior



contraction rates. Rates of convergence for optimization rather than Bayesian formulations of graph-based SSL have been recently established in Calder et al. [2020]. In a Bayesian framework, García Trillos and Sanz-Alonso [2018] and García Trillos et al. [2020b] show the continuum limit of posterior distributions as the size of the unlabeled dataset grows, without increasing the size of the labeled dataset. These papers did not investigate whether the posteriors contract around the truth, and they did not demonstrate the value of unlabeled data in boosting learning performance. The work Kirichenko and van Zanten [2017] studies fully-supervised function estimation on large graphs assuming that the truth changes with the size of the graph. In contrast, we investigate posterior contraction in a SSL setting with a *fixed* truth  $f_0$  defined on the underlying manifold  $\mathcal{M}$ .

Our analysis uses tools from Bayesian nonparametrics and spectral analysis of graph-Laplacians. We will provide the necessary background on posterior contraction in Section 4.3, and we refer to [Ghosal and Van der Vaart, 2017, Chapters 6, 8 and 11] for a comprehensive introduction to this subject. While numerous results on spectral convergence of graph-Laplacians can be found in the literature, our analysis of posterior contraction requires bounds in  $L^\infty$ -metric with rates, recently developed in Dunson et al. [2021] and Calder et al. [2022].

#### 4.1.2 *Main Contributions and Scope*

Our main result, Theorem 4.2.1, is to our knowledge the first to establish posterior contraction rates for graph-based SSL. In doing so, we provide novel understanding on the relative value of labeled and unlabeled data, and we set forth a rigorous quantitative framework in which to analyze this question. We point out, however, two caveats. First, our theory is non-adaptive in the sense that the Bayesian methodology we analyze only achieves optimal contraction rates when a priori information on the smoothness of the truth is available. Our analysis is motivated by existing graph-based techniques, and the development and analy-

sis of adaptive graph-based methods for SSL is an important research direction stemming from our work, but beyond the scope of this chapter. Second, our results build on existing spectral convergence rates of graph-Laplacians that may be suboptimal; as a consequence, the required sample size of the unlabeled data that we establish may not be sharp. Due to the plug-in nature of our analysis, improved spectral convergence rates will automatically translate into a sharper bound on the sample complexity. Despite these caveats, our theory provides evidence that letting the unlabeled dataset grow polynomially with the labeled dataset, as suggested by Theorem 4.2.1, is a fundamental requirement for standard graph-based methods to achieve optimal contraction.

A central part of our proof is devoted to analyzing the convergence of a discretely indexed Gauss-Markov random field in an unstructured data cloud to a Matérn-type Gaussian field on  $\mathcal{M}$ . This is formalized in Theorem 4.4.6, a result that we believe to be of independent interest. Finally, our work contributes to the Bayesian nonparametrics literature on manifolds [Castillo et al., 2014].

**Outline** The rest of this chapter is organized as follows. Section 4.2 introduces the construction of the graph-based prior and states our main result. Section 4.3 provides the necessary background on posterior contraction and outlines our analysis. Section 4.4 proves our main result, and we close in Section 4.5.

## 4.2 Prior Construction and Main Result

In this section we introduce the graph-based prior  $\Pi_n(\cdot | \mathcal{X}_{N_n})$  and we state the main result of this chapter. Before doing so, we formalize our setting. We assume to be given labeled data  $(X_1, Y_1), \dots, (X_n, Y_n) \stackrel{i.i.d.}{\sim} \mathcal{L}(X, Y)$  and unlabeled data  $X_{n+1}, \dots, X_{N_n} \stackrel{i.i.d.}{\sim} \mathcal{L}(X)$ , where  $\{X_i\}_{i=n+1}^{N_n}$  are independent from  $\{X_i\}_{i=1}^n$ . Recall that the goal is to estimate the regression

function

$$f_0(x) = \mathbb{E}(Y|X = x)$$

at the given features  $\mathcal{X}_{N_n} = \{X_1, \dots, X_{N_n}\}$ . We suppose that  $\mu := \mathcal{L}(X)$  is supported on an  $m$ -dimensional smooth, connected, compact manifold  $\mathcal{M}$  without boundary embedded in  $\mathbb{R}^d$ , with the absolute value of the sectional curvature bounded and with Riemannian metric inherited from the embedding. For technical reasons, we further assume that  $m \geq 2$  and that  $\mathcal{M}$  is a homogeneous manifold (the group of isometries acts transitively on  $\mathcal{M}$ , e.g. spheres and tori) normalized so that its volume is 1. We assume for simplicity that  $\mu$  is the uniform distribution on  $\mathcal{M}$ . As discussed in Section 4.5, our results can be generalized to nonuniform marginal density.

#### 4.2.1 Graph-based Prior

We now describe the construction of the graph-based prior  $\Pi_n(\cdot | \mathcal{X}_{N_n})$  on  $f_0$  restricted to the given features  $\mathcal{X}_{N_n}$ . The priors we consider have the general form

$$\Pi_n(\cdot | \mathcal{X}_{N_n}) = \mathcal{L}(\Phi(W_n) | \mathcal{X}_{N_n}), \quad (4.4)$$

where  $W_n$  is a Gauss-Markov random field in  $\mathbb{R}^{N_n}$  whose covariance will be defined in terms of a graph-Laplacian [Von Luxburg, 2007] of  $\mathcal{X}_{N_n}$ . For the regression problem,  $\Phi$  is taken to be the identity. For the classification problem, where  $f_0$  takes values in  $(0, 1)$ ,  $\Phi : \mathbb{R} \rightarrow (0, 1)$  is a link function, which we assume throughout to be invertible and twice differentiable with  $\Phi' / (\Phi(1 - \Phi))$  uniformly bounded and  $\int (\Phi'')^2 / \Phi' d\gamma < \infty$ . The logistic function, for instance, satisfies all these standard requirements.

In the remainder of this subsection we introduce and motivate our construction of the Gauss-Markov random field  $W_n$ . The starting point is to define a similarity matrix  $H \in$

$\mathbb{R}^{N_n \times N_n}$  whose entries  $H_{ij} \geq 0$  represent the closeness between features  $X_i$  and  $X_j$ . For reasons discussed in Section 4.2.1.1 below, we set

$$H_{ij} := \frac{2(m+2)}{N_n \nu_m \zeta_{N_n}^{m+2}} \mathbf{1}\{|X_i - X_j| < \zeta_{N_n}\}, \quad (4.5)$$

where  $|\cdot|$  is the Euclidean distance in  $\mathbb{R}^d$ ,  $\nu_m$  is the volume of the  $m$ -dimensional unit ball, and  $\zeta_{N_n}$  is the *connectivity* of the graph, a user-specified parameter. Recall that a graph-Laplacian is a positive semi-definite matrix obtained by suitably transforming a similarity matrix. For concreteness, we work with the unnormalized graph-Laplacian matrix  $\Delta_{N_n} := D - H$ , where  $D$  is the diagonal matrix with entries  $D_{ii} = \sum_{j=1}^{N_n} H_{ij}$ . For a vector  $v \in \mathbb{R}^{N_n}$ , naturally identified with a function on the data cloud  $\mathcal{X}_{N_n}$ , we have

$$v^T \Delta_{N_n} v = \frac{1}{2} \sum_{i=1}^{N_n} H_{ij} |v(X_i) - v(X_j)|^2. \quad (4.6)$$

We see that indeed  $\Delta_{N_n}$  is positive semi-definite and that functions  $v$  that change slowly with respect to the similarity  $H$  have a small value of  $v^T \Delta_{N_n} v$ . This naturally suggests considering the Gaussian distribution  $\mathcal{N}(0, (I_{N_n} + \Delta_{N_n})^{-s})$  since the likelihood for the samples will then have a term similar to (4.6). Based on this idea we shall define the Gauss-Markov random field  $W_n$  as the principal components of it. Precisely, let  $\{(\lambda_i^{(N_n)}, \psi_i^{(N_n)})\}_{i=1}^{N_n}$  be the ordered eigenpairs of  $\Delta_{N_n}$  and define

$$W_n = \sum_{i=1}^{k_{N_n}} \left[1 + \lambda_i^{(N_n)}\right]^{-\frac{s}{2}} \xi_i \psi_i^{(N_n)}, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (4.7)$$

where  $k_{N_n} \ll N_n$  is to be determined. Notice that  $W_n$  as in (4.7) is a truncation of the Karhunen-Loève expansion of samples from  $\mathcal{N}(0, (I_{N_n} + \Delta_{N_n})^{-s})$  and therefore the terminology “principal components”.

The law of  $W_n$  depends on three parameters: the graph *connectivity*  $\zeta_{N_n}$  used to define the

similarity matrix  $H$ , the *smoothness* parameter  $s$ , and the principal components *truncation* parameter  $k_{N_n}$ . The connectivity  $\zeta_{N_n}$  determines the sparsity of the precision matrix of  $W_n$ , and it should be taken to decrease with  $N_n$  to better resolve the geometry of  $\mathcal{M}$  as more unlabeled data are available. The smoothness  $s$  controls the level of regularization. Larger  $s$  leads to faster decay of the coefficients in (4.7) and smoother samples. The truncation parameter  $k_{N_n} \ll N_n$  allows us to keep only the components of the graph-Laplacian that contain useful geometric information on  $\mathcal{M}$ . Suitable choices and scalings of these three parameters, as well as further insights on their interpretation, will be given as we develop our theory. Importantly, the construction does not require knowledge of the underlying manifold  $\mathcal{M}$ , but only of its dimension. These data-driven Gauss-Markov fields have been used within various intrinsic approaches to Bayesian learning, see e.g. García Trillos and Sanz-Alonso [2018], García Trillos et al. [2019b, 2020b], Harlim et al. [2020].

#### 4.2.1.1 Interpretation as a Discretely Indexed Matérn Field

The Gauss-Markov field  $W_n$  can be interpreted as a discretely indexed Matérn Gaussian field [Sanz-Alonso and Yang, 2022a]. Consider the Gaussian measure  $\mathcal{N}(0, (I - \Delta)^{-s})$ , where  $-\Delta$  denotes the Laplace-Beltrami operator on  $\mathcal{M}$  and the fractional order operator is defined spectrally. Then, draws from  $\mathcal{N}(0, (I - \Delta)^{-s})$  admit a representation

$$W^{\mathcal{M}} = \sum_{i=1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (4.8)$$

where  $(\lambda_i, \psi_i)$ 's are the ordered eigenpairs of the Laplace-Beltrami operator. Note the analogy with (4.7). The field  $W^{\mathcal{M}}$  is a generalization of the Matérn model to compact manifolds [Lindgren et al., 2011, Sanz-Alonso and Yang, 2022a]. An important step in our analysis of posterior contraction will be to show the convergence of  $W_n$  towards  $W^{\mathcal{M}}$ , in a sense to be made precise in Section 4.4.3, provided that the connectivity, smoothness and truncation

parameters of  $W_n$  are suitably chosen. The similarity matrix defined in (4.5) enables this convergence result, but other choices (e.g.  $K$ -nearest neighbors) are possible. Key to showing convergence of  $W_n$  to  $W^{\mathcal{M}}$  is the spectral convergence of  $\Delta_{N_n}$  to  $-\Delta$ . The truncation parameter  $k_{N_n}$  is motivated by the fact that only the first eigenpairs of  $\Delta_{N_n}$  accurately approximate those of  $-\Delta$ ; see Proposition 4.4.9 below for a precise statement.

#### 4.2.1.2 Reconciling Optimization and Bayesian Perspectives

To further motivate the definition of  $W_n$ , we relate our prior construction to the optimization literature. To streamline the discussion, let us focus on the regression problem where the link function  $\Phi$  in (4.4) is the identity map. Classical graph-based optimization recovers  $f_0$  at  $\mathcal{X}_{N_n}$  using  $f^T \Delta_{N_n} f$  as a regularizer and an appropriate data-fidelity term to match the labeled data, for instance,

$$\hat{f}_0 := \arg \min_{f \in \mathbb{R}^{N_n}} \sum_{i=1}^n |Y_i - f(X_i)|^2 + \lambda f^T \Delta_{N_n} f, \quad (4.9)$$

where  $\lambda$  controls the level of regularization. The solution to (4.9) is conceptually equivalent to the *maximum a posteriori* estimator in the Bayesian approach when the prior is chosen to be  $\mathcal{N}(0, \Delta_{N_n}^{-1})$ , where  $\Delta_{N_n}^{-1}$  denotes the pseudo-inverse of  $\Delta_{N_n}$ . The paper Nadler et al. [2009] shows that (4.9) is not well-posed when  $m \geq 2$ , in the sense that as the number of unlabeled data points increases the solution degenerates to a noninformative function. The authors suggest that this issue can be alleviated by defining the regularization term as  $\lambda f^T \Delta_{N_n}^s f$  with  $s > \frac{m}{2}$ , so that higher order “derivatives” of  $f$  are controlled. Similar behavior has been observed with  $p$ -Laplacian regularizations [El Alaoui et al., 2016]. The parameter  $s$  in (4.7) plays the exact same role, and we shall see that suitably scaling  $s$  with  $m$  is needed to warrant consistent learning in the limit of large unlabeled datasets.

### 4.2.2 Main Result

Now we are ready to state our main result. Let  $\delta > 0$  be arbitrary. Let  $p_m = \frac{3}{4}$  if  $m = 2$  and  $p_m = \frac{1}{m}$  otherwise. Let  $\alpha_s = \frac{6m+6}{2s-3m+1}$  if  $s \leq \frac{9}{2}m + \frac{5}{2}$  and  $\alpha_s = 1$  otherwise.

**Theorem 4.2.1.** *Suppose  $\Phi^{-1}(f_0) \in B_{\infty, \infty}^\beta$ . Let  $\Pi_n$  be the prior defined by (4.4) and (4.7) with  $s > \frac{3}{2}m - \frac{1}{2}$ . Consider the following scaling for  $\zeta_{N_n}, k_{N_n}$  and  $N_n$ .*

$$1. \ m \leq 4: \zeta_{N_n} \asymp N_n^{-\frac{1}{m+4+\delta}} (\log N_n)^{\frac{pm}{2}}, k_{N_n} \asymp N_n^{\frac{m}{(m+4+\delta)(2s-3m+1)\alpha_s}} (\log N_n)^{-\frac{mpm}{(4s-6m+2)\alpha_s}}$$

with

$$N_n \asymp n^{(m+4+\delta)\alpha_s} (\log n)^{\frac{(m+4+\delta)pm}{2}}. \quad (4.10)$$

$$2. \ m \geq 5: \zeta_{N_n} \asymp N_n^{-\frac{1}{2m}} (\log N_n)^{\frac{pm}{2}}, k_{N_n} \asymp N_n^{\frac{1}{(4s-6m+2)\alpha_s}} (\log N_n)^{-\frac{mpm}{(4s-6m+2)\alpha_s}} \text{ with}$$

$$N_n \asymp n^{2m\alpha_s} (\log n)^{mpm}. \quad (4.11)$$

Then, for  $\varepsilon_n$  a multiple of  $n^{-\frac{(s-\frac{m}{2})\wedge\beta}{2s}} (\log n)^{\frac{(s-\frac{m}{2})\wedge\beta}{4s-2m}}$  and all  $M_n \rightarrow \infty$ ,

$$\mathbb{E}_X \mathbb{E}_{f_0} \Pi_n (f \in L^\infty(\mu_{N_n}) : \|f - f_0\|_n \geq M_n \varepsilon_n \mid \mathcal{X}_{N_n}, \mathcal{Y}_n) \xrightarrow{n \rightarrow \infty} 0, \quad (4.12)$$

where  $\mu_{N_n}$  is the empirical measure of  $\mathcal{X}_{N_n}$  and  $\|f - f_0\|_n^2 := \frac{1}{n} \sum_{i=1}^n |f(X_i) - f_0(X_i)|^2$ . Here  $\mathbb{E}_{f_0}$  denotes expectation with respect to the conditional distribution of  $\mathcal{Y}_n \mid X_1, \dots, X_n$  determined by  $f_0$  and  $\mathbb{E}_X$  denotes expectation with respect to the marginal distribution of  $\mathcal{X}_{N_n}$ .

Theorem 4.2.1 presents the posterior contraction rates with respect to the priors constructed in Section 4.2.1 under suitable choices of the parameters. Several remarks are in order.

The truth  $f_0$  is assumed to belong to a Besov-type space  $B_{\infty,\infty}^\beta$  on manifolds, defined in Section 4.4.2 in analogy to the usual Besov space  $B_{\infty,\infty}^\beta(\mathbb{R}^m)$  on Euclidean space. We recall that  $B_{\infty,\infty}^\beta(\mathbb{R}^m)$  coincides with the Hölder space  $\mathcal{C}^\beta(\mathbb{R}^m)$  (the space of functions with  $\lfloor\beta\rfloor$  continuous derivatives whose  $\lfloor\beta\rfloor$ -th derivative is  $\beta - \lfloor\beta\rfloor$ -Hölder continuous) when  $\beta \notin \mathbb{N}$  and contains  $\mathcal{C}^\beta(\mathbb{R}^m)$  when  $\beta \in \mathbb{N}$  (see e.g. Triebel [1992][Theorem 1.3.4]). Therefore,  $B_{\infty,\infty}^\beta$  can be interpreted as an analog of the Hölder space with parameter  $\beta$  on manifolds, and in particular represents a space of  $\beta$ -regular functions.

The key implication of Theorem 4.2.1 is that when  $s - \frac{m}{2} = \beta$ , we recover the rate  $n^{-\beta/(2\beta+m)}(\log n)^{1/4}$ , which is minimax optimal up to a logarithmic factor for  $\beta$ -regular functions. The assumption  $s > \frac{3}{2}m - \frac{1}{2}$  then requires  $\beta > m - \frac{1}{2}$ , so optimal contraction rates can only be attained if  $f_0$  is not too rough. Theorem 4.2.1 can be extended to hold for all  $s > m$  if the eigenfunctions of the Laplace-Beltrami operator on  $\mathcal{M}$  are uniformly bounded, which holds for the family of flat manifolds [Toth and Zelditch, 2002] that include for instance the tori. As mentioned in Section 4.1.2, the above choice of  $s$  requires knowing the regularity of  $f_0$  and is not adaptive.

Another key ingredient of the result is the scaling for  $N_n$  as in (4.10) and (4.11), which are both larger than a multiple of  $n^{2m}$  since  $\alpha_s \geq 1$  in all cases. In other words, the required sample size of the unlabeled data should grow polynomially with respect to the sample size of the labeled data in order to achieve the near optimal contraction rates described above, thereby justifying the claim that unlabeled data help.

We further remark that Theorem 4.2.1 only gives an upper bound for the required sample size  $N_n$ , whose proof (see Section 4.4) in fact has a plug-in nature. Suppose the sequence of Gaussian-Markov fields  $W_n$  in (4.7) converges in some semimetric  $d_n$  towards  $W^\mathcal{M}$  in (4.8) with rate  $\mathbb{E}_\xi d_n(W_n, W^\mathcal{M}) = R(N_n)$  for some function  $R$ . The required sample size is then obtained by matching  $R(N_n)$  with  $n^{-1}$ . In particular, any improvement of the rate  $R(N_n)$  shown in Section 4.4.3 will lead to a reduction of the required sample size. However,



the spectral convergence rate in Proposition 4.4.9 and hence  $R(N_n)$  should not be faster than the resolution of the point cloud, which is shown to scale like  $N_n^{-1/m}$  in Proposition 4.4.5. Therefore by matching  $N_n^{-1/m}$  with  $n^{-1}$  we expect a polynomial dependence of  $N_n$  on  $n$  to be necessary to achieve optimal contraction rate, further demonstrating the need of unlabeled data in our specific setting.

### 4.3 Posterior Contraction: Background and Set-up

In this section we present necessary background on posterior contraction rates. In Section 4.3.1 we review the main results on posterior contraction that our theory builds on, and in Section 4.3.2 we explain how these general results are used in our proof.

It will be important to note that our prior is constructed with the  $X_i$ 's, and we shall view our observations as the  $Y_i$ 's only, as the notation in (4.1) suggests. In particular, the  $Y_i$ 's are independent but non-identically distributed (i.n.i.d.) and hence we will apply general results from Ghosal and Van Der Vaart [2007]. This also explains the double expectation in (4.12), where the randomness of the  $X_i$ 's is treated separately.

#### 4.3.1 General Principles

Here we review general posterior contraction theory for the problem of estimating  $f_0$  at the continuum level from i.n.i.d data, following Ghosal and Van Der Vaart [2007]. Suppose the data  $\{Y_i\}_{i=1}^n$  are generated according to a density  $P_{f_0}^{(n)} = \prod_{i=1}^n p_{f_0,i}$  for some ground truth parameter  $f_0$ , where  $p_{f_0,i}$  is the individual density for each observation. Let  $\Pi_n$  be a sequence of priors over  $f_0$  that is supported on some parameter space  $\mathcal{F}$  equipped with a semimetric  $d_n$ . Theorem 4 of Ghosal and Van Der Vaart [2007] states that

$$\mathbb{E}_{f_0} \Pi_n (f \in \mathcal{F} : d_n(f, f_0) \geq M_n \varepsilon_n \mid \mathcal{Y}_n) \xrightarrow{n \rightarrow \infty} 0, \quad (4.13)$$

provided that there exists sets  $\mathcal{F}_n \subset \mathcal{F}$  and positive real numbers  $\varepsilon_n, K$  so that

$$\log N(\varepsilon_n, \mathcal{F}_n, d_n) \leq n\varepsilon_n^2, \quad (4.14)$$

$$\Pi_n(\mathcal{F}_n^c) \leq e^{-(K+4)n\varepsilon_n^2}, \quad (4.15)$$

$$\Pi_n(B_n^*(f_0, \varepsilon_n)) \geq e^{-Kn\varepsilon_n^2}, \quad (4.16)$$

where  $N(\varepsilon_n, \mathcal{F}_n, d_n)$  is the minimum number of  $d_n$ -balls of radius  $\varepsilon_n$  needed to cover  $\mathcal{F}_n$ .

Here

$$B_n^*(f_0, \varepsilon_n) := \left\{ f \in \mathcal{F} : \frac{1}{n} \sum_{i=1}^n d_{\text{KL}}(p_{f,i}, p_{f_0,i}) \leq \varepsilon_n^2, \quad \frac{1}{n} \sum_{i=1}^n v_{\text{KL}}(p_{f,i}, p_{f_0,i}) \leq C\varepsilon_n^2 \right\}, \quad (4.17)$$

where  $C$  is a universal constant,  $d_{\text{KL}}(p, q) = \int p \log(p/q) d\gamma$  is the Kullback-Leibler divergence and  $v_{\text{KL}}(p, q) = \int p |\log(p/q) - d_{\text{KL}}(p, q)|^2 d\gamma$ . Conditions (4.15) and (4.14) roughly state that there are approximating sieves  $\mathcal{F}_n$  which capture most of the prior probability while not being too large. Condition (4.16) requires sufficient prior mass near the truth  $f_0$  and together with (4.14) further indicates that the priors should be “uniformly spread”. In fact the three conditions are stronger than those in [Ghosal and Van Der Vaart, 2007, Theorem 4] but will suffice in our case for Gaussian process priors to be discussed shortly below. We shall refer to [Ghosal et al., 2000, Section 2] and [Ghosal and Van Der Vaart, 2007, Section 2] for further discussion on the interpretation and relaxation of these conditions.

The general theory can be used to analyze a wide range of statistical models but does not give a recipe for constructing the sieves  $\mathcal{F}_n$  and  $\varepsilon_n$ . However, when the priors are Gaussian there exists a simple relation between  $\varepsilon_n$  and the priors. Suppose in addition that  $\Pi_n = \mathcal{L}(w_n)$  are Gaussian measures on some Banach space  $(\mathbb{B}, \|\cdot\|_{\mathbb{B}})$  that converge to a fixed Gaussian measure  $\Pi = \mathcal{L}(w)$  on the same Banach space with  $10\mathbb{E}\|w_n - w\|_{\mathbb{B}}^2 \leq n^{-1}$ . Theorem 2.2 in van der Vaart and van Zanten [2008a] then states that the contraction rate

$\varepsilon_n$  can be obtained by studying the *concentration function* of the limit prior  $\Pi$ , defined as

$$\phi_{f_0}(\varepsilon) := \inf_{h \in \mathbb{H}: \|h - f_0\|_{\mathbb{B}} < \varepsilon} \|h\|_{\mathbb{H}}^2 - \log \mathbb{P}(\|w\|_{\mathbb{B}} < \varepsilon),$$

where  $(\mathbb{H}, \|\cdot\|_{\mathbb{H}})$  is the *reproducing kernel Hilbert space* (RKHS) of  $\Pi$  (see e.g. van der Vaart and van Zanten [2008b] for more details on Gaussian measures and the associated RKHSs). More precisely, if  $\varepsilon_n$  satisfies  $\phi_{f_0}(\varepsilon_n) \leq n\varepsilon_n^2$ , then for the same  $\varepsilon_n$  and some  $\mathcal{F}_n$ , the three conditions (4.14), (4.15) and (4.16) are satisfied (possibly for different proportion constants in front of  $n\varepsilon_n^2$ ) with  $d_n$  and  $B_n^*(f_0, \varepsilon_n)$  replaced by  $\|\cdot\|_{\mathbb{B}}$  and  $\{f : \|f - f_0\|_{\mathbb{B}} < \varepsilon_n\}$ , respectively.

### 4.3.2 Application to Our Setting

In this subsection we discuss how we utilize the general theory outlined above, and provide a road map for the proof of Theorem 4.2.1. Notice that in (4.12) the sequence of posteriors are supported on the discrete space  $L^\infty(\mu_{N_n})$ , whose size changes with  $n$ . To alleviate this issue we will reduce the analysis to a sequence of posteriors  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n)$  supported on the same continuum space  $L^\infty(\mu)$  in Section 4.4.1. The corresponding sequence of priors will turn out to have the form  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}) = \mathcal{L}(\Phi(\mathcal{I}W_n))$ , with  $W_n$  defined in (4.7) and  $\mathcal{I}$  a suitable interpolation map so that  $\mathcal{I}W_n$  is a Gaussian process on  $\mathcal{M}$  that approximates  $W^{\mathcal{M}}$ . Our analysis will then consist of the following steps. We will establish the three conditions (4.14), (4.15), (4.16) for  $W^{\mathcal{M}}$  in Section 4.4.2, followed by a convergence rate analysis of  $\mathcal{I}W_n$  towards  $W^{\mathcal{M}}$  in Section 4.4.3, so that the three conditions are inherited by  $\mathcal{I}W_n$ . The assumptions on the link function  $\Phi$  will then allow us to conclude similar conditions for  $\Phi(\mathcal{I}W_n)$ .

The discrepancy measure  $d_n$  can be taken as the empirical  $L^2$ -norm  $\|\cdot\|_n$  in the regression case (see e.g. [Ghosal and Van der Vaart, 2017, Section 8.3.2]) and the root average square

Hellinger distance in the classification setting [Ghosal and Van Der Vaart, 2007, Section 3], defined as

$$d_{n,H}^2(f, f') = \frac{1}{n} \sum_{i=1}^n \int \left( \sqrt{p_{f,i}} - \sqrt{p_{f',i}} \right)^2 d\gamma. \quad (4.18)$$

In the latter case, one can show that  $d_{n,H}$  is equivalent to  $\|\cdot\|_n$ . Indeed, since the densities are uniformly bounded,  $\|f - f'\|_n$  is bounded above by a multiple of (4.18). Furthermore (4.18) is upper bounded by a multiple of  $\|f - f'\|_n$  by [Ghosal and Van der Vaart, 2017, Lemma 2.8(ii)] given our assumption on  $\Phi$ . Hence this explains the choice of  $\|\cdot\|_n$  in (4.12).

A natural choice for the Banach space in our setting is then  $\mathbb{B} = L^\infty(\mu)$  since (i)  $d_n = \|\cdot\|_n \leq \|\cdot\|_{L^\infty(\mu)}$ ; and (ii)  $B_n^*(f_0, \varepsilon_n) \supset \{f \in \mathcal{F} : \|f - f_0\|_{L^\infty(\mu)} \leq \tilde{C}\varepsilon_n\}$  for some universal constant  $\tilde{C}$ . The first point is clear and an upper bound on the metric entropy (4.14) in  $\|\cdot\|_{L^\infty(\mu)}$  will automatically yield an upper bound in  $\|\cdot\|_n$ . The second point follows from the fact that one can upper bound the two quantities in (4.17) by  $\|f - f_0\|_n$  and hence  $\|f - f_0\|_{L^\infty(\mu)}$  in both the regression and classification setting. Therefore the prior mass condition (4.16) for  $W^{\mathcal{M}}$  in  $L^\infty(\mu)$  balls is sufficient to give the corresponding condition for  $B^*(f_0, \varepsilon_n)$ . This motivates us to consider the continuum Gaussian field defined in (4.8) as an element in  $L^\infty(\mu)$  in Section 4.4.2 and the  $L^\infty(\mu)$  convergence rate in Section 4.4.3.

## 4.4 Proof of the Main Result

In this section we prove Theorem 4.2.1. An important part of the proof is to formalize the convergence of the Gauss-Markov random field  $W_n$  in (4.7), defined in the data cloud  $\mathcal{X}_{N_n}$ , to the Matérn field  $W^{\mathcal{M}}$  in (4.8), defined in  $\mathcal{M}$ . To that end, we introduce an interpolation

of  $W_n$  into  $\mathcal{M}$

$$W_n^{\mathcal{M}} := \sum_{i=1}^{k_{N_n}} \left[ 1 + \lambda_i^{(N_n)} \right]^{-\frac{s}{2}} \xi_i \psi_i^{(N_n)} \circ T_{N_n}, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (4.19)$$

where  $T_{N_n} : \mathcal{M} \rightarrow \{X_1, \dots, X_{N_n}\}$  are transport maps to be specified in Proposition 4.4.5 below. In Section 4.4.1 we show that it suffices to establish posterior contraction with respect to the prior

$$\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}) := \mathcal{L}(\Phi(W_n^{\mathcal{M}}) | \mathcal{X}_{N_n}). \quad (4.20)$$

In Section 4.4.2 we show concentration properties of the limiting Gaussian field  $W^{\mathcal{M}}$ . In Section 4.4.3 we establish the  $L^\infty(\mu)$ -convergence rate of  $W_n^{\mathcal{M}}$  towards  $W^{\mathcal{M}}$ . Combining the facts that  $W_n^{\mathcal{M}}$  is close to  $W^{\mathcal{M}}$  with the contraction properties of  $W^{\mathcal{M}}$ , we deduce that  $\mathcal{L}(W_n^{\mathcal{M}} | \mathcal{X}_{N_n})$  satisfies the three conditions (4.14), (4.15), (4.16). We complete the proof in Section 4.4.4 by lifting these conditions to  $\mathcal{L}(\Phi(W_n^{\mathcal{M}}) | \mathcal{X}_{N_n})$  in both the regression and classification problems.

#### 4.4.1 Reduction via Interpolation

Here we show that in order to establish Theorem 4.2.1 it suffices to prove posterior contraction with respect to the continuum prior  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n})$  defined in (4.20). Let  $\mathcal{I} := \mathcal{I}_{N_n} : L^\infty(\mu_{N_n}) \rightarrow L^\infty(\mu)$  be the interpolation map defined by  $\mathcal{I}u := u \circ T_{N_n}$ . Specifically, we claim that, in order to establish Theorem 4.2.1, it suffices to show that

$$\mathbb{E}_X \mathbb{E}_{f_0} \Pi_n^{\mathcal{M}}(f \in \mathcal{I}(L^\infty(\mu_{N_n})) : \|f - f_0\|_n \geq M_n \varepsilon_n | \mathcal{X}_{N_n}, \mathcal{Y}_n) \xrightarrow{n \rightarrow \infty} 0.$$

The only property of the maps  $T_{N_n}$  in (4.19) that we will use in this subsection is that  $T_{N_n}(X_i) = X_i$ . In order to establish the claim, let  $A_n = \{f \in L^\infty(\mu_{N_n}) : \|f - f_0\|_n \geq M_n \varepsilon_n\}$ .

Since  $A_n \subset \mathcal{I}^{-1}(\mathcal{I}(A_n))$  we have

$$\Pi_n(A_n | \mathcal{X}_{N_n}, \mathcal{Y}_n) \leq \Pi_n(\mathcal{I}^{-1}(\mathcal{I}(A_n)) | \mathcal{X}_{N_n}, \mathcal{Y}_n) = \mathcal{I}_\#[\Pi_n(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n)](\mathcal{I}(A_n)),$$

where  $\mathcal{I}_\#$  denotes push-forward through the map  $\mathcal{I}$ . (Recall that the push-forward of a measure  $\nu$  through the map  $\mathcal{I}$  is defined by the relationship  $(\mathcal{I}_\# \nu)(B) := \nu(\mathcal{I}^{-1}(B))$ .) Therefore, it follows from Lemma 4.4.1 below that

$$\Pi_n(A_n | \mathcal{X}_{N_n}, \mathcal{Y}_n) \leq \Pi_n^{\mathcal{M}}(\mathcal{I}(A_n) | \mathcal{X}_{N_n}, \mathcal{Y}_n), \quad (4.21)$$

and hence it suffices to bound the right-hand side of (4.21). Since  $T_{N_n}(X_i) = X_i$ , we have

$$\begin{aligned} \mathcal{I}(A_n) &= \{f \circ T_{N_n} : \|f - f_0\|_n \geq M_n \varepsilon_n\} = \{f \circ T_{N_n} : \|f \circ T_{N_n} - f_0\|_n \geq M_n \varepsilon_n\} \\ &= \{f \in \mathcal{I}(L^\infty(\mu_{N_n})) : \|f - f_0\|_n \geq M_n \varepsilon_n\}, \end{aligned}$$

and the claim is established.

**Lemma 4.4.1.** *It holds that  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n) = \mathcal{I}_\#[\Pi_n(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n)]$ .*

*Proof. Step 1:* First we show that  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}) = \mathcal{I}_\#[\Pi_n(\cdot | \mathcal{X}_{N_n})]$ . Since  $\mathcal{I}$  is linear, we have that  $W_n^{\mathcal{M}} = \mathcal{I}W_n$ . Then observe that  $\mathcal{I}(\Phi(W_n)) = \Phi(\mathcal{I}(W_n))$ . Indeed, for  $x \in T_{N_n}^{-1}(\{X_i\})$ , we have

$$\mathcal{I}(\Phi(W_n))(x) = \Phi(W_n)(X_i) = \Phi(W_n(X_i)) = \Phi(\mathcal{I}(W_n)(x)) = \Phi(\mathcal{I}(W_n))(x).$$

Therefore, we have that

$$\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}) = \mathcal{L}(\Phi(W_n^{\mathcal{M}}) | \mathcal{X}_{N_n}) = \mathcal{L}(\mathcal{I}(\Phi(W_n)) | \mathcal{X}_{N_n}).$$

Finally, note that  $\mathcal{I}_\#[\Pi_n(\cdot | \mathcal{X}_{N_n})] = \mathcal{L}(\mathcal{I}(\Phi(W_n)) | \mathcal{X}_{N_n})$ , since

$$\mathcal{I}_\# \Pi_n(B | \mathcal{X}_{N_n}) = \Pi_n(\mathcal{I}^{-1}(B) | \mathcal{X}_{N_n}) = \mathbb{P}(\Phi(W_n) \in \mathcal{I}^{-1}(B) | \mathcal{X}_{N_n}) = \mathbb{P}(\mathcal{I}(\Phi(W_n)) \in B | \mathcal{X}_{N_n}).$$

**Step 2:** Now we show that  $\Pi_n^{\mathcal{M}}(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n) = \mathcal{I}_\#[\Pi_n(\cdot | \mathcal{X}_{N_n}, \mathcal{Y}_n)]$ . By definition of pushforward measure, it suffices to show that, for any measurable  $B$ ,

$$\Pi_n^{\mathcal{M}}(B | \mathcal{X}_{N_n}, \mathcal{Y}_n) = \Pi_n(\mathcal{I}^{-1}(B) | \mathcal{X}_{N_n}, \mathcal{Y}_n).$$

Using Step 1, the left-hand side equals

$$\Pi_n^{\mathcal{M}}(B | \mathcal{X}_{N_n}, \mathcal{Y}_n) = \frac{\int_B \prod_{i=1}^n L_{Y_i | X_i}(f) d\mathcal{I}_\#[\Pi_n(f | \mathcal{X}_{N_n})]}{\int_{\mathcal{I}(L^\infty(\mu_{N_n}))} \prod_{i=1}^n L_{Y_i | X_i}(f) d\mathcal{I}_\#[\Pi_n(f | \mathcal{X}_{N_n})]}, \quad (4.22)$$

where  $L_{Y_i | X_i}(f)$  is given in (4.2) and (4.3). Note that pointwise values of  $f$  are well-defined since  $\mathcal{I}_\# \Pi_n$  is supported on  $\mathcal{I}(L^\infty(\mu_{N_n}))$ . By the change-of-variable formula for pushforward measures,

$$(4.22) = \frac{\int_{\mathcal{I}^{-1}(B)} \prod_{i=1}^n L_{Y_i | X_i} \circ \mathcal{I}(f_n) d\Pi_n(f_n | \mathcal{X}_{N_n})}{\int_{L^\infty(\mu_{N_n})} \prod_{i=1}^n L_{Y_i | X_i} \circ \mathcal{I}(f_n) d\Pi_n(f_n | \mathcal{X}_{N_n})},$$

which equals (4.1) with  $B$  replaced by  $\mathcal{I}^{-1}(B)$ , by noticing that  $L_{Y_i | X_i} \circ \mathcal{I}(f_n)$  is exactly the same as in (4.2) and (4.3). The result follows.  $\square$

#### 4.4.2 Regularity and Contraction Properties of the Limiting Field

In this subsection we study the limit Gaussian measure  $\pi = \mathcal{N}(0, (I - \Delta)^{-s}) = \mathcal{L}(W^{\mathcal{M}})$ .

Recall that the samples admit Karhunen-Loève expansion

$$W^{\mathcal{M}} = \sum_{i=1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (4.23)$$

where  $(\lambda_i, \psi_i)$ 's are the eigenpairs of the Laplace-Beltrami operator  $-\Delta$  on  $\mathcal{M}$ . Notice that a larger  $s$  leads to a faster decay of the coefficients and hence more regular sample paths. From Weyl's law that  $\lambda_i \asymp i^{\frac{2}{m}}$  (see e.g. [Canzani, 2013, Theorem 72]), setting  $s > \frac{m}{2}$  makes  $\pi$  a well-defined measure on  $L^2(\mu)$ . To ensure almost sure continuity of the samples, so that point evaluations are well-defined and  $\pi$  is supported over  $L^\infty(\mu)$ , we need the stronger assumption that  $s > m - \frac{1}{2}$ .

**Lemma 4.4.2.** *If  $s > m - \frac{1}{2}$ , then samples of  $\pi$  are almost surely continuous.*

*Proof.* By [Lang et al., 2016, Corollary 4.5], it suffices to show that

$$\mathbb{E}_\xi |W^\mathcal{M}(x) - W^\mathcal{M}(y)|^2 \leq C d_\mathcal{M}(x, y)^\eta,$$

for some  $\eta > 0$ . We have

$$\begin{aligned} \mathbb{E}_\xi |W^\mathcal{M}(x) - W^\mathcal{M}(y)|^2 &= \mathbb{E}_\xi \left( \sum_{i=1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i (\psi_i(x) - \psi_i(y)) \right)^2 \\ &= \sum_{i=1}^{\infty} (1 + \lambda_i)^{-s} |\psi_i(x) - \psi_i(y)|^2, \end{aligned}$$

since  $\mathbb{E}_\xi \xi_i \xi_j = 0$  for  $i \neq j$ . By Proposition 4.4.12,

$$\begin{aligned} |\psi_i(x) - \psi_i(y)| &\leq C \lambda_i^{\frac{m-1}{4}}, \\ |\psi_i(x) - \psi_i(y)| &\leq \|\nabla \psi_i\|_{L^\infty(\mu)} d_\mathcal{M}(x, y) \leq C \lambda_i^{\frac{m+1}{2}} d_\mathcal{M}(x, y). \end{aligned}$$



Using Weyl's law, we have

$$\begin{aligned}
\mathbb{E}_\xi |W^\mathcal{M}(x) - W^\mathcal{M}(y)|^2 &\lesssim \sum_{i=1}^{\infty} (1 + \lambda_i)^{-s} \min \left\{ \lambda_i^{m+1} d_{\mathcal{M}}(x, y)^2, \lambda_i^{\frac{m-1}{2}} \right\} \\
&\lesssim \sum_{i=1}^{\infty} i^{-\frac{2s}{m}} \min \left\{ i^{\frac{2m+2}{m}} d_{\mathcal{M}}(x, y)^2, i^{\frac{m-1}{m}} \right\} \\
&\lesssim d_{\mathcal{M}}(x, y)^2 \int_1^K z^{-\frac{2s+2m+2}{m}} dz + \int_K^\infty z^{-\frac{2s+m-1}{m}} dz.
\end{aligned}$$

Setting  $K = d_{\mathcal{M}}(x, y)^{-\frac{2m}{m+3}}$ , we have

$$\mathbb{E}_\xi |W^\mathcal{M}(x) - W^\mathcal{M}(y)|^2 \lesssim d_{\mathcal{M}}(x, y)^{\frac{4s-4m+2}{m+3}}.$$

The result follows since  $s > m - \frac{1}{2}$ . □

**Remark 4.4.3.** *The argument above relies on the worst case  $L^\infty(\mu)$  bound on the eigenfunctions of  $-\Delta$  given in Proposition 4.4.12. If the eigenfunctions are uniformly bounded (which holds for the family of flat manifolds [Toth and Zelditch, 2002] that includes, for instance, the torus), then continuity can be guaranteed for  $s > \frac{m}{2}$ .*

From now on, we shall consider  $\pi$  as a measure over  $L^\infty(\mu)$  with continuous sample paths, where pointwise evaluation is well-defined. Using the series representation (4.23), the reproducing kernel Hilbert space  $\mathbb{H}$  associated with  $\pi$  has the following characterization

$$\mathbb{H} = \left\{ h = \sum_{i=1}^{\infty} h_i \psi_i : \|h\|_{\mathbb{H}}^2 := \sum_{i=1}^{\infty} h_i^2 (1 + \lambda_i)^s < \infty \right\}. \quad (4.24)$$

From the general theory in van der Vaart and van Zanten [2008a], the concentration properties of  $W$  can be characterized by the concentration function

$$\phi_{w_0}(\varepsilon) := \inf_{h \in \mathbb{H}: \|h - w_0\|_{L^\infty(\mu)} < \varepsilon} \|h\|_{\mathbb{H}}^2 - \log \mathbb{P}(\|W\|_{L^\infty(\mu)} < \varepsilon), \quad (4.25)$$

where  $w_0$  belongs to the support of  $\pi$ . For our purpose,  $w_0$  will be set as  $\Phi^{-1}(f_0)$ . The three conditions (4.14), (4.15) and (4.16) hold for the  $\varepsilon_n$  that satisfies  $\phi_{w_0}(\varepsilon_n) \leq n\varepsilon_n^2$ .

Before stating our main result in this subsection, we follow Castillo et al. [2014] and Coulhon et al. [2012] to define a Besov space which will be needed to characterize the regularity of  $w_0$ . Let  $\Psi$  be an even function in the Schwartz space  $\mathcal{S}(\mathbb{R})$  with

$$0 \leq \Psi \leq 1, \quad \Psi \equiv 1 \text{ on } \left[-\frac{1}{2}, \frac{1}{2}\right], \quad \text{supp}(\Psi) \subset [-1, 1].$$

Define the Besov space

$$B_{\infty, \infty}^{\beta} := \left\{ w : \|w\|_{B_{\infty, \infty}^{\beta}}^{\beta} := \sup_{j \in \mathbb{N}} 2^{\beta j} \|\Psi_j(\sqrt{\Delta})w(\cdot) - w(\cdot)\|_{L^{\infty}(\mu)} < \infty \right\},$$

where  $\Psi_j(\cdot) = \Psi(2^{-j}\cdot)$  and, for  $w = \sum_{i=1}^{\infty} w_i \psi_i$ ,

$$\Psi_j(\sqrt{\Delta})w := \sum_{i=1}^{\infty} \Psi_j(\sqrt{\lambda_i})w_i \psi_i.$$

It was shown in [Coulhon et al., 2012, Proposition 6.2] that the definition is independent of the choice of  $\Psi$ .

**Theorem 4.4.4.** *Suppose  $w_0 \in B_{\infty, \infty}^{\beta}$ . Consider the prior  $\pi = \mathcal{N}(0, (I - \Delta)^{-s})$  with  $s > \beta \wedge m - \frac{1}{2}$ . Then there exists sets  $B_n \subset L^{\infty}(\mu)$  so that*

$$\begin{aligned} \log N(3\varepsilon_n, B_n, \|\cdot\|_{L^{\infty}(\mu)}) &\leq 6Cn\varepsilon_n^2, \\ \pi(B_n^c) &\leq e^{-Cn\varepsilon_n^2}, \\ \pi(\|w - w_0\|_{L^{\infty}(\mu)} < 2\varepsilon_n) &\geq e^{-n\varepsilon_n^2}, \end{aligned}$$

where  $\varepsilon_n$  is a multiple of  $n^{-\frac{(s-\frac{m}{2})\wedge\beta}{2s}} (\log n)^{\frac{(s-\frac{m}{2})\wedge\beta}{4s-2m}}$  and  $C > 1$  is a constant satisfying  $e^{-Cn\varepsilon_n^2} < \frac{1}{2}$ .

*Proof.* As noted above, it suffices to show that the above  $\varepsilon_n$  satisfies  $\phi_{w_0}(\varepsilon_n) \leq n\varepsilon_n^2$ , where  $\phi_{w_0}$  is defined in (4.25). To bound the small ball probability, it suffices by general results from Li and Linde [1999] to bound the metric entropy of  $\mathbb{H}_1$ , the unit ball in  $\mathbb{H}$ . Notice that (4.24) has the form of a Sobolev ball of regularity  $s$ . Classical results for Sobolev spaces on  $\mathbb{R}^m$  such as [Edmunds and Triebel, 1996, Theorem 3.3.2] have shown that the entropy in the  $L^\infty$  metric is bounded by a multiple of  $\varepsilon^{-m/s}$ . The manifold case has been treated in [Kushpel and Levesley, 2015, Theorem 3.12] assuming homogeneity, which includes an additional logarithmic factor, i.e.,

$$\log N(\varepsilon, \mathbb{H}_1, \|\cdot\|_{L^\infty(\mu)}) \lesssim \varepsilon^{-\frac{m}{s}} \left( \log \frac{1}{\varepsilon} \right)^{\frac{1}{2}}. \quad (4.26)$$

The original theorem was about entropy numbers but can be translated to the above statement on metric entropy and its proof suggested that when  $p = 2$  the theorem holds with  $s > \frac{m}{2}$ . In particular their Sobolev class is defined as

$$I_s U_2 := \left\{ h = h_1 + \sum_{i=2}^{\infty} \lambda_i^{-\frac{s}{2}} h_i \psi_i : \sum_{i=1}^{\infty} h_i^2 \leq 1 \right\},$$

which is compatible with our  $\mathbb{H}$ . (We remark that although Kushpel and Levesley [2015] defined their Sobolev class with a further mean-zero condition, their proof actually applied to the general case.) Now by (4.26) and [Li and Linde, 1999, Theorem 1.2] we have

$$-\log \mathbb{P}(\|w\|_{L^\infty(\mu)} < \varepsilon) \lesssim \varepsilon^{-\frac{2m}{2s-m}} \left( \log \frac{1}{\varepsilon} \right)^{\frac{s}{2s-m}}. \quad (4.27)$$

For the decentering function, let  $C_0 = \|w_0\|_{B_{\infty,\infty}^\beta}$  and consider  $h = \Psi_J(\sqrt{\Delta})w_0$  with  $J$  large enough so that  $C_0 2^{-\beta J} \leq \varepsilon$ . Since  $w_0 \in B_{\infty,\infty}^\beta$ , we have

$$\|\Psi_j(\sqrt{\Delta})w_0 - w_0\|_{L^\infty(\mu)} \leq C_0 2^{-\beta j}, \quad j \in \mathbb{N}. \quad (4.28)$$

In particular  $\|h - w_0\|_{L^\infty(\mu)} \leq C_0 2^{-\beta J} \leq \varepsilon$ . Suppose  $w_0$  has the representation  $w_0 = \sum_{i=1}^{\infty} w_i \psi_i$ . We then have  $h = \sum \Psi_J(\sqrt{\lambda_i}) w_i \psi_i \in \mathbb{H}$  since  $\Psi_J(\sqrt{\lambda_i}) = 0$  for  $\sqrt{\lambda_i} > 2^J$ . Moreover, since  $\Psi_J \leq 1$ ,

$$\|h\|_{\mathbb{H}}^2 \leq \sum_{\sqrt{\lambda_i} \leq 2^J} w_i^2 (1 + \lambda_i)^s \leq \sum_{\sqrt{\lambda_i} \leq 1} 2^s w_i^2 + \sum_{j=1}^J \sum_{2^{j-1} < \sqrt{\lambda_i} \leq 2^j} w_i^2 (1 + \lambda_i)^s. \quad (4.29)$$

By (4.28), we have

$$\begin{aligned} \sum_{2^{j-1} < \sqrt{\lambda_i} \leq 2^j} w_i^2 (1 + \lambda_i)^s &\lesssim 2^{2js} \sum_{2^{j-1} < \sqrt{\lambda_i}} w_i^2 \leq 2^{2js} \|\Psi_{j-1}(\sqrt{\Delta}) w_0 - w_0\|_2^2 \\ &\lesssim 2^{2js} \|\Psi_{j-1}(\sqrt{\Delta}) w_0 - w_0\|_\infty^2 \lesssim 2^{2(s-\beta)j}. \end{aligned}$$

Therefore recalling that  $C_0 2^{-\beta J} \leq \varepsilon$ ,

$$(4.29) \lesssim \sum_{\sqrt{\lambda_i} \leq 1} 2^s w_i^2 + \sum_{j=1}^J 2^{2(s-\beta)j} \lesssim 2^{2(s-\beta)J} \lesssim \varepsilon^{-\frac{2(s-\beta)}{\beta}},$$

since the first sum remains bounded as  $J \rightarrow \infty$ . Combining the above with (4.27), we get

$$\phi_{w_0}(\varepsilon) \lesssim \varepsilon^{-\frac{2m}{2s-m}} \left( \log \frac{1}{\varepsilon} \right)^{\frac{s}{2s-m}} + \varepsilon^{-\frac{2(s-\beta)}{\beta}},$$

which implies

$$\frac{\phi_{w_0}(\varepsilon)}{\varepsilon^2} \lesssim \varepsilon^{-\frac{2s}{s-\frac{m}{2}}} \left( \log \frac{1}{\varepsilon} \right)^{\frac{s}{2s-m}} + \varepsilon^{-\frac{2s}{\beta}}.$$

Therefore by choosing

$$\varepsilon_n = C n^{-\frac{(s-\frac{m}{2}) \wedge \beta}{2s}} (\log n)^{\frac{(s-\frac{m}{2}) \wedge \beta}{4s-2m}}$$

for a large enough constant  $C$ , the condition  $\phi_{w_0}(\varepsilon_n) \leq n\varepsilon_n^2$  is satisfied. The three assertions then follow from [van der Vaart and van Zanten, 2008a, Theorem 2.1].  $\square$

#### 4.4.3 Convergence of Gaussian Fields in $L^\infty(\mu)$

In this subsection we establish  $L^\infty(\mu)$  convergence of  $W_n^{\mathcal{M}}$  to  $W^{\mathcal{M}}$ . We shall denote  $N := N_n$  throughout the rest of this subsection to simplify our notation. Recall that

$$\begin{aligned} W_n^{\mathcal{M}} &= \sum_{i=1}^{k_N} \left[1 + \lambda_i^{(N)}\right]^{-\frac{s}{2}} \xi_i \psi_i^{(N)} \circ T_N, & \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \\ W^{\mathcal{M}} &= \sum_{i=1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i, & \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \end{aligned} \tag{4.30}$$

In order to show the convergence, the transportation maps  $T_N$  will be assumed to be close to the identity in the sense made precise in the following proposition from [García Trillos et al., 2019a, Theorem 2].

**Proposition 4.4.5.** *For  $\gamma > 1$ , there exists a transportation map  $T_N : \mathcal{M} \rightarrow \{X_1, \dots, X_N\}$  satisfying  $T_{N_n}(X_i) = X_i$  so that, with probability  $1 - O(N^{-\gamma})$ ,*

$$\rho_N := \sup_{x \in \mathcal{M}} d_{\mathcal{M}}(x, T_N(x)) \lesssim \frac{(\log N)^{p_m}}{N^{1/m}}, \tag{4.31}$$

where  $p_m = \frac{3}{4}$  if  $m = 2$  and  $p_m = \frac{1}{m}$  otherwise.

The transport map  $T_{N_n}$  is a measure-preserving transformation in the sense that  $\mu(T_{N_n}^{-1}(U)) = \mu_{N_n}(U)$  for all  $U \subset \mathcal{M}$  measurable. The additional requirement  $T_N(X_i) = X_i$  that was absent in [García Trillos et al., 2019a, Theorem 2] is valid here since modifying  $T_N$  on a set of  $\mu$ -measure 0 does not affect the measure preserving property and (4.31) still holds in this case. The scaling in (4.31) can be thought as the resolution of the point cloud, and is important in suitably defining the choice of connectivity  $\zeta_N$ , as we shall see.

The main result of this subsection is the following.

**Theorem 4.4.6.** *Let  $\delta > 0$  be arbitrary.*

1.  $\frac{3}{2}m - \frac{1}{2} < s \leq \frac{9}{2}m + \frac{5}{2}$ . Set

$$\begin{cases} \zeta_N \asymp N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{(m+4+\delta)(6+\frac{6}{m})}} (\log N)^{-\frac{mpm}{12m+12}}, & m \leq 4, \\ \zeta_N \asymp N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{12m+12}} (\log N)^{-\frac{mpm}{12m+12}}, & m \geq 5. \end{cases}$$

Then, with probability tending to 1,

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \begin{cases} N^{-\frac{2s-3m+1}{(m+4+\delta)(6m+6)}} (\log N)^{\frac{pm(2s-3m+1)}{12m+12}}, & m \leq 4, \\ N^{-\frac{2s-3m+1}{m(12m+12)}} (\log N)^{\frac{pm(2s-3m+1)}{12m+12}}, & m \geq 5. \end{cases}$$

2.  $s > \frac{9}{2}m + \frac{5}{2}$ . Set

$$\begin{cases} \zeta_N \asymp N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{m}{(m+4+\delta)(2s-3m+1)}} (\log N)^{-\frac{mpm}{4s-6m+2}}, & m \leq 4, \\ \zeta_N \asymp N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{4s-6m+2}} (\log N)^{-\frac{mpm}{4s-6m+2}}, & m \geq 5. \end{cases}$$

Then, with probability tending to 1,

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \begin{cases} N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & m \leq 4, \\ N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & m \geq 5. \end{cases}$$

We remark that the statement “with probability tending to 1” refers to the randomness coming from the  $X_i$ ’s. By solving for  $N$  so that the rate matches  $n^{-1}$ , we get the following.

**Corollary 4.4.7.** *Let  $\delta > 0$  be arbitrary. Let*

$$N \asymp n^{\alpha_1} (\log n)^{\alpha_2},$$

where

$$\left\{ \begin{array}{ll} \alpha_1 = \frac{(m+4+\delta)(6m+6)}{2s-3m+1}, \quad \alpha_2 = \frac{p_m(m+4+\delta)}{2}, & \text{if } m \leq 4, \quad \frac{3}{2}m - \frac{1}{2} < s \leq \frac{9}{2}m + \frac{5}{2}, \\ \alpha_1 = \frac{m(12m+12)}{2s-3m+1}, \quad \alpha_2 = mp_m, & \text{if } m \geq 5, \quad \frac{3}{2}m - \frac{1}{2} < s \leq \frac{9}{2}m + \frac{5}{2}, \\ \alpha_1 = m + 4 + \delta, \quad \alpha_2 = \frac{p_m(m+4+\delta)}{2}, & \text{if } m \leq 4, \quad s > \frac{9}{2}m + \frac{5}{2}, \\ \alpha_1 = 2m, \quad \alpha_2 = mp_m, & \text{if } m \geq 5, \quad s > \frac{9}{2}m + \frac{5}{2}. \end{array} \right.$$

Let  $\zeta_N$  and  $k_N$  have the same scaling as in Theorem 4.4.6. Then, with probability tending to 1,

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim n^{-1}.$$

**Remark 4.4.8.** For flat manifolds the results of Theorem 4.4.6 and Corollary 4.4.7 can be shown to hold for  $s > m$  with corresponding modifications in the scaling of the parameters.

The key to show the above results is to derive convergence rates for

$$|\lambda_i^{(N)} - \lambda_i|, \quad \text{and} \quad \|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^\infty(\mu)}.$$

We shall build our analysis on the existing results from the literature. Recall that  $\zeta_N$  is the connectivity of the graph and the resolution  $\rho_N$  defined in (4.31). Assuming we are in the event that (4.31) holds, the following results from [Sanz-Alonso and Yang, 2022a, Theorem 4.6 & 4.7] bound the spectral approximations.

**Proposition 4.4.9** (Spectral Approximation). *Suppose  $\rho_N \ll \zeta_N$  and  $\zeta_N \sqrt{\lambda_{k_N}} \ll 1$  for  $N$  large. Then there exists orthonormalized eigenfunctions  $\{\psi_i^{(N)}\}_{i=1}^N$  and  $\{\psi_i\}_{i=1}^\infty$  so that, for*

$i = 1, \dots, k_N,$

$$|\lambda_i^{(N)} - \lambda_i| \lesssim \lambda_i \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right), \quad (4.32)$$

$$\|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^2(\mu)}^2 \lesssim i^3 \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right). \quad (4.33)$$

**Remark 4.4.10.** *Proposition 4.4.9 bounds the eigenfunction approximation in  $L^2(\mu)$  norm and we need to lift such result to the  $L^\infty(\mu_N)$  norm. We remark that  $L^\infty(\mu_N)$  convergence rates have already been established in the literature (see e.g. Dunson et al. [2021] and Calder et al. [2022]). However these results do not contain an explicit proportion constant in terms of the index  $i$  as in (4.33). Instead of going through their details to find the explicit constants, which would make our presentation much more involved, we directly show  $L^\infty(\mu_N)$  bounds based on Proposition 4.4.9, by following the same idea as in Calder et al. [2022]. Since we build our results from (4.33), which is not the sharpest bound in the literature, our results in Theorem 4.4.6 and Corollary 4.4.7 suffer the same drawback. Nevertheless, the goal of this chapter is to demonstrate the idea that unlabeled data helps and hence the finding sharpest scaling of  $N_n$  is less essential.*

Below we record four results from [Calder and García Trillos, 2022, Theorem 3.3], [Donnelly, 2006, Theorem 1.2] and [Xu, 2006, Equation (2.10)], [Calder et al., 2020, Corollary 2.5] that will be needed.

**Proposition 4.4.11** (Pointwise Error of  $\Delta_N$ ). *Let  $f \in \mathcal{C}^3(\mathcal{M})$ . Then with probability  $1 - 2n \exp(-cN\zeta_N^{m+4})$ ,*

$$\|\Delta_N f(x) - \Delta f(x)\|_{L^\infty(\mu_N)} \leq C(1 + \|f\|_{\mathcal{C}^3(\mathcal{M})})\zeta_N.$$

**Proposition 4.4.12** (Bounds on Eigenfunctions and Their Gradients). *Let  $\psi$  be an  $L^2(\mu)$ -*



normalized eigenfunction of  $-\Delta$  associated with eigenvalue  $\lambda \neq 0$ . Then, for  $k \in \mathbb{N}$ ,

$$\begin{aligned}\|\psi\|_{L^\infty(\mu)} &\leq C\lambda^{\frac{m-1}{4}}, \\ \|\nabla^k \psi\|_{L^\infty(\mu)} &\leq C\lambda^{k+\frac{m-1}{2}}.\end{aligned}$$

**Proposition 4.4.13** ( $L^\infty$  Bound In Terms of  $L^1$  Bound). *Suppose  $\zeta_N \leq \frac{c}{\Lambda+1}$  where  $c$  is a sufficiently small constant depending only on  $\mathcal{M}$ . Then with probability  $1 - C\zeta_N^{-6m} \exp(-cN\zeta_N^{m+4}) - 2N \exp(-cN(\Lambda+1)^{-m})$*

$$\|u\|_{L^\infty(\mu_N)} \leq C(\Lambda+1)^{m+1} \|u\|_{L^1(\mu_N)}$$

for all  $u \in L^2(\mu_N)$  with  $\lambda_u < \Lambda$ , where

$$\lambda_u := \frac{\|\Delta_N u\|_{L^\infty(\mu_N)}}{\|u\|_{L^\infty(\mu_N)}}. \quad (4.34)$$

Notice that the high probability condition is satisfied only if

$$\zeta_N^{-6m} \exp(-cN\zeta_N^{m+4}) - 2N \exp(-cN(\Lambda+1)^{-m}) \rightarrow 0$$

and this requires some care in setting the scaling of  $\zeta_N$  and  $k_N$ . In particular a sufficient condition for  $\zeta_N$  is  $\zeta_N \gtrsim N^{-\frac{1}{m+4+\delta}}$  for an arbitrarily small  $\delta > 0$ . On the other hand the condition for  $\Lambda$  reduces to the scaling of  $k_N$ . Indeed if  $u = \psi_i^{(N)}$ , then  $\lambda_u = \lambda_i^{(N)}$  and therefore  $\Lambda$  can be chosen as a multiple of  $\lambda_i$  given the spectral approximation in Proposition 4.4.9. We will show that the same choice of  $\Lambda$  suffices when  $u = \psi_i^{(N)} - \psi_i$ . Since we are interested in bounding the first  $k_N$  eigenfunctions, we can set  $\Lambda$  to be multiple of  $\lambda_{K_N}$ . Therefore in order to make  $2N \exp(-cN(\Lambda+1)^{-m}) \ll 1$ , it is sufficient to have  $\lambda_{k_N} \lesssim N^{\frac{1-\delta}{m}}$ , i.e., if  $k_N \lesssim N^{\frac{1-\delta}{2}}$  by Weyl's law. Therefore we should keep the following in

mind when we set the scaling later

$$N^{-\frac{1}{m+4+\delta}} \lesssim \zeta_N, \quad k_N \lesssim N^{\frac{1-\delta}{m}}, \quad \zeta_N k_N^{\frac{2}{m}} \lesssim 1, \quad (4.35)$$

where the third requirement corresponds to the condition that  $\zeta_N(\Lambda + 1) \leq c$ .

Now we are ready to show the  $L^\infty(\mu_N)$  bound of eigenfunction approximations.

**Lemma 4.4.14.** *Under the same conditions and the intersection of the high probability events as in Propositions 4.4.9 and 4.4.13 we have, for  $i = 1, \dots, k_N$ ,*

$$\|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^\infty(\mu)} \lesssim \lambda_i^{m+1} i^{\frac{3}{2}} \sqrt{\frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i}}.$$

*Proof.* We will first modify the  $L^2(\mu)$ -bound (4.33) to an  $L^2(\mu_N)$ -bound using the regularity of eigenfunctions of  $-\Delta$ , which is then lifted to an  $L^\infty(\mu_N)$ -bound using Proposition 4.4.11 and 4.4.13. Finally we use regularity of the eigenfunctions again to transfer the  $L^\infty(\mu_N)$ -bound to an  $L^\infty(\mu)$ -bound. To start, we notice that the transport  $T_N$  induces a partition of  $\mathcal{M}$  by the sets  $\{U_i = T_N^{-1}(\{X_i\})\}_{i=1}^N$ . Furthermore, the measure preserving property gives  $\mu(U_i) = \frac{1}{N}$  and Proposition 4.4.5 implies that  $U_i \subset B_{\mathcal{M}}(X_i, \rho_N)$ . We then have

$$\begin{aligned} \|\psi_i^{(N)} - \psi_i\|_{L^2(\mu_N)}^2 &= \frac{1}{N} \sum_{i=1}^N |\psi_i^{(N)}(X_i) - \psi_i(X_i)|^2 = \sum_{i=1}^N \int_{U_i} |\psi_i^{(N)}(X_i) - \psi_i(X_i)|^2 d\mu(x) \\ &\leq \sum_{i=1}^N \int_{U_i} 2|\psi_i^{(N)}(X_i) - \psi_i(x)|^2 + 2|\psi_i(x) - \psi_i(X_i)|^2 d\mu(x), \end{aligned}$$

where

$$\int_{U_i} |\psi_i(x) - \psi_i(X_i)|^2 d\mu(x) \leq \frac{1}{N} \|\nabla \psi_i\|_\infty^2 \rho_N^2.$$

Hence by Proposition 4.4.12 and Weyl's law

$$\|\psi_i^{(N)} - \psi_i\|_{L^2(\mu_N)}^2 \leq 2\|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^2(\mu)}^2 + Ci^{\frac{m+1}{m}} \rho_N^2.$$

Since  $i^{\frac{m}{m+1}} \rho_N^2$  is of higher order than (4.33), we conclude that

$$\|\psi_i^{(N)} - \psi_i\|_{L^2(\mu_N)}^2 \lesssim i^3 \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right). \quad (4.36)$$

Now let  $g = \psi_i^{(N)} - \tilde{\psi}_i$  with  $\tilde{\psi}_i = \psi_i|_{\{X_1, \dots, X_N\}}$ . We have by Proposition 4.4.11

$$\begin{aligned} \|\Delta_N g\|_{L^\infty(\mu_N)} &\leq \|\Delta_N \psi_i^{(N)} - \Delta \psi_i\|_{L^\infty(\mu_N)} + \|\Delta \psi_i - \Delta_N \tilde{\psi}_i\|_{L^\infty(\mu_N)} \\ &\leq \|\lambda_i^{(N)} \psi_i^{(N)} - \lambda_i \psi_i\|_{L^\infty(\mu_N)} + C(1 + \|\psi_i\|_{\mathcal{C}^3(\mathcal{M})}) \zeta_N \\ &\leq \lambda_i^{(N)} \|\psi_i^{(N)} - \psi_i\|_{L^\infty(\mu_N)} + |\lambda_i^{(N)} - \lambda_i| \|\psi_i\|_{L^\infty(\mu_N)} + C(1 + \|\psi_i\|_{\mathcal{C}^3(\mathcal{M})}) \zeta_N \\ &\leq \lambda_i^{(N)} \|g\|_{L^\infty(\mu_N)} + C \lambda_i^{\frac{m+5}{2}} \left( \frac{\rho_N}{\zeta_N} + \zeta_N \right), \end{aligned}$$

where we have used (4.32) and Proposition 4.4.12 in the last step. Therefore recalling the definition in (4.34), we have

$$\lambda_g \leq \lambda_i^{(N)} + C \|g\|_{L^\infty(\mu_N)}^{-1} \lambda_i^{\frac{m+5}{2}} \left( \frac{\rho_N}{\zeta_N} + \zeta_N \right).$$

If  $\lambda_g \geq \lambda_i^{(N)} + 1$ , then we have

$$\|g\|_{L^\infty(\mu_N)} \leq C \lambda_i^{\frac{m+5}{2}} \left( \frac{\rho_N}{\zeta_N} + \zeta_N \right). \quad (4.37)$$

Otherwise if  $\lambda_g \leq \lambda_i^{(N)} + 1$ , then Proposition 4.4.13 and (4.36) implies

$$\|g\|_{L^\infty(\mu_N)} \leq C \left[ \lambda_i^{(N)} + 1 \right]^{m+1} \|g\|_{L^1(\mu_N)} \leq C \lambda_i^{m+1} \|g\|_{L^2(\mu_N)} \leq C \lambda_i^{m+1} i^{\frac{3}{2}} \sqrt{\frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i}}, \quad (4.38)$$

where we have used the fact that  $\lambda_i^{(N)} \leq C \lambda_i$  and  $\|g\|_{L^1(\mu_N)} \leq \|g\|_{L^2(\mu_N)}$ . Comparing (4.38) with (4.37), we see that the error in (4.38) dominates. To finish, we need again to lift the  $L^\infty(\mu_N)$ -bound to a  $L^\infty(\mu)$ -bound using regularity of the  $\psi_i$ 's. Notice that

$$\begin{aligned} \|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^\infty(\mu)} &= \max_{1 \leq j \leq N} \sup_{x \in U_j} |\psi_i^{(N)} \circ T_N(x) - \psi_i(x)| \\ &\leq \max_{1 \leq j \leq N} \sup_{x \in U_j} \left( |\psi_i^{(N)}(X_j) - \psi_i(X_j)| + |\psi_i(X_j) - \psi_i(x)| \right) \\ &\leq \|\psi_i^{(N)} - \psi_i\|_{L^\infty(\mu_N)} + \|\nabla \psi_i\|_{L^\infty(\mu)} \rho_N \\ &\lesssim \lambda_i^{m+1} i^{\frac{3}{2}} \sqrt{\frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i}}, \end{aligned}$$

where in the last step  $\|\nabla \psi_i\|_{L^\infty(\mu)} \rho_N$  is a higher order term that we drop.  $\square$

Now we are finally ready to show the  $L^\infty(\mu)$  convergence of  $W_n^{\mathcal{M}}$  towards  $W^{\mathcal{M}}$ .

*Proof of Theorem 4.4.6.* Recall that

$$\begin{aligned} W_n^{\mathcal{M}} &= \sum_{i=1}^{k_N} \left[ 1 + \lambda_i^{(N)} \right]^{-\frac{s}{2}} \xi_i \psi_i^{(N)} \circ T_N, \\ W^{\mathcal{M}} &= \sum_{i=1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i. \end{aligned}$$

Consider two intermediate quantities:

$$\begin{aligned}\widetilde{W}_n^{\mathcal{M}} &= \sum_{i=1}^{k_N} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i^{(N)} \circ T_N, \\ \widehat{W}_n^{\mathcal{M}} &= \sum_{i=1}^{k_N} (1 + \lambda_i)^{-\frac{s}{2}} \xi_i \psi_i.\end{aligned}$$

By the triangle inequality,

$$\begin{aligned}\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 &\leq \mathbb{E}_\xi \left( \|W_n^{\mathcal{M}} - \widetilde{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)} + \|\widetilde{W}_n^{\mathcal{M}} - \widehat{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)} + \|\widehat{W}_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)} \right)^2 \\ &\leq 3 \left( \mathbb{E}_\xi \|W_n^{\mathcal{M}} - \widetilde{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)}^2 + \mathbb{E}_\xi \|\widetilde{W}_n^{\mathcal{M}} - \widehat{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)}^2 + \mathbb{E}_\xi \|\widehat{W}_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \right),\end{aligned}$$

so it suffices to bound each term. First, by Proposition 4.4.12 and Weyl's law,

$$\begin{aligned}\mathbb{E}_\xi \|\widehat{W}_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 &\leq \mathbb{E}_\xi \left[ \sum_{i=k_N+1}^{\infty} \sum_{j=k_N+1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} (1 + \lambda_j)^{-\frac{s}{2}} |\xi_i| |\xi_j| \|\psi_i\|_{L^\infty(\mu)} \|\psi_j\|_{L^\infty(\mu)} \right] \\ &\lesssim \sum_{i=k_N+1}^{\infty} \sum_{j=k_N+1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} (1 + \lambda_j)^{-\frac{s}{2}} \lambda_i^{\frac{m-1}{4}} \lambda_j^{\frac{m-1}{4}} \\ &= \left[ \sum_{i=k_N+1}^{\infty} (1 + \lambda_i)^{-\frac{s}{2}} \lambda_i^{\frac{m-1}{4}} \right]^2 \lesssim \left[ \int_{k_N}^{\infty} x^{-\frac{2s+m-1}{2m}} dx \right]^2 \lesssim k_N^{-\frac{2s+3m-1}{m}}.\end{aligned}\tag{4.39}$$

Similarly,

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - \widetilde{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \left[ \sum_{i=1}^{k_N} \left| \left[ 1 + \lambda_i^{(N)} \right]^{-\frac{s}{2}} - \left[ 1 + \lambda_i \right]^{-\frac{s}{2}} \right| \|\psi_i^{(N)} \circ T_N\|_{L^\infty(\mu)} \right]^2.$$

By Lipschitz continuity of  $x^{-s/2}$  on  $[1, \infty)$  and (4.32), for  $i = 1, \dots, k_N$ ,

$$\begin{aligned} \left| \left[1 + \lambda_i^{(N)}\right]^{-\frac{s}{2}} - \left[1 + \lambda_i\right]^{-\frac{s}{2}} \right| &\leq \left[ (1 + \lambda_i^{(N)}) \wedge (1 + \lambda_i) \right]^{-\frac{s}{2}-1} \left| \lambda_i^{(N)} - \lambda_i \right| \\ &\lesssim \left[ (1 + \lambda_i^{(N)}) \wedge (1 + \lambda_i) \right]^{-\frac{s}{2}-1} \lambda_i \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right) \\ &\lesssim \lambda_i^{-\frac{s}{2}} \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right). \end{aligned}$$

By (4.33), if we choose  $k_N$  so that

$$k_N^3 \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_{k_N}} \right) \ll 1, \quad (4.40)$$

then we have, for  $i = 1, \dots, k_N$ ,

$$\|\psi_i^{(N)} \circ T_N\|_{L^\infty(\mu)} \leq \|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^\infty(\mu)} + \|\psi_i\|_{L^\infty(\mu)} \lesssim \lambda_i^{\frac{m-1}{4}}.$$

Therefore

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - \widetilde{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \left[ \sum_{i=1}^{k_N} \lambda_i^{-\frac{s}{2}} \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i} \right) \lambda_i^{\frac{m-1}{4}} \right]^2 \lesssim \left( \frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_{k_N}} \right)^2, \quad (4.41)$$

where we have used that  $\lambda_i^{\frac{-2s+m-1}{4}}$  is summable for  $s > \frac{3}{2}m - \frac{1}{2}$ . Lastly, by Lemma 4.4.14,

$$\begin{aligned} \mathbb{E}_\xi \|\widetilde{W}_n^{\mathcal{M}} - \widehat{W}_n^{\mathcal{M}}\|_{L^\infty(\mu)}^2 &\lesssim \left[ \sum_{i=1}^{k_N} (1 + \lambda_i)^{-\frac{s}{2}} \|\psi_i^{(N)} \circ T_N - \psi_i\|_{L^\infty(\mu)} \right]^2 \\ &\lesssim \left[ \sum_{i=1}^{k_N} (1 + \lambda_i)^{-\frac{s}{2}} \lambda_i^{m+1} i^{\frac{3}{2}} \sqrt{\frac{\rho_N}{\zeta_N} + \zeta_N \sqrt{\lambda_i}} \right]^2 \\ &\lesssim \left[ 1 \vee k_N^{\frac{-2s+9m+5}{m}} \right] \left( \frac{\rho_N}{\zeta_N} + \zeta_N \right). \end{aligned} \quad (4.42)$$

Combining (4.39), (4.41), (4.42), we have

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim k_N^{\frac{-2s+3m-1}{m}} + \left[1 \vee k_N^{\frac{-2s+9m+5}{m}}\right] \left(\frac{\rho_N}{\zeta_N} + \zeta_N\right). \quad (4.43)$$

Now we will set  $\zeta_N$  and  $k_N$  based on the dimension  $m$  and the smoothness parameter  $s$ .

**Case 1:**  $s \leq \frac{9}{2}m + \frac{5}{2}$ . Let  $\delta > 0$  be arbitrary and set

$$\begin{cases} \zeta_N \asymp N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{(m+4+\delta)(6+\frac{6}{m})}} (\log N)^{-\frac{mpm}{12m+12}}, & m \leq 4, \\ \zeta_N \asymp N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{12m+12}} (\log N)^{-\frac{mpm}{12m+12}}, & m \geq 5. \end{cases}$$

One can check that the conditions in (4.35) and (4.40) are satisfied and we get

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \begin{cases} N^{-\frac{2s-3m+1}{(m+4+\delta)(6m+6)}} (\log N)^{\frac{pm(2s-3m+1)}{12m+12}}, & m \leq 4, \\ N^{-\frac{2s-3m+1}{m(12m+12)}} (\log N)^{\frac{pm(2s-3m+1)}{12m+12}}, & m \geq 5. \end{cases}$$

**Case 2:**  $s > \frac{9}{2}m + \frac{5}{2}$ . Now (4.43) simplifies to

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim k_N^{\frac{-2s+3m-1}{m}} + \left(\frac{\rho_N}{\zeta_N} + \zeta_N\right).$$

Therefore by setting

$$\begin{cases} \zeta_N \asymp N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{m}{(m+4+\delta)(2s-3m+1)}} (\log N)^{-\frac{mpm}{4s-6m+2}}, & m \leq 4, \\ \zeta_N \asymp N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & k_N \asymp N^{\frac{1}{4s-6m+2}} (\log N)^{-\frac{mpm}{4s-6m+2}}, & m \geq 5, \end{cases}$$

we have

$$\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \lesssim \begin{cases} N^{-\frac{1}{m+4+\delta}} (\log N)^{\frac{pm}{2}}, & m \leq 4, \\ N^{-\frac{1}{2m}} (\log N)^{\frac{pm}{2}}, & m \geq 5. \end{cases}$$

Again one can check that the conditions in (4.35) and (4.40) are satisfied. This finishes the proof. □

#### 4.4.4 Putting Everything Together

Now we are ready to prove Theorem 4.2.1.

*Proof of Theorem 4.2.1.* By Section 4.4.1, it suffices to show that

$$\mathbb{E}_X \mathbb{E}_{f_0} \Pi_n^{\mathcal{M}}(f \in \mathcal{I}(L^\infty(\mu_{N_n})) : \|f - f_0\|_n \geq M_n \varepsilon_n \mid \mathcal{X}_{N_n}, \mathcal{Y}_n) \xrightarrow{n \rightarrow \infty} 0,$$

with the said  $\varepsilon_n$ . Let  $A_n$  be the high probability event in Corollary 4.4.7. Denote  $F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n) := \Pi_n^{\mathcal{M}}(f \in \mathcal{I}(L^\infty(\mu_{N_n})) : \|f - f_0\|_n \geq M_n \varepsilon_n \mid \mathcal{Y}_n)$ . We have

$$\mathbb{E}_X \mathbb{E}_{f_0} F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n) = \mathbb{E}_X [\mathbb{E}_{f_0} F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n)] \mathbf{1}_{A_n} + \mathbb{E}_X [\mathbb{E}_{f_0} F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n)] \mathbf{1}_{A_n^c}.$$

Since  $F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n) \leq 1$ , the second term is upper bounded by  $\mathbb{P}_X(A_n^c) \rightarrow 0$ . It then suffices to show  $\mathbb{E}_{f_0} F_n(\mathcal{X}_{N_n}, \mathcal{Y}_n) \rightarrow 0$  in the event of  $A_n$ . By Corollary 4.4.7 we can have  $10\mathbb{E}_\xi \|W_n^{\mathcal{M}} - W^{\mathcal{M}}\|_{L^\infty(\mu)}^2 \leq n^{-1}$  if we set a large enough proportion constant for  $N_n$ . This together with Theorem 4.4.4 and [van der Vaart and van Zanten, 2008a, Theorem 2.2] implies that there exists sets  $B_n \subset \mathcal{I}(L^\infty(\mu_{N_n}))$  (by taking the intersection of the sets  $B_n$  in Theorem 4.4.4 and  $\mathcal{I}(L^\infty(\mu_{N_n}))$ ) so that, for the same  $\varepsilon_n$  in the theorem statement,

$$\begin{aligned} \log N(6\varepsilon_n, B_n, \|\cdot\|_{L^\infty(\mu)}) &\leq 24Cn\varepsilon_n^2, \\ \pi_n^{\mathcal{M}}(B_n^c) &\leq e^{-4Cn\varepsilon_n^2}, \\ \pi_n^{\mathcal{M}}(\|w - \Phi^{-1}(f_0)\|_{L^\infty(\mu)} < 4\varepsilon_n) &\geq e^{-4n\varepsilon_n^2}, \end{aligned}$$

where  $\pi_n^{\mathcal{M}} = \mathcal{L}(W_n^{\mathcal{M}})$ .



**Case 1: Regression.** In the regression case, since  $\Phi$  is the identity, we have  $\Pi_n^{\mathcal{M}} = \pi_n^{\mathcal{M}}$  and the above three conditions are true with  $\pi_n^{\mathcal{M}}$  replaced by  $\Pi_n^{\mathcal{M}}$ . Furthermore, since  $\|\cdot\|_n$  is upper bounded by  $\|\cdot\|_{L^\infty(\mu)}$ , the above conditions remain true for the empirical norm. This together with the general results in [Ghosal and Van der Vaart, 2017, Section 8.3.2] proves the assertion.

**Case 2: Classification.** In the classification setting, by [Ghosal and Van der Vaart, 2017, Lemma 2.8] the average Kullback-Leibler divergence and variance in (4.17) between the densities after applying the link function  $\Phi$  are upper bounded by a multiple of the empirical norm. In particular if we set  $\mathcal{B}_n = \Phi(B_n)$ , then the conditions in (4.14), (4.15) and (4.16) hold, for a possibly different proportion constant for  $\varepsilon_n$ . Moreover, as discussed before, the root average square Hellinger distance  $d_n$  in (4.18) is equivalent to the empirical norm. Hence the result follows by [Ghosal and Van Der Vaart, 2007, Theorem 4].  $\square$

## 4.5 Discussion

In this chapter we have analyzed graph-based SSL using recent results on spectral convergence of graph Laplacians and standard Bayesian nonparametrics techniques. We show that, for a suitable choice of prior constructed with sufficiently many unlabeled data, the posterior contracts around the truth at a rate that is minimax optimal up to logarithmic factor. Our theory applies to both regression and classification.

We have assumed throughout that the  $X_i$ 's are uniformly distributed on  $\mathcal{M}$ . Our results can be generalized to nonuniform positive density  $q$  with respect to the volume form. In such a case, the continuum field is the Gaussian measure  $\mathcal{N}(0, (I - \Delta_q))^{-s}$  where  $-\Delta_q := -\frac{1}{q}\text{div}(q^2\nabla)$  is a weighted Laplacian-Beltrami operator, for which spectral convergence results can be found in García Trillos et al. [2019a]. Since we do not have an explicit dependence of the proportion constant on the index  $i$  as in Proposition 4.4.9, we have chosen to present our result in the uniform case. However, similar conclusions should be expected to hold in

the general case.

The Bayesian methodology we analyzed is inspired by popular existing graph-based optimization methods for SSL [Belkin et al., 2004, Zhu, 2005]. In order to achieve optimal contraction, the prior smoothness parameter needs to match the regularity of the truth, which is rarely available in applications. An important research direction stemming from our work is the development and analysis of adaptive Bayesian SSL methodologies that can achieve optimal contraction without a priori smoothness information. We expect that the existing results on adaptive estimation on manifolds [Castillo et al., 2014] and on large graphs [Kirichenko and van Zanten, 2017] will be an important stepping stone in this direction.

# CHAPTER 5

## FINITE ELEMENT REPRESENTATIONS OF GAUSSIAN PROCESSES: BALANCING NUMERICAL AND STATISTICAL ACCURACY

### 5.1 Introduction

Gaussian processes (GPs) are an important model for prior distributions over functions, and play a central role in spatial statistics, machine learning, Bayesian inverse problems, and a variety of other scientific and engineering applications [Stein, 1999a, Williams and Rasmussen, 2006, Stuart, 2010, Sullivan, 2015, Owhadi and Scovel, 2019]. However, GP methodology often suffers from the *big N problem*: conditioning a GP to  $N$  observations requires to factorize an  $N \times N$  covariance matrix, with a general cost of  $O(N^3)$ . Numerous approaches to address this challenge have been developed [Heaton et al., 2019]. The aim of this chapter is to provide novel understanding of the popular stochastic partial differential equation (SPDE) approach [Lindgren et al., 2011, 2022] for GP regression and classification with large datasets.

Let  $u(\mathbf{x})$  be a Matérn-type GP (see e.g. Bolin et al. [2020] or Section 5.2.1 below) on a bounded domain  $\mathcal{D} \subset \mathbb{R}^D$ . The SPDE approach approximates  $u$  with a GP  $u_h$  of the form

$$u_h(\mathbf{x}) = \sum_{i=1}^{n_h} w_i e_i(\mathbf{x}), \quad \mathbf{x} \in \mathcal{D}, \quad (5.1)$$

where  $e_i : \mathcal{D} \rightarrow \mathbb{R}$  are finite element (FE) basis functions and  $\mathbf{w} := (w_1, \dots, w_{n_h})^\top \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1})$  with sparse precision matrix  $\mathbf{Q} \in \mathbb{R}^{n_h \times n_h}$ . The dimension  $n_h$  of the basis  $\{e_i\}_{i=1}^{n_h}$  is determined by a mesh-size parameter  $h > 0$ . Previous work sets  $h$  so that  $n_h \approx N$ , and the  $O(N^3)$  computational cost is reduced by exploiting the local support of the FE basis functions  $e_i$  and the sparsity of  $\mathbf{Q}$ , see e.g. Lindgren et al. [2011], Bolin and Lindgren [2013], Bolin

and Kirchner [2020] or Section 5.2.3. However, the choice  $n_h \approx N$  has not been theoretically or empirically investigated. In particular, it is not clear if the computational gain achieved with  $n_h \approx N$  comes at the price of larger estimation error. In this chapter we shall introduce a framework for selecting  $n_h$  based on the posterior estimation performance achieved when using GP prior  $u_h$ . Our theory implies that, under certain smoothness assumptions, choosing  $n_h \ll N$  can indeed be sufficient in the large  $N$  asymptotics, as otherwise the statistical errors inherent to the regression or classification tasks dominate the numerical error in the approximation  $u_h \approx u$ . Therefore, in addition to the computational gain facilitated by sparsity, there is a second computational and memory gain: the *dimension* of the matrices that need to be factorized can be reduced in large  $N$  regimes without hindering the estimation accuracy. Numerical experiments will illustrate the applicability of our theory and the effect of the prior lengthscale in the pre-asymptotic regime.

The SPDE approach is part of a trend in GP methodology that seeks to leverage sparsity for computational efficiency [Quinonero-Candela and Rasmussen, 2005]. In this spirit, one can construct sparse approximations of the covariance matrix of the observations (a procedure known as *tapering* or *localization* [Gaspari and Cohn, 1999, Furrer et al., 2006]), or of the precision matrix [Datta et al., 2016] and its Cholesky factor [Schäfer et al., 2021, Kang and Katzfuss, 2021]. Other approaches exploiting sparsity include Vecchia approximations [Vecchia, 1988, Katzfuss and Guinness, 2021] and methods based on the screening effect [Stein, 2002]. These techniques are well established in several applications and are essential, for example, in the practical implementation of data assimilation algorithms for numerical weather forecasting [Houtekamer and Mitchell, 2001]. A complementary line of work relies on *smoothness* rather than sparsity for computational expediency. For instance, truncated Karhunen-Loève expansions in Bayesian inverse problems rely on a representation of the form (5.1) with small dimension  $n$ , spectral basis functions, and stochastic weights with diagonal covariance [Stuart, 2010]. These low-rank representations [Solin and Särkkä,

2020, Greengard and O’Neil, 2021] have been claimed to remove fine-scale variations of the process [Bolin and Kirchner, 2020], but can be accurate if the underlying process is smooth. Our work blends sparsity and smoothness demonstrating that, for regression and classification with large data-sets, sparse methods can benefit from a significant dimension reduction under mild smoothness assumptions.

To propose a criterion for choosing  $n_h$  with respect to  $N$ , we will exploit the concept of posterior contraction rates [Ghosal et al., 2000], which is discussed in Section 5.2.4. Roughly speaking, we consider a scaling sufficient if the posterior constructed with GP prior  $u_h$  contracts at the same rate as the posterior constructed with the true GP prior  $u$ . The Bayesian nonparametrics framework in van der Vaart and van Zanten [2008a] guarantees that if the rate of convergence of the GP prior approximation  $u_h \approx u$  is fast enough, then the corresponding posteriors contract at the same rate. Establishing convergence rates for approximations  $u_h \approx u$  is an active research area on numerical analysis of FE solution of fractional SPDEs [Bolin et al., 2020, Bolin and Kirchner, 2020, Cox and Kirchner, 2020]. As part of our analysis, we derive a crude estimate of the approximation error  $\mathbb{E}\|u_h - u\|_\infty^2$  for a particular FE discretization when  $\mathcal{D}$  is a hyperrectangle. The result holds for general dimension  $D$  while being less sharp than the one-dimensional result in Cox and Kirchner [2020], which also allows for more general domains. However, our main objective is to illustrate that the plug-in character of the framework [van der Vaart and van Zanten, 2008a] allows to seamlessly translate  $L^\infty$  and  $L^2$  error bounds for the approximation  $u_h \approx u$  into sufficient choices of  $n_h$  in terms of  $N$  in regression and classification settings. As we shall see, even crude error bounds suggest that, in the large  $N$  asymptotics,  $n_h \ll N$  can be sufficient under mild smoothness assumptions.

Numerical simulations in the regression setting will complement our theoretical analysis. Our experiments illustrate that (i) the qualitative theoretical behavior suggested by our large  $N$  asymptotic analysis is in agreement with the behavior observed with moderate sample-

size; (ii) if the truth is not smooth and has a short lengthscale, choosing  $n_h \gg N$  may indeed be necessary for the SPDE approach to match the estimation accuracy of the ground truth prior model; and (iii) outside the large  $N$  asymptotic regime, the prior lengthscale plays an important role in determining appropriate choice of  $n_h$  in terms of  $N$ . This last point is also partly explained by our theory, where the lengthscale appears as a prefactor in the error bound for the FE prior representations. We believe these findings together with our theoretical results provide useful insights for calibrating the FE approach in practice.

The study of fixed-domain, large  $N$  asymptotics [Stein, 1999a, 1990b,a, 1999b, Du et al., 2009, Wang and Loh, 2011] is motivated by applications in environmental science, ecology, climate, and hydrology, where  $N$  is often in the order of hundreds of thousands or larger. At a high level, our criterion resembles the in-fill asymptotic analysis of tapered covariance functions in Furrer et al. [2006], where the authors give conditions on the taper function that guarantee large-data asymptotic equivalence of the mean-squared prediction error of the true and tapered covariance models. As in Furrer et al. [2006], we may interpret  $u_h$  as defining a misspecified covariance model and then, similar to Furrer et al. [2006], Stein [1993], our criterion guarantees that the misspecification is inconsequential in a large-data regime. On the other hand, even if the Matérn-type GP  $u$  is not interpreted as a ground truth prior model, our analysis suggests that over-discretizing the FE representations  $u_h$  should be avoided, as there is a threshold beyond which further discretizing increases the computational cost without improving the estimation accuracy. Similar ideas permeate the study of the value of unlabeled data in semi-supervised learning [Sanz-Alonso and Yang, 2022b] with graph representations of Matérn GPs [Sanz-Alonso and Yang, 2022a].

The rest of this chapter is organized as follows. We provide all necessary background and formalize our problem setting in Section 5.2. Our main results are in Section 5.3 and complementary numerical experiments in Section 5.4. We close in Section 5.5 with possible extensions of our main results and open directions that stem from our work. All the proofs

are deferred to Section 5.6.

## 5.2 Background and Problem Setting

To make our presentation self-contained, we introduce in this section all necessary background and formalize our problem setting. Matérn-type GPs and their connection with the classical Matérn covariance function are discussed in Section 5.2.1. Section 5.2.2 reviews FE representations of Matérn-type GPs. Our regression and classification problem settings are formalized in Section 5.2.3, where we also summarize how FE representations of Matérn-type GPs allow to speed up computations. Finally, Section 5.2.4 overviews the Bayesian nonparametrics framework that we employ as our criterion to identify sufficient scalings of  $n_h$  with respect to  $N$ .

### 5.2.1 The Matérn Covariance Function and SPDE Representations

Recall that the Matérn covariance function is defined by

$$c_{\text{Mat}}(\mathbf{x}, \mathbf{x}') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} (\kappa |\mathbf{x} - \mathbf{x}'|)^\nu K_\nu(\kappa |\mathbf{x} - \mathbf{x}'|), \quad \mathbf{x}, \mathbf{x}' \in \mathbb{R}^D, \quad (5.2)$$

where  $|\cdot|$  is the Euclidean distance on  $\mathbb{R}^D$ ,  $\Gamma$  is the gamma function and  $K_\nu$  is the modified Bessel function of the second kind. The parameters  $\sigma$ ,  $\nu$ ,  $\kappa$  control, respectively, the marginal variance, smoothness of the sample paths, and correlation lengthscale. Due to its flexibility, the Matérn model is widely used in spatial statistics [Stein, 1999a, Gelfand et al., 2010], machine learning [Williams and Rasmussen, 2006], and uncertainty quantification [Sullivan, 2015], with applications in various scientific fields [Guttorp and Gneiting, 2006, Cameletti et al., 2013]. The connection between the Matérn covariance and SPDEs has long been

noticed [Whittle, 1954]. Consider formally the equation

$$(\kappa^2 - \Delta)^{s/2}u = \kappa^{s-D/2}\mathcal{W} \quad \text{in } \mathcal{D}, \quad (5.3)$$

where  $s = \nu + D/2$ ,  $\Delta$  is a Laplacian and  $\mathcal{W}$  is a spatial white noise. (Here and below we will ignore the marginal variance which acts only as a scaling factor.) If  $\mathcal{D} := \mathbb{R}^D$ , then the unique stationary solution to (5.3), suitably interpreted [Whittle, 1954], has covariance function (5.2).

Following Lindgren et al. [2011], we will define *Matérn-type GPs* by solution of (5.3) in a *bounded* domain  $\mathcal{D} \subset \mathbb{R}^D$ , interpreting the SPDE (5.3) as in Bolin et al. [2020]. We outline here the main ideas and refer to Bolin et al. [2020] for further details. Let  $\mathcal{L} := \kappa^2 - \Delta$  be equipped with homogeneous Dirichlet or Neumann boundary condition. The eigenfunctions  $\{\psi_i\}_{i=1}^\infty$  of the Dirichlet (or Neumann) Laplacian form an orthonormal basis of  $L^2(\mathcal{D})$ , where the associated ordered eigenvalues  $\{\lambda_i\}_{i=1}^\infty$  satisfy  $\lambda_i \asymp i^{2/D}$  by Weyl's law (see e.g. [Davies, 1996, Theorem 6.3.1]). The fractional power operator  $\mathcal{L}^{s/2}$  in (5.3) is then defined by

$$\mathcal{L}^{s/2}u := \sum_{i=1}^{\infty} (\kappa^2 + \lambda_i)^{s/2} \langle u, \psi_i \rangle \psi_i$$

with domain  $\{u \in L^2(\mathcal{D}) : \sum_{i=1}^{\infty} (\kappa^2 + \lambda_i)^s \langle u, \psi_i \rangle^2 < \infty\}$ , where  $\langle \cdot, \cdot \rangle$  denotes the  $L^2(\mathcal{D})$ -inner product. The white noise in (5.3) is formally defined by the series  $\mathcal{W} = \sum_{i=1}^{\infty} \xi_i \psi_i$ , with  $\xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$  set on a complete probability space  $(\Omega, \mathcal{A}, \mathbb{P})$ . As rigorously shown in [Bolin et al., 2020, Lemma 2.1], existence and uniqueness of solutions to (5.3) in  $L^2(\Omega; L^2(\mathcal{D}))$  is guaranteed for  $s > D/2$ . Moreover, the solution can be represented as a series expansion

$$u(\mathbf{x}) = \kappa^{s-D/2} \sum_{i=1}^{\infty} (\kappa^2 + \lambda_i)^{-s/2} \xi_i \psi_i(\mathbf{x}), \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad \mathbf{x} \in \mathcal{D}, \quad (5.4)$$

where the assumption  $s > D/2$  together with Weyl's law guarantees that  $u \in L^2(\mathcal{D})$  almost



surely. We refer to  $u$  defined by (5.4) as a Matérn-type GP. The covariance function of Matérn-type GPs no longer agrees with the classical Matérn covariance model (5.2), but approximates it well away from the boundary —see for instance Proposition 5.3.1 below.

### 5.2.2 Finite Element Representations of Matérn-type Gaussian Processes

Let  $\mathcal{D} \subset \mathbb{R}^D$  be a bounded domain and let  $\{V_h\}_{h \in (0,1)}$  be a family of subspaces of  $H^1(\mathcal{D})$  (the space of functions whose weak derivatives belong to  $L^2(\mathcal{D})$ ) with finite dimensions  $n_h := \dim(V_h) < \infty$ . In subsequent developments  $h$  will play the role of a mesh-size parameter and  $n_h \asymp h^{-D}$ . Consider the Galerkin discretization  $-\Delta_h : V_h \rightarrow V_h$  of  $-\Delta$  defined as

$$\langle -\Delta_h u_h, v_h \rangle = \langle -\Delta u_h, v_h \rangle \quad \forall u_h, v_h \in V_h.$$

Let  $\{(\lambda_{h,i}, \psi_{h,i})\}_{i=1}^{n_h}$  be the eigenpairs of  $-\Delta_h$  satisfying

$$\langle -\Delta_h \psi_{h,i}, v_h \rangle = \lambda_{h,i} \langle \psi_{h,i}, v_h \rangle \quad \forall v_h \in V_h,$$

where we assume the  $\lambda_{h,i}$ 's are in increasing order and the  $\psi_{h,i}$ 's are orthonormal. We then define a discretization of the SPDE (5.3) by

$$\mathcal{L}_h^{s/2} u_h := (\kappa^2 - \Delta_h)^{s/2} u_h = \kappa^{s-D/2} \mathcal{W}_h, \quad \mathcal{W}_h := \sum_{i=1}^{n_h} \xi_i \psi_{h,i}, \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1). \quad (5.5)$$

We refer to the solution  $u_h$  as a FE representation of the Matérn-type GP  $u$ . Note that

$$u_h(\mathbf{x}) = \kappa^{s-D/2} \sum_{i=1}^{n_h} (\kappa^2 + \lambda_{h,i})^{-s/2} \xi_i \psi_{h,i}(\mathbf{x}), \quad \xi_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad \mathbf{x} \in \mathcal{D}. \quad (5.6)$$

Inspection of (5.4) and (5.6) suggests that the error in the approximation  $u_h \approx u$  is largely determined by the FE error in the approximations  $\lambda_{h,i} \approx \lambda_i$  and  $\psi_{h,i} \approx \psi_i$ . We will pursue this idea in our error analysis in Section 5.3. However, the Karhunen-Loève representation (5.6) is not in general useful for practical implementation, as the eigenpairs  $\{(\lambda_{h,i}, \psi_{h,i})\}_{i=1}^{n_h}$  can be expensive to compute and the eigenfunctions do not have compact support. The following result from Lindgren et al. [2011] shows that the solution to (5.5) admits an equivalent representation in terms of a FE basis, as foreshadowed in (5.1).

**Proposition 5.2.1.** *Let  $\{e_{h,i}\}_{i=1}^{n_h}$  be a FE basis of  $V_h$ , and denote by  $\mathbf{M}$  and  $\mathbf{G}$  the mass and stiffness matrices with entries  $\mathbf{M}_{ij} = \langle e_{h,i}, e_{h,j} \rangle$  and  $\mathbf{G}_{ij} = \langle \nabla e_{h,i}, \nabla e_{h,j} \rangle$ . For  $0 \neq s \in \mathbb{N}$ , the FE representation  $u_h$  of the Matérn-type GP  $u$  admits the characterization*

$$u_h(\mathbf{x}) = \sum_{i=1}^{n_h} w_i e_{h,i}(\mathbf{x}), \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1}), \quad (5.7)$$

where  $\mathbf{Q} = (\kappa^2 \mathbf{M} + \mathbf{G}) [\mathbf{M}^{-1} (\kappa^2 \mathbf{M} + \mathbf{G})]^{s-1}$ .

Notice that (5.7) does not involve the eigenpairs. Moreover, the matrices  $\mathbf{M}$  and  $\mathbf{G}$  are sparse for standard FE basis  $e_{h,i}$ , e.g. tent functions. Lumping the mass matrix  $\mathbf{M}$  ensures sparsity of  $\mathbf{Q}$  and gives a Gauss-Markov approximation to the Matérn-type GP  $u$  [Lindgren et al., 2011]. For  $s \notin \mathbb{N}$ , the rational SPDE approach can be adopted [Bolin and Kirchner, 2020].

### 5.2.3 Gaussian Process Regression and Classification: Finite Element

#### Representations

Here we introduce the regression and classification models we consider, and describe briefly how FE representations of Matérn-type GPs can alleviate the computational burden of these tasks. Given  $N$  pairs of data  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^N$  we are interested in inferring  $f_0(\mathbf{x}) = \mathbb{E}[Y|\mathbf{X} = \mathbf{x}]$  under the following data-generating mechanisms:

- Fixed design regression:  $Y_i = f_0(\mathbf{X}_i) + \eta_i$ , where the  $\mathbf{X}_i$ 's are fixed (and distinct) covariates and  $\eta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \tau^2)$  with  $\tau$  known.
- Binary classification:  $\mathbb{P}(Y_i = 1 | \mathbf{X}_i) = f_0(\mathbf{X}_i)$ , where  $\mathbf{X}_i \stackrel{i.i.d.}{\sim} \mu$  for some distribution  $\mu$  over  $\mathcal{D}$ .

For simplicity we shall assume for the rest of this chapter that  $\mu$  is the uniform distribution over  $\mathcal{D}$ , but we note that it suffices to assume that  $\mu$  admits a Lebesgue density bounded above and below by positive constants.

For fixed design regression, we set a FE Matérn-type GP prior  $u_h$  on  $f_0$ . The posterior of the weights  $\mathbf{w}$  is given by

$$\mathbf{w} | \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \sim \mathcal{N}((\mathbf{S}^\top \mathbf{S} + \tau^2 \mathbf{Q})^{-1} \mathbf{S}^\top \mathbf{y}, (\tau^{-2} \mathbf{S}^\top \mathbf{S} + \mathbf{Q})^{-1}),$$

where  $\mathbf{S} \in \mathbb{R}^{N \times n_h}$  has entries  $\mathbf{S}_{ij} = e_j(\mathbf{X}_i)$  and  $\mathbf{y} = (Y_1, \dots, Y_N)^\top$ . The main computational cost for posterior inference is in factorizing the  $n_h \times n_h$  matrix  $\mathbf{S}^\top \mathbf{S} + \tau^2 \mathbf{Q}$ . This factorization can be efficiently computed since the local support of standard FE basis functions ensures sparsity of  $\mathbf{S}$ , and  $\mathbf{Q}$  can be made sparse as discussed in Section 5.2.2.

For binary classification, let  $\Phi$  be the logistic function and consider a wrapped GP prior  $\Phi \circ u_h$  over  $f_0$ . The posterior log-density is given by

$$\begin{aligned} \log \mathbb{P}(\mathbf{w} | \{(\mathbf{X}_i, Y_i)\}_{i=1}^N) &= \sum_{i=1}^N Y_i \log \Phi((\mathbf{S}\mathbf{w})_i) + (1 - Y_i) \log(1 - \Phi((\mathbf{S}\mathbf{w})_i)) \\ &\quad - \frac{1}{2} \mathbf{w}^\top \mathbf{Q} \mathbf{w} + \text{const}, \end{aligned} \tag{5.8}$$

where  $(\mathbf{S}\mathbf{w})_i$  denotes the  $i$ -th entry of  $\mathbf{S}\mathbf{w}$ . Two standard procedures for posterior inference are *maximum a posteriori* (MAP) estimation and Markov chain Monte Carlo (MCMC) sampling. To compute the MAP estimate, (5.8) is optimized to recover the weights with highest posterior density. This optimization problem can be efficiently solved using the

Hessian of the objective function, which takes the form  $\mathbf{S}^\top \mathbf{D} \mathbf{S} - \mathbf{Q}$ , where  $\mathbf{D}$  is a diagonal matrix with

$$\mathbf{D}_{ii} = \Phi''((\mathbf{S}\mathbf{w})_i) \left[ \frac{Y_i}{\Phi((\mathbf{S}\mathbf{w})_i)} - \frac{1 - Y_i}{1 - \Phi((\mathbf{S}\mathbf{w})_i)} \right] - [\Phi'((\mathbf{S}\mathbf{w})_i)]^2 \left[ \frac{Y_i}{\Phi((\mathbf{S}\mathbf{w})_i)^2} + \frac{1 - Y_i}{(1 - \Phi((\mathbf{S}\mathbf{w})_i))^2} \right].$$

Therefore, the computational cost is largely determined by the sparsity of the matrix  $\mathbf{S}^\top \mathbf{D} \mathbf{S} - \mathbf{Q}$ , which in turn depends on the sparsity of  $\mathbf{S}$  and  $\mathbf{Q}$ . On the other hand, MCMC algorithms for posterior inference with GP priors have been widely studied [Neal, 1998, Cotter et al., 2010b, Cui et al., 2016, García Trillos et al., 2020b, Sanz-Alonso et al., 2019], and a key idea behind these methods is to employ a proposal mechanism  $\mathbf{w} \mapsto \mathbf{w}'$  of the form

$$\mathbf{w}' = \theta \mathbf{w} + (1 - \theta)^{1/2} \boldsymbol{\gamma}, \quad \boldsymbol{\gamma} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1}), \quad (5.9)$$

which leaves the prior distribution  $\mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1})$  of the weights invariant. In order to sample  $\boldsymbol{\gamma} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1})$  with large  $n_h$  it is important to leverage sparsity of  $\mathbf{Q}$  [Rue and Held, 2005].

#### 5.2.4 Our Criterion: Matching Posterior Contraction Rates

The FE approach outlined above involves a user-chosen hyperparameter  $h$  that affects both the estimation performance and computational cost. Smaller  $h$  leads to better approximation of the Matérn-type GP  $u$  by  $u_h$  and possibly enhanced inference, but renders a larger  $n_h$  that increases the computational cost. Since  $u_h$  is supposed to approximate the Matérn-type GP  $u$ , a natural choice for  $h$  is so that the estimation performance of using  $u_h$  as the prior is “comparable” to that of  $u$ . In this section we shall formalize such intuition with the notion of posterior contraction rates.

To begin with, recall that the goal is to infer the conditional expectation  $f_0(\mathbf{x}) = \mathbb{E}[Y | \mathbf{X} = \mathbf{x}]$  from data  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^N$ . We shall adopt a frequentist Bayesian perspective by putting a sequence of priors  $\Pi_N$  over  $f_0$  and assuming that the data are indeed generated from a

fixed  $f_0$  which we interpret as the ground truth. Following Ghosal et al. [2000], we say that the sequence of posteriors with respect to  $\Pi_N$  contracts around  $f_0$  with rate  $\varepsilon_N$  if, for any sufficiently large  $M > 0$ ,

$$\mathbb{E}_{f_0} \Pi_N \left( f : d_N(f, f_0) \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \xrightarrow{N \rightarrow \infty} 1. \quad (5.10)$$

Here the expectation is taken with respect to the data distribution of  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^N$  determined by  $f_0$  and the marginal of the  $\mathbf{X}_i$ 's, and  $d_N$  is a suitable discrepancy measure. Roughly speaking,  $\varepsilon_N$  is the rate at which one can shrink the radius of a ball centered around the truth while at the same time capturing almost all the posterior mass. The condition (5.10) implies that asymptotically the sequence of posteriors will be nearly supported on a ball of radius  $O(\varepsilon_N)$  around  $f_0$ . Therefore,  $\varepsilon_N$  can be loosely interpreted as the convergence rate of the posteriors towards the truth. An important consequence [Ghosal et al., 2000, Theorem 2.5] is that the point estimator defined as

$$\hat{f}_N = \arg \max_g \left[ \Pi_N \left( f : d_N(f, g) \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \right],$$

converges (in probability) to  $f_0$  with the same rate  $\varepsilon_N$ . Therefore the contraction rate serves as a natural criterion for quantifying the estimation performance of the posteriors.

Following Section 5.2.3, the sequence of priors is taken as  $\Pi_N = \text{Law}(u_{h_N})$  (resp.  $\text{Law}(\Phi(u_{h_N}))$ ) for fixed design regression (resp. binary classification). The selection criterion for  $h_N$  that we propose is to choose  $h_N$  so that the sequence of posteriors with respect to  $\Pi_N$  contracts at the same rate as if  $\Pi_N \equiv \Pi := \text{Law}(u)$  (resp.  $\text{Law}(\Phi(u))$ ), where  $u$  is the Matérn-type GP that  $u_{h_N}$  is approximating. It turns out that there is a simple condition on the approximation accuracy of  $u_{h_N}$  towards  $u$  that guarantees this matching of posterior contraction rates, which we make precise below.

We start by reviewing the key ingredients of the theory when a single prior is adopted, i.e.,

when  $\Pi_N \equiv \Pi$  in the above. Consider now  $u$  as a GP taking values in  $(L^\infty(\mathcal{D}), \|\cdot\|_\infty)$  (see e.g. Lemma 5.3.6 below for conditions under which this is valid) for fixed design regression and in  $(L^2(\mathcal{D}), \|\cdot\|_2)$  for binary classification. By [van der Vaart and van Zanten, 2008a, Theorems 3.2 and 3.3], the contraction rate with respect to  $\Pi$  in the fixed design regression (resp. binary classification) setting can be characterized as the sequence  $\varepsilon_N$  that satisfies  $\phi_{f_0}(\varepsilon_N; u, \|\cdot\|_\infty) \leq N\varepsilon_N^2$  (resp.  $\phi_{\Phi^{-1}(f_0)}(\varepsilon_N; u, \|\cdot\|_2) \leq N\varepsilon_N^2$ ), where

$$\phi_{\omega_0}(\varepsilon; u, \|\cdot\|_{\mathbb{B}}) := \inf_{g \in \mathbb{H}: \|g - \omega_0\|_{\mathbb{B}} < \varepsilon} \|g\|_{\mathbb{H}}^2 - \log \mathbb{P}(\|u\|_{\mathbb{B}} < \varepsilon), \quad (5.11)$$

and  $(\mathbb{H}, \|\cdot\|_{\mathbb{H}})$  denotes the *reproducing kernel Hilbert space* (RKHS) of  $\Pi$  (see e.g. van der Vaart and van Zanten [2008b] for more details). Under such circumstances, the sequence of posteriors with respect to  $\Pi$  contracts around  $f_0$  with rate  $\varepsilon_N$  in the sense of (5.10) with  $d_N = \|\cdot\|_N$  the empirical norm defined as  $\|f\|_N^2 = N^{-1} \sum_{i=1}^N |f(\mathbf{X}_i)|^2$  for fixed design regression and  $d_N = \|\cdot\|_2$  for binary classification. In other words, the posterior contraction rate can be determined by analyzing the so-called *concentration function* (5.11) of the prior. Now when a sequence of priors  $\Pi_N$  is used instead, it is reasonable to expect that if  $\Pi_N$  approximates  $\Pi$  sufficiently well, the concentration functions of  $\Pi_N$  will be close to that of  $\Pi$  so that the same contraction rate can be achieved. Indeed this is implied by [van der Vaart and van Zanten, 2008a, Theorems 2.2, 3.2 and 3.3], which we record as a proposition.

**Proposition 5.2.2.** *1. Fixed design regression: Let  $\Pi_N = \text{Law}(u_{h_N})$ . Suppose  $\varepsilon_N$  is a sequence of real numbers satisfying  $\phi_{f_0}(\varepsilon_N; u, \|\cdot\|_\infty) \leq N\varepsilon_N^2$  and*

$$10\mathbb{E}\|u_{h_N} - u\|_\infty^2 \leq N^{-1}. \quad (5.12)$$

*Then, for any sufficiently large  $M > 0$ ,*

$$\mathbb{E}_{f_0} \Pi_N \left( f : \|f - f_0\|_N \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \xrightarrow{N \rightarrow \infty} 1.$$

2. *Binary classification:* Let  $\Pi_N = \text{Law}(\Phi(u_{h_N}))$ . Suppose  $\varepsilon_N$  is a sequence of real numbers satisfying  $\phi_{\Phi^{-1}(f_0)}(\varepsilon_N; u, \|\cdot\|_2) \leq N\varepsilon_N^2$  and

$$10\mathbb{E}\|u_{h_N} - u\|_2^2 \leq N^{-1}. \quad (5.13)$$

Then, for any sufficiently large  $M > 0$ ,

$$\mathbb{E}_{f_0} \Pi_N \left( f : \|f - f_0\|_2 \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \xrightarrow{N \rightarrow \infty} 1.$$

Proposition 5.2.2 shows that the posteriors constructed with prior  $\Pi_N$  and with prior  $\Pi$  contract at the same rate, provided that the prior approximation is sufficiently accurate. Therefore it suffices to choose  $h_N$  so that (5.12) or (5.13) is satisfied, giving a simple criterion for setting  $h_N$ . In particular, if the error  $\mathbb{E}\|u_{h_N} - u\|_\infty^2$  or  $\mathbb{E}\|u_{h_N} - u\|_2^2$  decreases sufficiently fast, then a slowly decaying  $h_N$  is enough and leads to  $n_{h_N} \asymp h_N^{-D} \ll N$ . We will show in Section 5.3 for a simple linear FE method in a concrete setting that this is indeed the case under certain smoothness assumptions, and demonstrate such behavior through simulation studies in Section 5.4. Several possible extensions will be discussed in Section 5.5, building on the key idea of using Proposition 5.2.2 to balance the numerical error in the prior approximation with the statistical errors in regression and classification tasks.

### 5.3 Main Results

In this section we obtain sufficient scalings of  $n_h$  with respect to  $N$  using spectral error analysis for FE eigenvalue problems and our criterion outlined in Section 5.2.4. We assume throughout that  $\mathcal{D} = (0, L_1) \times \cdots \times (0, L_D)$  is a hyperrectangle and that the Laplacian in (5.3) is supplemented with Neumann boundary condition, so that we have the following

explicit expressions for its eigenvalues and eigenfunctions

$$\Lambda_{\mathbf{i}} = \sum_{d=1}^D \frac{i_d \pi}{L_d}, \quad \Psi_{\mathbf{i}}(\mathbf{x}) = C_{\mathbf{i}} \prod_{d=1}^D \cos\left(\frac{i_d \pi x_d}{L_d}\right), \quad (5.14)$$

where  $\mathbf{i} = (i_1, \dots, i_D) \in \mathbb{N}^D$  is a multi-index and  $C_{\mathbf{i}}$ 's are constants so that the  $\Psi_{\mathbf{i}}$ 's are  $L^2(\mathcal{D})$ -normalized. The Matérn-type GP (5.3) can then be written as

$$u = \kappa^{s-D/2} \sum_{\mathbf{i} \in \mathbb{N}^D} (\kappa^2 + \Lambda_{\mathbf{i}})^{-s/2} \xi_{\mathbf{i}} \Psi_{\mathbf{i}}, \quad \xi_{\mathbf{i}} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1). \quad (5.15)$$

The explicit expressions for the eigenpairs in (5.14) allow us to establish the following result [Khristenko et al., 2019, Theorem 2.1], which shows that the covariance function of (5.15) is nearly indistinguishable from the classical Matérn covariance function (5.2) away from the boundary.

**Proposition 5.3.1.** *Let  $c(\mathbf{x}, \mathbf{x}')$  denote the covariance function of the Matérn-type GP (5.15) and let  $c_{\text{Mat}}(\mathbf{x}, \mathbf{x}')$  be the Matérn covariance function (5.2) with*

$$\sigma^2 = \frac{\Gamma(s - \frac{D}{2})}{(4\pi)^{D/2} \Gamma(s)}. \quad (5.16)$$

Then

$$c(\mathbf{x}, \mathbf{x}') = \sum_{\mathbf{k} \in \mathbb{Z}^D} \sum_{\mathbf{T} \in \mathcal{T}} c_{\text{Mat}}(\mathbf{T}\mathbf{x}, \mathbf{x}' - 2\mathbf{k}\mathbf{L}), \quad \mathbf{x}, \mathbf{x}' \in (0, L_1) \times \dots \times (0, L_D), \quad (5.17)$$

where  $\mathcal{T}$  is the collection of all  $D \times D$  diagonal matrices whose diagonal entries are either 1 or  $-1$ ,  $\mathbf{T}\mathbf{x}$  denotes matrix-vector multiplication and  $\mathbf{k}\mathbf{L}$  denotes  $(k_1 L_1, \dots, k_D L_D)$ .

Note that if the correlation range  $\rho = \sqrt{8\nu}/\kappa$  (where, recall,  $\nu = s - D/2$ ) is much smaller than  $\min_d L_d$ , and in addition  $\mathbf{x}, \mathbf{x}'$  are at a distance larger than  $2\rho$  from each side of the hyperrectangle, the only significant term that remains in (5.17) is  $c_{\text{Mat}}(\mathbf{x}, \mathbf{x}')$ . Therefore,



(5.15) gives a good approximation of the classical Matérn model away from the boundary. In practice one can choose a larger hyperrectangle than the domain of interest to reduce the boundary effect [Lindgren et al., 2011], see also Khristenko et al. [2019]. Our focus on hyperrectangles also facilitates the concrete FE construction and error analysis in the next subsection.

### 5.3.1 FEM Construction and Spectral Error Bounds

We shall construct the FE space on  $[0, L_1] \times \cdots \times [0, L_D]$  as the tensor product of FE spaces on each interval  $[0, L_d]$ . To begin with, let  $P$  be a uniform partition of  $[0, L_d]$  into  $K + 1$  points with width  $h = L/K$  and let  $V_h$  be the space of continuous piecewise linear functions with respect to  $P$ . To simplify the notation we drop the dependence on  $d$  below. Precisely, a basis of  $V_h$  consists of

$$e_{h,k} = \begin{cases} h^{-1}x - k + 1 & x \in [(k-1)h, kh] \\ -h^{-1}x + k + 1 & x \in [kh, (k+1)h], \quad k = 1, \dots, K-1, \\ 0 & \text{otherwise} \end{cases}$$

with  $e_{h,0} = (-h^{-1}x + 1)\mathbf{1}_{[0,h]}$  and  $e_{K,h} = (h^{-1}x - K + 1)\mathbf{1}_{[(K-1)h, Kh]}$ . Let  $\mathcal{J}_h$  be the Galerkin discretization of  $\kappa^2 - \frac{d^2}{dx^2}$  over  $V_h$ . The eigenvalues  $\{\lambda_{h,i}\}_{i=0}^K$  and eigenfunctions  $\{\psi_{h,i}\}_{i=0}^K$  of  $\mathcal{J}_h$  can be found by solving the generalized eigenvalue problem

$$\mathbf{Gz} = \lambda \mathbf{Mz},$$

where  $\mathbf{z}$  represents the coordinates of  $\psi_{h,i}$  in terms of the  $e_{h,k}$ 's and  $\mathbf{G}, \mathbf{M} \in \mathbb{R}^{(K+1) \times (K+1)}$  are matrices with entries

$$\mathbf{G}_{ij} = \frac{1}{h} \cdot \begin{cases} 2 & i = j \notin \{1, K+1\} \\ 1 & i = j \in \{1, K+1\} \\ -1 & |i - j| = 1 \\ 0 & \text{otherwise} \end{cases}, \quad \mathbf{M}_{ij} = h \cdot \begin{cases} 2/3 & i = j \notin \{1, K+1\} \\ 1/3 & i = j \in \{1, K+1\} \\ 1/6 & |i - j| = 1 \\ 0 & \text{otherwise} \end{cases}.$$

One can check that

$$\lambda_{h,i} = \frac{6}{h^2} \frac{1 - \cos(i\pi h/L)}{2 + \cos(i\pi h/L)}, \quad \psi_{h,i} = c_i \sum_{k=0}^K \cos\left(\frac{ki\pi h}{L}\right) e_{h,k}, \quad i, k \in [K], \quad (5.18)$$

where  $c_i$ 's are normalizing constants so that  $\psi_{h,i}$  has  $L^2(\mathcal{D})$  norm one, and  $[K] = \{0, \dots, K\}$ .

We then have the following error estimates:

**Lemma 5.3.2.** *Let  $\{(\lambda_i, \psi_i)\}_{i=1}^\infty$  be the eigenvalues and  $L^2(\mathcal{D})$ -orthonormal eigenfunctions of  $\kappa^2 - \frac{d^2}{dx^2}$  over  $(0, L)$  with Neumann boundary condition. There is a constant  $C$  so that, for  $i \in [K]$ ,*

$$|\lambda_{h,i} - \lambda_i| \leq C\lambda_i^2 h^2, \quad \|\psi_{h,i} - \psi_i\|_\infty \leq C\lambda_i h^2.$$

Furthermore the  $\psi_{h,i}$ 's are also  $L^2(\mathcal{D})$ -orthonormal.

**Remark 5.3.3.** *Eigenvalue estimates and eigenfunction estimates in  $L^2$  norm can be found for instance in [Strang and Fix, 1973, Theorems 6.1 and 6.2], where more general elliptic operators and domains are considered. However, for our subsequent developments we need eigenfunction estimates in  $L^\infty$  norm, and for this reason we include an elementary proof of Lemma 5.3.2 in Section 5.6.  $\square$*

For Galerkin discretization of  $\kappa^2 - \Delta$  on  $[0, L_1] \times \cdots \times [0, L_D]$ , let  $\mathcal{P}$  be the uniform grid constructed by uniformly partitioning each interval with  $K_d + 1$  nodes so that  $h_d = L_d/K_d$  in each dimension. Define for  $\mathbf{h} = (h_1, \dots, h_D)$  the FE space

$$\mathcal{V}_{\mathbf{h}} = V_{h_1} \otimes \cdots \otimes V_{h_D} := \left\{ v(\mathbf{x}) = \prod_{d=1}^D v_{h_d}(x_d) : v_{h_d} \in V_{h_d} \right\},$$

where  $V_{h_d}$  is the FE space on  $[0, L_d]$  constructed above. It can be shown that the eigenvalues  $\Lambda_{\mathbf{h}, \mathbf{i}}$  and eigenfunctions  $\Psi_{\mathbf{h}, \mathbf{i}}$  of  $\mathcal{L}_{\mathbf{h}}$  (the Galerkin discretization of  $\kappa^2 - \Delta$ ) are

$$\Lambda_{\mathbf{h}, \mathbf{i}} = \sum_{d=1}^D \lambda_{h_d, i_d}, \quad \Psi_{\mathbf{h}, \mathbf{i}}(\mathbf{x}) = \prod_{d=1}^D \psi_{h_d, i_d}(x_d), \quad \mathbf{i} \in [K_1] \times \cdots \times [K_D],$$

where the  $\lambda_{h_d, i_d}$ 's and  $\psi_{h_d, i_d}$ 's are as in (5.18). Indeed for  $v_{\mathbf{h}}(\mathbf{x}) = \prod_{d=1}^D v_{h_d}(x_d) \in \mathcal{V}_{\mathbf{h}}$  we have that

$$\begin{aligned} \langle \nabla \Psi_{\mathbf{h}, \mathbf{i}}, \nabla v_{\mathbf{h}} \rangle &= \int_{\mathcal{D}} \sum_{d=1}^D \left( \psi'_{h_d, i_d} v'_{h_d} \prod_{\ell \neq d} \psi_{h_\ell, i_\ell} v_{h_\ell} \right) d\mathbf{x} \\ &= \sum_{d=1}^D \langle \psi'_{h_d, i_d}, v'_{h_d} \rangle \prod_{\ell \neq d} \langle \psi_{h_\ell, i_\ell}, v_{h_\ell} \rangle \\ &= \sum_{d=1}^D \lambda_{h_d, i_d} \prod_{\ell=1}^D \langle \psi_{h_\ell, i_\ell}, v_{h_\ell} \rangle = \sum_{d=1}^D \lambda_{h_d, i_d} \langle \Psi_{\mathbf{h}, \mathbf{i}}, v_{\mathbf{h}} \rangle, \end{aligned}$$

where the primes denote weak derivatives. Moreover the  $\Psi_{\mathbf{h}, \mathbf{i}}$ 's are orthonormal since the  $\psi_{h_d, i_d}$ 's are and hence they form a complete set of eigenbasis for  $\mathcal{L}_{\mathbf{h}}$ . The following error estimates are immediate, where we recall that the true eigenpairs are given in (5.14):

**Lemma 5.3.4.** *For  $\mathbf{i} \in [K_1] \times \cdots \times [K_D]$  we have*

$$|\Lambda_{\mathbf{h}, \mathbf{i}} - \Lambda_{\mathbf{i}}| \leq C \Lambda_{\mathbf{i}}^2 h^2, \quad \|\Psi_{\mathbf{h}, \mathbf{i}} - \Psi_{\mathbf{i}}\|_{\infty} \leq C \Lambda_{\mathbf{i}} h^2,$$

where  $h = \max_d h_d$  and  $C$  is a constant depending only on  $D$  and the  $L_d$ 's.

**Remark 5.3.5.** Since  $\mathcal{D}$  is a bounded domain, we obtain also the  $L^2(\mathcal{D})$  bound  $\|\Psi_{\mathbf{h},\mathbf{i}} - \Psi_{\mathbf{i}}\|_2 \leq C\Lambda_{\mathbf{i}}h^2$ .  $\square$

Since the approximation error in Lemma 5.3.4 depends on  $h = \max_d h_d$ , we shall from now on assume that the  $h_d$ 's are chosen so that they are of the same order, i.e.,  $\max_{j \neq k} \frac{h_j}{h_k} = O(1)$  as  $h \rightarrow 0$ , and treat only  $h$  as the mesh size. As a consequence the total number of grid points satisfies the following scaling

$$n_{\mathbf{h}} = \prod_{d=1}^D (L_d/h_d + 1) \asymp h^{-D}. \quad (5.19)$$

### 5.3.2 Balancing Numerical and Statistical Errors

Now we use the spectral error bounds in Lemma 5.3.4 to obtain  $L^2(\mathcal{D})$  and  $L^\infty(\mathcal{D})$  error bounds for FE representations of Matérn-type GP priors (Lemma 5.3.6). These prior bounds, combined with Proposition 5.2.2, will yield our main result (Theorem 5.3.8). Let

$$u_{\mathbf{h}} = \kappa^{s-D/2} \sum_{\mathbf{i} \in [K_1] \times \dots \times [K_D]} (\kappa^2 + \Lambda_{\mathbf{h},\mathbf{i}})^{-s/2} \xi_{\mathbf{i}} \Psi_{\mathbf{h},\mathbf{i}}, \quad \xi_{\mathbf{i}} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1),$$

be the FE representation of the Matérn-type GP  $u$  in (5.15). Recall that we are interested in estimating the function  $f_0(\mathbf{x}) = \mathbb{E}[Y|\mathbf{X} = \mathbf{x}]$  based on i.i.d. samples  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^N$  with prior  $\Pi_N = \text{Law}(u_{\mathbf{h}_N})$  for the fixed design regression setting and  $\Pi_N = \text{Law}(\Phi(u_{\mathbf{h}_N}))$  for the binary classification setting, where  $\mathbf{h}_N = (h_{N,1}, \dots, h_{N,D})$  is to be determined. Based on the discussion in Section 5.2.4, it suffices to quantify the approximation error of  $u$  defined in (5.15) by  $u_{\mathbf{h}_N}$ .

**Lemma 5.3.6.** *Recall that  $h = \max_d h_d$ . Suppose  $s > D/2$ . It holds that*

$$\mathbb{E}\|u_{\mathbf{h}} - u\|_2^2 \leq C\kappa^{2s-D}h^{(2s-D)\wedge 4},$$

where  $C$  is a constant independent of  $\kappa$  and  $h$ . Furthermore the Matérn-type GP  $u$  defined in (5.15) belongs almost surely to  $\mathcal{C}^\beta(\mathcal{D})$  for  $0 < \beta < 1 \wedge (s - D/2)$ . Moreover, for  $s > D$  it holds that

$$\mathbb{E}\|u_{\mathbf{h}} - u\|_\infty^2 \leq C\kappa^{2s-D}h^{(2s-2D)\wedge 4},$$

where  $C$  is a constant independent of  $\kappa$  and  $h$ .

**Remark 5.3.7.** *The  $L^2$  error bound has been shown to hold in greater generality, see e.g. [Bolin et al., 2020, Theorem 2.10] and [Cox and Kirchner, 2020, Theorem 2]. A sharper  $L^\infty$  error bound was shown in [Cox and Kirchner, 2020, Theorem 3] when  $D = 1$ , while our result holds for general dimension  $D$ .*

As a corollary of Proposition 5.2.2 we have the following main result, presented in terms of the scaling of  $h_N = \max_d h_{N,d}$ . Notice that the concentration function defined in (5.11) depends implicitly on  $s$  through  $u$ .

**Theorem 5.3.8.** *1. Fixed design regression: Consider the Matérn-type GP  $u$  defined by*

(5.4) *with  $s > D$ . Suppose  $\varepsilon_N$  satisfies  $\phi_{f_0}(\varepsilon_N; u, \|\cdot\|_\infty) \leq N\varepsilon_N^2$ . Set*

$$h_N \asymp N^{-\frac{1}{(2s-2D)\wedge 4}} \tag{5.20}$$

*with a large enough proportion constant. Then, for any sufficiently large  $M > 0$ ,*

$$\mathbb{E}_{f_0} \Pi_N \left( f : \|f - f_0\|_N \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \xrightarrow{N \rightarrow \infty} 1,$$

where we recall  $\|f\|_N^2 = N^{-1} \sum_{i=1}^N |f(\mathbf{X}_i)|^2$ .

2. *Binary classification:* Consider the Matérn-type GP  $u$  defined by (5.4) with  $s > D/2$ .

Suppose  $\varepsilon_N$  satisfies  $\phi_{\Phi^{-1}(f_0)}(\varepsilon_N; u, \|\cdot\|_2) \leq N\varepsilon_N^2$ . Set

$$h_N \asymp N^{-\frac{1}{(2s-D)\wedge 4}} \quad (5.21)$$

with a large enough proportion constant. Then, for any sufficiently large  $M > 0$ ,

$$\mathbb{E}_{f_0} \Pi_N \left( f : \|f - f_0\|_2 \leq M\varepsilon_N \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N \right) \xrightarrow{N \rightarrow \infty} 1.$$

**Remark 5.3.9.** *Theorem 5.3.8 provides a scaling of  $h_N$  so that the sequence of posteriors with respect to the FE prior  $\Pi_N$  achieves the same contraction rate as if the Matérn-type prior  $\Pi$  was used. We remark that a refined analysis of the rate at which the posterior probabilities go to 1 could be used to obtain similar conclusions for the posterior means under suitable assumptions, i.e.*

$$\begin{aligned} \mathbb{E}_{f_0} d_N(\hat{f}, f_0)^2 &\lesssim \varepsilon_N^2, & \hat{f} &= \int f d\Pi(f \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N), \\ \mathbb{E}_{f_0} d_N(\hat{f}_N, f_0)^2 &\lesssim \varepsilon_N^2, & \hat{f}_N &= \int f d\Pi_N(f \mid \{(\mathbf{X}_i, Y_i)\}_{i=1}^N). \end{aligned}$$

*In other words, the sequence of posterior means with respect to  $\Pi_N$  converges to  $f_0$  at the same rate as those with respect to  $\Pi$ , thereby giving a more interpretable conclusion. For fixed design regression, this follows from [van der Vaart and van Zanten, 2011, Theorem 1] and Jensen's inequality with  $d_N = \|\cdot\|_N$ . For binary classification, using again Jensen's*

inequality and the fact that  $|f| \leq 1$  we have

$$\begin{aligned} \|\widehat{f} - f_0\|_2^2 &\leq \int \|f - f_0\|_2^2 d\Pi(f|\{\mathbf{X}_i, Y_i\}_{i=1}^N) \\ &\leq M^2 \varepsilon_N^2 + 4|\mathcal{D}|\Pi(f : \|f - f_0\|_2 \geq M\varepsilon_N|\{\mathbf{X}_i, Y_i\}_{i=1}^N), \end{aligned}$$

where  $|\mathcal{D}|$  is the Lebesgue measure of  $\mathcal{D}$ . Therefore a rate faster than  $\varepsilon_N^2$  on the decay of the posterior probability suffices, which is satisfied under mild assumptions [Ghosal et al., 2000, Theorems 2.2 and 2.3].  $\square$

**Remark 5.3.10.** For the regression setting, (5.20) together with (5.19) gives the scaling for the total number of grid points needed,

$$n_{\mathbf{h}_N} \asymp N^{\frac{D}{(2s-2D)\wedge 4}}.$$

In particular when  $s > 3D/2$ ,  $D = 1, 2, 3$ , the exponent for  $N$  is less than one and we have  $n_{\mathbf{h}_N} \ll N$  asymptotically. For classification,  $s > D$  suffices. However, we remark that the proportion constant depends implicitly on  $\kappa$  and the  $L_d$ 's as can be seen from (5.19) and Lemma 5.3.6. In particular, if both  $\kappa$  and the  $L_d$ 's are large, which reflects the case of a rapidly changing field over a large spatial domain, then  $N$  may need to be large enough in order for  $n_{\mathbf{h}_N}$  to be smaller than  $N$ . We shall demonstrate through simulation studies in Section 5.4 that for moderate  $\kappa$  and  $L_d$ 's one can achieve  $n_{\mathbf{h}} < N$  when  $N = O(10^2)$  for a one-dimensional example and  $N = O(10^3)$  for a two-dimensional one, thereby suggesting that Theorem 5.3.8 has some practical implication.  $\square$

The scaling of  $h_N$  in Theorem 5.3.8 ensures that the numerical errors in the FE representations of a true Matérn-type GPs  $u$  do not impact the corresponding contraction rates. In the remainder of this section we give an example where the rates  $\varepsilon_N$  with respect to the true Matérn-type GP  $u$  can be explicitly computed under a smoothness assumption on the

truth  $f_0$ . For this purpose we introduce a notion of regularity of  $f_0$  based on the orthonormal basis  $\{\Psi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}^D}$ . Let  $S$  be an even function in the Schwartz space  $\mathcal{S}(\mathbb{R})$  satisfying

$$0 \leq S \leq 1, \quad S \equiv 1 \text{ on } \left[-\frac{1}{2}, \frac{1}{2}\right], \quad \text{supp}(S) \subset [-1, 1].$$

Define the space

$$B_{\infty, \infty}^{\beta} = \left\{ f = \sum_{\mathbf{i} \in \mathbb{N}^D} f_{\mathbf{i}} \Psi_{\mathbf{i}} : \|f\|_{B_{\infty, \infty}^{\beta}} = \sup_{j \in \mathbb{N}} 2^{\beta j} \|S_j(\sqrt{\Delta})f(\cdot) - f(\cdot)\|_{\infty} < \infty \right\},$$

where  $S_j(\cdot) = S(2^{-j}\cdot)$  and

$$S_j(\sqrt{\Delta})f = \sum_{\mathbf{i} \in \mathbb{N}^D} S_j(\sqrt{\Lambda_{\mathbf{i}}})f_{\mathbf{i}} \Psi_{\mathbf{i}}.$$

**Proposition 5.3.11.** *Suppose  $f_0 \in B_{\infty, \infty}^{\beta}$  and set  $s = \beta + \frac{D}{2}$  in the definition of  $u$ . Then for  $\varepsilon_N$  a large enough multiple of  $N^{-\beta/(2\beta+D)}$ , we have  $\phi_{f_0}(\varepsilon_N; u, \|\cdot\|_2) \leq N\varepsilon_N^2$  and  $\phi_{f_0}(\varepsilon_N; u, \|\cdot\|_{\infty}) \leq N\varepsilon_N^2$ .*

The space  $B_{\infty, \infty}^{\beta}$  can be seen as a Besov-type space tailored to our specific setting, where the prior support associated with the Matérn-type GP  $u$  consists of functions defined as series expansions in terms of the  $\Psi_{\mathbf{i}}$ 's. Similar function spaces have been considered in Castillo et al. [2014]. As the usual Besov spaces, functions in  $B_{\infty, \infty}^{\beta}$  should be understood to have regularity of order  $\beta$ , in which case the contraction rate  $N^{-\beta/(2\beta+D)}$  matches the usual minimax optimal rate for estimating  $\beta$ -regular functions.

## 5.4 Simulation Study

The aim of this section is to complement the understanding given by Theorem 5.3.8 through numerical simulations in the regression setting. We consider one and two-dimensional ex-



amples in Sections 5.4.1 and 5.4.2, respectively.

The general set up is as follows. Let  $\{\mathbf{X}_i\}_{i=1}^N$  be fixed design points in the domain  $\mathcal{D}$  and  $\{Y_i\}_{i=1}^N$  be noisy observations generated from

$$Y_i = f_0(\mathbf{X}_i) + \eta_i, \quad \eta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \tau^2),$$

where  $f_0$  is the ground truth and  $\tau$  is known. We compare two approaches for inferring  $f_0$ , namely the covariance function (CF) approach and the finite element (FE) approach with mass lumping. They can be summarized as follows:

$$\mathbf{y} \sim \mathcal{N}(\mathbf{f}_N, \tau^2 I_N), \quad \mathbf{f}_N \sim \mathcal{N}(\mathbf{0}, \Sigma) \implies \hat{\mathbf{f}}_{\text{CF}} = \Sigma(\Sigma + \tau^2 \mathbf{I}_N)^{-1} \mathbf{y},$$

where  $\Sigma = \{c_{\text{Mat}}(\mathbf{X}_i, \mathbf{X}_j)\}$ ; and

$$\mathbf{y} \sim \mathcal{N}(\mathbf{S}\mathbf{w}, \tau^2 I_N), \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1}) \implies \hat{\mathbf{f}}_{\text{FE}} = \mathbf{S}(\mathbf{S}^\top \mathbf{S} + \tau^2 \mathbf{Q})^{-1} \mathbf{S}^\top \mathbf{y},$$

where  $\mathbf{S}_{ij} = e_j(X_i)$  is as in Section 5.2.3 and  $\mathbf{Q} = (\kappa^2 \mathbf{M} + \mathbf{G})[\widetilde{\mathbf{M}}^{-1}(\kappa^2 \mathbf{M} + \mathbf{G})]^{s-1}$  as in Proposition 5.2.1 but with the lumped mass matrix  $\widetilde{\mathbf{M}}$  instead. As noted in Remark 5.3.9, we shall compare the error  $\|\hat{\mathbf{f}}_{\text{CF}} - \mathbf{f}_0\|_N$  and  $\|\hat{\mathbf{f}}_{\text{FE}} - \mathbf{f}_0\|_N$  when an increasing number of grid points ( $n_h$ ) is used in the FE approach, where  $\mathbf{f}_0 = (f_0(\mathbf{X}_1), \dots, f_0(\mathbf{X}_N))^\top$  and  $\|\cdot\|_N$  is the vector 2-norm normalized by  $1/\sqrt{N}$ . Note that the CF and FE approaches studied here are not exactly those analyzed in Theorem 5.3.8, i.e., the error of going from the CF approach to the Matérn-type prior (expected to be small by Proposition 5.3.1) and that of the lumped mass procedure were not accounted for. However, we remark that both errors do not lead to a significant difference in the numerical results and we will only focus on the CF and FE approaches, which are used in practice.

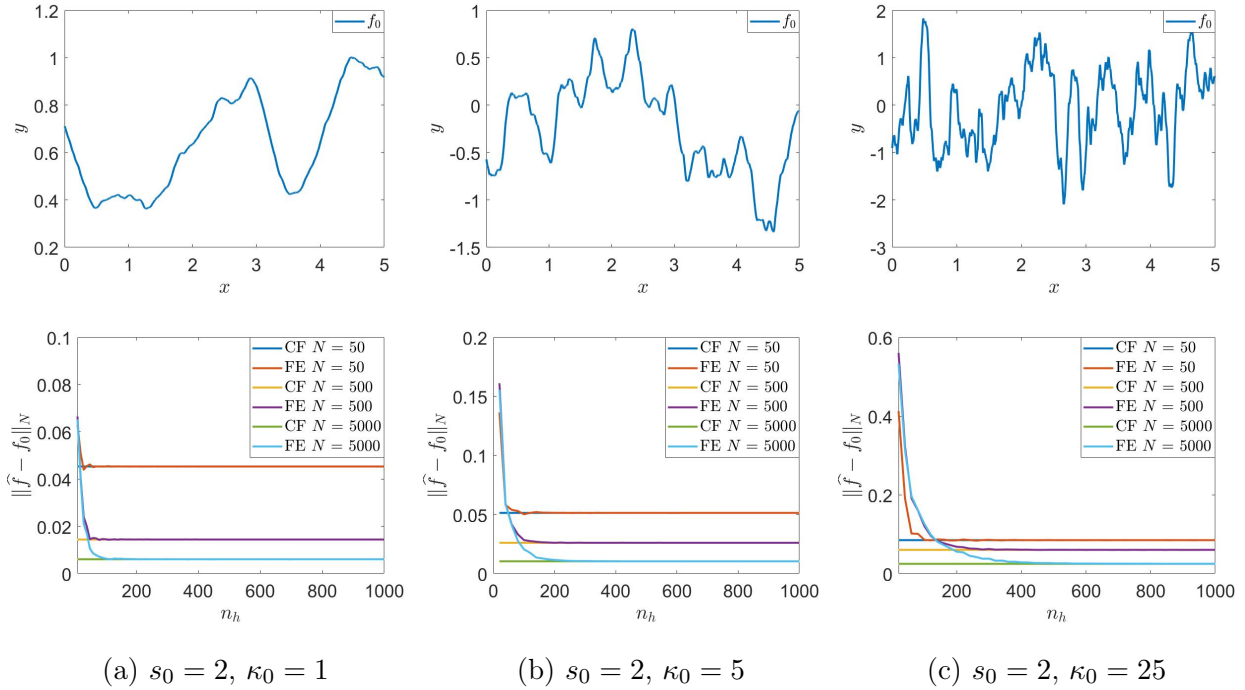


Figure 5.1: The three columns represent simulations for  $\kappa_0 = 1, 5$  and  $25$  respectively with  $s_0 = 2$  in all cases. The upper row shows plots of  $f_0$ . The lower row compares the estimation error  $\|\hat{\mathbf{f}} - \mathbf{f}_0\|_N$  between the covariance function (CF) approach and the finite element (FE) approach as  $n_h$  increases, for three levels of data  $N = 50, 500$  and  $5000$ .

#### 5.4.1 The One-dimensional Case

To start with, let  $\{\mathbf{X}_i\}_{i=1}^N$  be fixed design points generated from the uniform distribution over  $[0, L]$ . We shall consider  $f_0$ 's generated from the following series expansion (with a sufficiently high truncation)

$$f_0(x) \sim \frac{\kappa_0^{-1/2} \xi_0}{\sqrt{L}} + \frac{\sqrt{2} \kappa_0^{s_0-1/2}}{\sqrt{L}} \sum_{i=1}^{\infty} \left[ \kappa_0^2 + \left( \frac{i\pi}{L} \right)^2 \right]^{-s_0/2} \left[ \xi_i \cos \left( \frac{i\pi x}{L} \right) + \zeta_i \sin \left( \frac{i\pi x}{L} \right) \right], \quad (5.22)$$

where  $\xi_i, \zeta_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ . Notice that (5.22) is defined in the same spirit as (5.4) except that the full trigonometric basis is used, so that the random field (5.22) does not have a prescribed boundary condition. Our motivation to not consider here a Neumann boundary condition

is to illustrate that similar conclusions as those suggested by our theory can be expected in more general settings. Notice again that there are two parameters  $s_0$  and  $\kappa_0$ , which control the smoothness and correlation lengthscale respectively. We will vary both  $s_0$  and  $\kappa_0$  in the following simulations.

For both the CF and FE approaches, we use the same parameters  $s_0$  and  $\kappa_0$  that are used to generate  $f_0$ . In other words, we consider the Matérn covariance (5.2) with parameters  $\nu = s_0 - 1/2$ ,  $\kappa = \kappa_0$  and  $\sigma^2$  given in (5.16), and FE approximation (5.6) with  $s = s_0$  and  $\kappa = \kappa_0$ . For the FE approach, we construct the approximation over the larger interval  $[-\rho, L + \rho]$  where  $\rho = \sqrt{8\nu}/\kappa$  to reduce the boundary effects suggested in Proposition 5.3.1. Three levels of data  $N = 50, 500$  and  $5000$  are considered and, for each  $N$ , we study the performance for the FE approach as the number  $n_h$  of grid points increases. Finally we let  $L = 5$  and  $\tau = 0.1 \cdot \|\mathbf{f}_0\|_2/\sqrt{N}$ , which amounts to about 10% error.

Figure 5.1 shows the results when we fix the smoothness  $s_0 = 2$  and vary  $\kappa_0 = 1, 5$  and  $25$ . We see that the estimation error for the FE approach decreases to that of the CF approach after certain threshold  $n_h^*$ . In other words, discretization at the level of  $n_h^*$  for the FE approach is sufficient to yield the same estimation performance as the CF approach. The value of  $n_h^*$  is seen to be smaller than the sample size when  $N = 500$  and is of an order of magnitude smaller when  $N = 5000$ , in the same spirit as the scaling suggested in Theorem 5.3.8. The fact that  $n_h^*$  is larger than the sample size when  $N = 50$  can be explained by the large proportion constant in Remark 5.3.10. Furthermore such proportion constant increases with  $\kappa$ , as suggested by the larger  $n_h^*$  for a larger  $\kappa$ .

To further understand the effect of the smoothness  $s_0$ , we perform two more simulations for (a)  $s_0 = 1, \kappa_0 = 1$  and (b)  $s_0 = 3, \kappa_0 = 25$ . For (a) we see in Figure 5.2a that the  $n_h^*$ 's in this case are much larger than the  $s_0 = 2$  cases. This is due to the roughness of the truth and the prior used and hence a large number of grid points are needed for accurate approximation even if  $\kappa$  is small. On the other hand when  $s_0 = 3$ , Figure 5.2b shows qualitatively similar

results as in Figure 5.1 in the sense that  $n_h^*$  is asymptotically much smaller than  $N$ . Moreover the  $n_h^*$ 's are seen to be smaller than those when  $s_0 = 2$ ,  $\kappa_0 = 25$ , as the underlying field is smoother and the required scaling suggested by Theorem 5.3.8 is smaller.




Figure 5.2: Comparison of estimation error  $\|\hat{\mathbf{f}} - \mathbf{f}_0\|_N$  between the covariance function (CF) approach and the finite element (FE) approach on three data levels  $N = 50, 500, 5000$  for (a)  $s_0 = 1, \kappa_0 = 1$  and (b)  $s_0 = 3, \kappa_0 = 25$ .

#### 5.4.2 The Two-dimensional Case

Now we move on to the more practically relevant two-dimensional case following a similar set up as above. Let  $\{\mathbf{X}_i\}_{i=1}^N$  be fixed design points generated from the uniform distribution over the square  $[0, L]^2$  and  $f_0$  be generated similarly as (5.4) with  $\psi_i$ 's the full trigonometric basis (i.e. elements of the form  $\sin(10\pi x_1/L) \sin(5\pi x_2/L)$ ,  $\cos(3\pi x_1/L) \sin(9\pi x_2/L)$ , etc.) so that there is no prescribed boundary condition for  $f_0$ . We shall again compare the CF and FE approaches when  $f_0$  is generated with different values of  $s_0$  and  $\kappa_0$ .

The exact procedure for the CF and FE approaches will be completely analogous to the 1D case. In particular, the same parameters  $s_0$  and  $\kappa_0$  that generate  $f_0$  are used and furthermore the FE approach is carried out over the larger domain  $[-\rho, L + \rho]^2$  with  $\rho = \sqrt{8\nu}/\kappa$  to reduce boundary effects. However we remark that in the 2D case extending the domain has a larger impact on the performance of the FE approach than the 1D case. The reason is that to achieve the same mesh size within the domain  $[0, L]^2$ , the FE approach over  $[-\rho, L + \rho]^2$  will

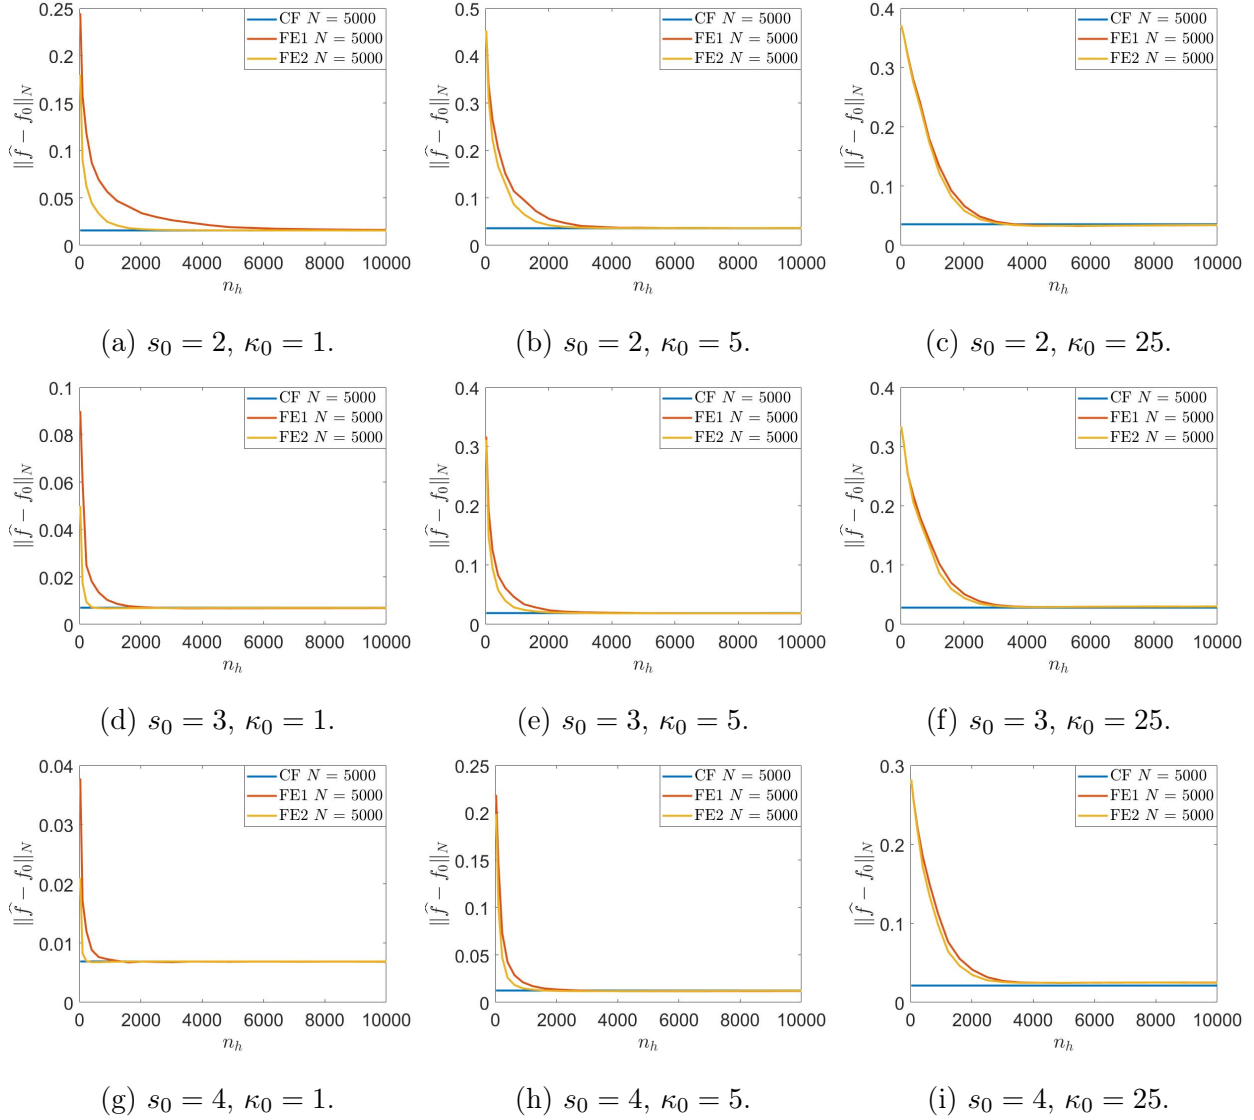


Figure 5.3: Comparison of the estimation error  $\|\hat{\mathbf{f}} - \mathbf{f}_0\|_N$  as  $n_h$  increases between the covariance function (CF) approach and two finite element approaches where FE1 is computed over  $[-\rho, L + \rho]^2$  and FE2 is computed over  $[-0.1\rho, L + 0.1\rho]^2$ . Simulation results for different combinations of  $s_0$  and  $\kappa_0$  are shown.

require many more grid points than over  $[0, L]^2$ . In particular if a uniform partition of mesh size  $h$  as in Section 5.3.1 is adopted, then the overall increment of number of grid points is

$$\left(\frac{L + 2\rho}{h}\right)^2 - \left(\frac{L}{h}\right)^2 = \left(\frac{2\rho}{h}\right)\left(\frac{2L}{h}\right) + \left(\frac{\rho}{h}\right)^2. \quad (5.23)$$

The factor  $2L/h$  makes (5.23) much larger than the increment  $2\rho/h$  in each dimension and leads to a much larger saturation threshold  $n_h^*$  (that we have introduced in Section 5.4.1). For this reason we consider an alternate FE approach carried out over the smaller domain  $[-0.1\rho, L+0.1\rho]^2$  and compare its performance with the other FE approach over  $[-\rho, L+\rho]^2$ . For the simulations that we are going to present, we fix  $N = 5000$ ,  $L = 5$ ,  $\tau = 0.1 \cdot \|\mathbf{f}_0\|_2 / \sqrt{N}$  and vary  $s_0 \in \{2, 3, 4\}$ ,  $\kappa_0 \in \{1, 5, 25\}$ , where we recall  $\mathbf{f}_0 = (f_0(\mathbf{X}_1), \dots, f_0(\mathbf{X}_N))^\top$ . Similar parameter settings were considered in Bolin and Lindgren [2013].

Figure 5.3 shows qualitatively similar results as those in Section 5.4.1, where the estimation error of both FE approaches decreases to that of the CF approach after certain threshold  $n_h^*$ . Although Theorem 5.3.8 suggests a smaller asymptotic scaling for  $n_h^*$  only when  $s > 3$ , the simulation results suggest that this is true for  $s = 3$  and even for  $s = 2$  when the FE approach is computed over  $[-0.1\rho, L + 0.1\rho]^2$ . Furthermore, no estimation accuracy is lost when this smaller domain is used and a smaller  $n_h^*$  suffices so that it is more favorable, especially when  $\kappa$  is small or equivalently when  $\rho$  is large. Finally we remark that  $s_0 = 2, \kappa_0 = 25$  corresponds to a very rapidly changing field and even in this case we have  $n_h^* < N$  when  $N = 5000$ , which is also a realistic amount of data relative to the domain size. Following the same intuition as provided in Figure 5.1, it is reasonable to expect that one can take  $n_h$  an order of magnitude smaller than  $N$  when e.g.  $N = 50000$ . Therefore we believe the results in Theorem 5.3.8 have practical implications for a wide range of moderate nonasymptotic regimes and can provide some meaningful insights for real world applications.

## 5.5 Discussion and Open Directions

In this chapter we have employed a Bayesian nonparametrics framework to provide new understanding on the choice of the dimension  $n_h$  in FE approaches to GP regression and classification. Our theory and simulation studies demonstrate that under mild smoothness assumptions one can take  $n_h \ll N$  for a wide range of practical scenarios without hindering

the estimation accuracy, leading to a second layer of computational gain on top of the well-celebrated sparsity provided by the FE approach.

One of the key elements in our analysis is the framework of van der Vaart and van Zanten [2008a] which allows to translate prior approximation guarantees to the posteriors. In the context of GP regression and classification, this boils down to controlling, respectively, the error  $\mathbb{E}\|u_{h_N} - u\|_\infty^2$  and  $\mathbb{E}\|u_{h_N} - u\|_2^2$  as in Proposition 5.2.2. In Sections 5.3 and 5.4 we have used this framework to analyze the Matérn covariance approach and linear FE approximations thereof on a hyperrectangle. We remark that the applicability of our framework goes beyond this simple setting with the following possible extensions.

**General Elliptic Operators.** One can define a nonstationary Matérn-type GP similarly as in (5.3) by replacing the operator  $\kappa^2 - \Delta$  with a more general elliptic operator  $\kappa^2 - \nabla \cdot (\mathbf{H} \nabla)$  on a general domain, where  $\kappa$  and  $\mathbf{H}$  are smooth functions taking values in real numbers and matrices, respectively. The error analysis in Bolin et al. [2020], Cox and Kirchner [2020] on FE approximations of these random fields together with Proposition 5.2.2 would give a sufficient scaling of the mesh size  $h_N$ . We remark that in this case there would not be an easily computable covariance function approach to compare with, but one can still arrive at the conclusion that there is no need to discretize beyond the threshold implied by Proposition 5.2.2.

**Higher Order FEM.** Higher order finite elements may be employed when the smoothness parameter  $s$  is large. In particular, results from Bolin et al. [2020], Cox and Kirchner [2020] show that the  $L^2$  approximation rate in Lemma 5.3.6 can be improved to  $h^{(2s-D) \wedge (2p+2)}$  when polynomials of order  $p$  are used. As a result, the scaling for  $h_N$  in Theorem 5.3.8 can be improved accordingly. The  $L^\infty$  approximation rate for  $p > 1$  remains an interesting open question.

**Rational Approximation.** When the smoothness parameter  $s$  is not an integer, the favorable sparsity of the FE approximation is lost. For this reason, Bolin and Kirchner [2020] proposed a rational approximation of the fractional operator which retains sparsity. The resulting approximate field is shown to satisfy a similar  $L^2$  approximation error bound as in Lemma 5.3.6, with an additional term coming from the rational approximation which can be made as small as desired. Following a similar argument as in Bolin and Kirchner [2020], our proof for the  $L^\infty$  bound can also be extended to the rational approximate field. These prior approximation rates can again be combined with Proposition 5.2.2 to yield a sufficient scaling of the mesh size.

**Learning the Lengthscale.** A novel aspect of our error bounds for FE prior representations is that we keep track of the (inverse) lengthscale parameter  $\kappa$ . Our theory and numerical experiments help explain the need of finer discretization when the lengthscale is shorter. An interesting direction for further research is the design of algorithms for the simultaneous learning of (i) adaptive FE meshes for GP representations; and (ii) spatially-variable lengthscale parameters  $\kappa(x)$  in nonstationary Matérn-type models.

**Beyond Regression and Classification.** Lastly, we also envision that the framework we have introduced can be adopted in other problems such as density estimation [Ghosal et al., 2000] and nonlinear Bayesian inverse problems [Stuart, 2010]. The results from van der Vaart and van Zanten [2008a] readily extend our framework to density estimation problems and it is an interesting direction to extend Proposition 5.2.2 to Bayesian inverse problem settings building, for instance, on Nickl et al. [2020], Giordano and Nickl [2020].



## 5.6 Proof of Main Results

*Proof of Proposition 5.2.1.* It suffices to find the coordinates of (5.6) in terms of the finite element basis  $e_{h,i}$ 's. Taking inner product of (5.6) with  $e_{h,j}$ , we get

$$\sum_{i=1}^{n_h} (\kappa^2 + \lambda_{h,i})^{-s/2} \xi_i \langle \psi_{h,i}, e_{h,j} \rangle = \sum_{i=1}^{n_h} w_i \langle e_{h,i}, e_{h,j} \rangle \quad j = 1, \dots, n_h,$$

and the system

$$\mathbf{R}(\kappa^2 \mathbf{I}_{n_h} + \mathbf{\Lambda})^{-s/2} \boldsymbol{\xi} = \mathbf{M} \mathbf{w}, \quad (5.24)$$

where  $\mathbf{R}_{ij} = \langle e_{h,i}, \psi_{h,j} \rangle$  and  $\mathbf{\Lambda}$  is the diagonal matrix with entries  $\mathbf{\Lambda}_{ii} = \lambda_{h,i}$ . It now remains to relate  $\mathbf{\Lambda}$  with the matrices  $\mathbf{R}, \mathbf{M}, \mathbf{G}$ . Since the  $\psi_{h,i}$ 's form an orthonormal basis, we have  $e_{h,i} = \sum_{j=1}^{n_h} \langle e_{h,i}, \psi_{h,j} \rangle \psi_{h,j} = \sum_{j=1}^{n_h} \mathbf{R}_{ij} \psi_{h,j}$ , which implies that  $\psi_{h,i} = \sum_{j=1}^{n_h} (\mathbf{R}^{-1})_{ij} e_{h,j}$ . The fact that  $\psi_{h,i}$ 's are (variational) eigenvectors of  $-\Delta_h$  with corresponding eigenvalues  $\lambda_{h,i}$  gives  $\mathbf{G} \mathbf{R}^{-\top} = \mathbf{M} \mathbf{R}^{-\top} \mathbf{\Lambda}$ , which together with the fact that  $\mathbf{R} \mathbf{R}^{\top} = \mathbf{M}$  further implies  $\mathbf{\Lambda} = \mathbf{R}^{-1} \mathbf{G} \mathbf{R}^{-\top}$ . The result then follows by plugging such representation for  $\mathbf{\Lambda}$  into (5.24).  $\square$

*Proof of Lemma 5.3.2.* First note that we have  $\lambda_0 = 0$ ,  $\psi_0 \equiv 1/L$  and

$$\lambda_i = \left( \frac{i\pi}{L} \right)^2, \quad \psi_i = \frac{2}{L} \cos \left( \frac{i\pi x}{L} \right), \quad i = 1, 2, \dots$$

We then have

$$\begin{aligned} |\lambda_{h,i} - \lambda_i| &\leq \frac{1}{2 + \cos(i\pi h/L)} \left| \frac{6}{h^2} \left( 1 - \cos \frac{i\pi h}{L} \right) - \left( \frac{i\pi}{L} \right)^2 \left( 2 + \cos \frac{i\pi h}{L} \right) \right| \\ &\leq \left| \frac{6}{h^2} \left( 1 - \cos \frac{i\pi h}{L} \right) - \left( \frac{i\pi}{L} \right)^2 \left( 2 + \cos \frac{i\pi h}{L} \right) \right|. \end{aligned}$$

Expanding the last expression based on the Taylor series of  $\cos x$  we obtain

$$|\lambda_{h,i} - \lambda_i| \leq \left| \frac{3}{h^2} \left( \frac{i\pi h}{L} \right)^2 + \frac{1}{h^2} O \left( \frac{i\pi h}{L} \right)^4 - 3 \left( \frac{i\pi}{L} \right)^2 + \left( \frac{i\pi}{L} \right)^2 O \left( \frac{i\pi h}{L} \right)^2 \right| \leq C \lambda_i^2 h^2.$$

For the approximation error of the eigenfunctions, we first compute the normalizing constants  $c_i$ 's. For  $i = 0$  we notice that  $\psi_{h,0}$  is constant and hence  $c_0 = 1/L$ . For general  $i$ 's, we denote  $\psi_{h,i} = \sum_{k=0}^K z_{h,i,k} e_{h,k}$  and compute

$$\begin{aligned} \langle \psi_{h,i}, \psi_{h,j} \rangle &= \sum_{k=0}^K \sum_{\ell=0}^K z_{h,i,k} z_{h,j,\ell} \langle e_{h,k}, e_{h,\ell} \rangle \\ &= z_{h,i,0} \sum_{\ell=0}^K z_{h,j,\ell} \langle e_{h,0}, e_{h,\ell} \rangle + z_{h,i,K} \sum_{\ell=0}^K z_{h,j,\ell} \langle e_{h,K}, e_{h,\ell} \rangle \\ &\quad + \sum_{k=1}^{K-1} z_{h,i,k} \sum_{\ell=0}^K z_{h,j,\ell} \langle e_{h,k}, e_{h,\ell} \rangle \\ &= z_{h,i,0} \left( \frac{z_{h,j,0}}{3} + \frac{z_{h,j,1}}{6} \right) h + z_{h,i,K} \left( \frac{z_{h,j,K}}{3} + \frac{z_{h,j,K-1}}{6} \right) h \\ &\quad + \sum_{k=1}^{K-1} z_{h,i,k} \left( \frac{z_{h,j,k-1}}{6} + \frac{2z_{h,j,k}}{3} + \frac{z_{h,j,k+1}}{6} \right) h \\ &= c_i c_j h \left( \frac{1}{6} \cos \frac{j\pi h}{L} + \frac{1}{3} \right) \left[ 1 + (-1)^{i+j} + 2 \sum_{k=1}^{K-1} \cos \left( \frac{k i \pi h}{L} \right) \cos \left( \frac{k j \pi h}{L} \right) \right], \end{aligned}$$

where we have used that  $\cos(a-t) + \cos(a+t) = 2 \cos(a) \cos(t)$ . Using further the fact that  $2 \cos(a) \cos(b) = \cos(a+b) + \cos(a-b)$ , we have

$$2 \sum_{k=1}^{K-1} \cos \left( \frac{k i \pi h}{L} \right) \cos \left( \frac{k j \pi h}{L} \right) = \sum_{k=1}^{K-1} \cos \left( \frac{k(i+j)\pi h}{L} \right) + \cos \left( \frac{k(i-j)\pi h}{L} \right).$$

Now letting  $t = (i + j)\pi h/L$  and denoting  $\iota$  as the imaginary unit, we have

$$\begin{aligned} \sum_{k=1}^{K-1} \cos(kt) &= \frac{1}{2} \sum_{k=1}^{K-1} e^{\iota kt} + e^{-\iota kt} = \frac{1}{2} \left[ \frac{e^{\iota t}(1 - e^{\iota(K-1)t})}{1 - e^{\iota t}} + \frac{e^{-\iota t}(1 - e^{-\iota(K-1)t})}{1 - e^{-\iota t}} \right] \\ &= -\frac{1}{2} \left[ 1 + (-1)^{i+j} \right], \end{aligned}$$

where we have used that  $Kt = (i + j)\pi$ . Similarly, we have

$$\sum_{k=1}^{K-1} \cos\left(\frac{k(i-j)\pi h}{L}\right) = \begin{cases} -\frac{1}{2} [1 + (-1)^{i-j}] & i \neq j \\ K-1 & i = j \end{cases}.$$

Therefore we have

$$\langle \psi_{h,i}, \psi_{h,j} \rangle = c_i c_j L \left( \frac{1}{6} \cos \frac{j\pi h}{L} + \frac{1}{3} \right) \delta_{ij},$$

where  $\delta_{ij}$  denotes the Kronecker delta. Therefore the  $\psi_{h,i}$ 's are orthonormal with

$$c_i = \left[ L \left( \frac{1}{6} \cos \frac{i\pi h}{L} + \frac{1}{3} \right) \right]^{-1/2}.$$

Now to bound the eigenfunction approximation error, we have

$$\|\psi_{h,i} - \psi_i\|_\infty \leq \|\psi_{h,i} - \tilde{\psi}_{h,i}\|_\infty + \|\tilde{\psi}_{h,i} - \psi_i\|_\infty,$$

where  $\tilde{\psi}_{h,i} = \frac{1}{c_i} \sqrt{\frac{2}{L}} \psi_{h,i}$ . Since  $\|\psi_{h,i}\|_\infty \leq 1$ , we have

$$\|\psi_{h,i} - \tilde{\psi}_{h,i}\|_\infty \leq \left| 1 - \frac{1}{c_i} \sqrt{\frac{2}{L}} \right| \leq C \left| L \left( \frac{1}{6} \cos \frac{i\pi h}{L} + \frac{1}{3} \right) - \frac{L}{2} \right| \leq C \left( \frac{i\pi h}{L} \right)^2 = C \lambda_i h^2,$$

where  $C$  is constant depending only on  $L$ . To bound  $\|\tilde{\psi}_{h,i} - \psi_i\|_\infty$ , notice that after the

rescaling,  $\tilde{\psi}_{h,i}$  is a linear interpolant of  $\psi_i$  over the nodes. In particular, denoting  $x_k = kh$  we have

$$\tilde{\psi}_{h,i}(x) = \psi_i(x_k) + \frac{x - x_k}{x_{k+1} - x_k} [\psi_i(x_{k+1}) - \psi_i(x_k)] \quad \text{on } [x_k, x_{k+1}].$$

Taylor expanding at  $x$  we have

$$\begin{aligned} \psi_i(x_{k+1}) &= \psi_i(x) + \psi'_i(x)(x_{k+1} - x) + \frac{\psi''_i(\eta_1)}{2}(x_{k+1} - x)^2 & x < \eta_1 < x_{k+1} \\ \psi_i(x_k) &= \psi_i(x) + \psi'_i(x)(x_k - x) + \frac{\psi''_i(\eta_2)}{2}(x_k - x)^2 & x_k < \eta_2 < x. \end{aligned}$$

Therefore

$$\tilde{\psi}_{h,i}(x) = \psi_i(x) + \frac{\psi''_i(\eta_2)}{2}(x_k - x)^2 + \frac{x - x_k}{x_{k+1} - x_k} \left[ \frac{\psi''_i(\eta_1)}{2}(x_{k+1} - x)^2 + \frac{\psi''_i(\eta_2)}{2}(x_k - x)^2 \right]$$

and since  $\|\psi''_i\|_\infty \leq \lambda_i$  we have

$$\sup_{x \in [x_k, x_{k+1}]} |\tilde{\psi}_{h,i}(x) - \psi_i(x)| \leq C\lambda_i h^2.$$

Therefore  $\|\tilde{\psi}_{h,i} - \psi_i\|_\infty \leq C\lambda_i h^2$  and the result follows.  $\square$

*Proof of Lemma 5.3.4.* Notice that

$$|\Lambda_{\mathbf{h},\mathbf{i}} - \Lambda_{\mathbf{i}}| \leq \sum_{d=1}^D |\lambda_{h_d, i_d} - \lambda_{i_d}|$$

and

$$\begin{aligned} |\Psi_{\mathbf{h},\mathbf{i}}(\mathbf{x}) - \Psi_{\mathbf{i}}(\mathbf{x})| &= \left| \prod_{d=1}^D \psi_{h_d, i_d}(x_d) - \prod_{d=1}^D \psi_{i_d}(x_d) \right| \\ &\leq \sum_{d=1}^D \left( |\psi_{h_d, i_d}(x_d) - \psi_{i_d}(x_d)| \prod_{\ell=1}^{d-1} |\psi_{i_\ell}(x_\ell)| \prod_{\ell=d+1}^D |\psi_{h_\ell, i_\ell}(x_\ell)| \right). \end{aligned}$$

Therefore the result follows from the one-dimensional estimates in Lemma 5.3.2.  $\square$

*Proof of Lemma 5.3.6.* We shall abuse the notation and order the multi index  $\mathbf{i} \in \mathbb{N}^D$  as a single sequence  $i \in \mathbb{N}$  so that

$$\begin{aligned} u &= \kappa^{s-D/2} \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} \xi_i \Psi_i, \\ u_{\mathbf{h}} &= \kappa^{s-D/2} \sum_{i=1}^{n_{\mathbf{h}}} (\kappa^2 + \Lambda_{\mathbf{h},i})^{-\frac{s}{2}} \xi_i \Psi_{\mathbf{h},i}. \end{aligned}$$

**Bound for  $\mathbb{E}\|u_{\mathbf{h}} - u\|_2^2$ .** This can be proven using the techniques in [Bolin et al., 2020, Theorem 2.10] and in the  $L^\infty$  bound that we will establish below.

**Hölder continuity.** We have

$$\begin{aligned} \mathbb{E}|u(\mathbf{x} + \mathbf{h}) - u(\mathbf{x})|^2 &\lesssim \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-s} |\Psi_i(\mathbf{x} + \mathbf{h}) - \Psi_i(\mathbf{x})|^2 \\ &\lesssim \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-s} \min\{\|\nabla \Psi_i\|_\infty^2 h^2, 1\} \\ &\lesssim \sum_{i=1}^{\infty} i^{-\frac{2s}{D}} \min\{i^{\frac{2}{D}} h^2, 1\} \\ &\lesssim \int_{x \geq 1} x^{-\frac{2s}{D}} \min\{x^{\frac{2}{D}} h^2, 1\} dx \\ &\lesssim \left[ \int_{1 \leq x \leq h^{-D}} h^2 x^{-\frac{2}{D} - \frac{2s}{D}} dx + \int_{x > h^{-D}} x^{-\frac{2s}{D}} dx \right] \lesssim h^{2s-D}, \end{aligned}$$

where we have used that  $\|\nabla\Psi_i\|_\infty \leq \sqrt{\Lambda_i}$  and Weyl's law  $\Lambda_i \asymp i^{2/D}$ . Then by [Stuart, 2010, Corollary 6.8] we have

$$\mathbb{E}|u(\mathbf{x}) - u(\mathbf{x}')|^{2p} \leq C_p |\mathbf{x} - \mathbf{x}'|^{(2s-D)p}$$

for all  $p \in \mathbb{N}$ . Kolmogorov continuity theorem [Stuart, 2010, Theorem 6.24] implies that  $u$  is  $\beta$ -Hölder for  $\beta < \frac{(2s-D)p-2D}{2p}$ . Letting  $p \rightarrow \infty$  gives the desired result.

**Bound for  $\mathbb{E}\|u_{\mathbf{h}} - u\|_\infty^2$ .** Consider two intermediate quantities

$$\tilde{u} = \kappa^{s-D/2} \sum_{i=1}^{n_{\mathbf{h}}} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} \xi_i \Psi_i, \quad (5.25)$$

$$\tilde{u}_{\mathbf{h}} = \kappa^{s-D/2} \sum_{i=1}^{n_{\mathbf{h}}} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} \xi_i \Psi_{\mathbf{h},i}. \quad (5.26)$$

We have

$$\begin{aligned} \mathbb{E}\|u - u_{\mathbf{h}}\|_\infty^2 &\leq \mathbb{E}(\|u - \tilde{u}\|_\infty + \|\tilde{u} - \tilde{u}_{\mathbf{h}}\|_\infty + \|\tilde{u}_{\mathbf{h}} - u_{\mathbf{h}}\|_\infty)^2 \\ &\leq 2 \left( \mathbb{E}\|u - \tilde{u}\|_\infty^2 + \mathbb{E}\|\tilde{u} - \tilde{u}_{\mathbf{h}}\|_\infty^2 + \mathbb{E}\|\tilde{u}_{\mathbf{h}} - u_{\mathbf{h}}\|_\infty^2 \right) \end{aligned}$$

and it suffices to bound each term. Since the  $\Psi_i$ 's are uniformly bounded and that  $\xi$  has bounded first moment, we have

$$\begin{aligned} \mathbb{E}\|u - \tilde{u}\|_\infty^2 &\leq \kappa^{2s-D} \mathbb{E} \left[ \sum_{i=n_{\mathbf{h}}+1}^{\infty} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} |\xi_i| \|\Psi_i\|_\infty \right]^2 \\ &\lesssim \kappa^{2s-D} \left[ \sum_{i=n_{\mathbf{h}}+1}^{\infty} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} \right]^2 \\ &\lesssim \kappa^{2s-D} \left( \sum_{i=n_{\mathbf{h}}+1}^{\infty} i^{-\frac{s}{D}} \right)^2 \lesssim \kappa^{2s-D} n_{\mathbf{h}}^{2-\frac{2s}{D}} \asymp \kappa^{2s-D} h^{2s-2D}. \end{aligned} \quad (5.27)$$

Similarly, by Lemma 5.3.4

$$\begin{aligned}
\mathbb{E}\|\tilde{u} - \tilde{u}_{\mathbf{h}}\|_{\infty}^2 &\lesssim \kappa^{2s-D} \left[ \sum_{i=1}^{n_{\mathbf{h}}} (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} \|\Psi_i - \Psi_{\mathbf{h},i}\|_{\infty} \right]^2 \\
&\lesssim \kappa^{2s-D} h^4 \left( \sum_{i=1}^{n_{\mathbf{h}}} \Lambda_i^{1-\frac{s}{2}} \right)^2 \lesssim \kappa^{2s-D} h^4 \left[ 1 \vee n_{\mathbf{h}}^{2+\frac{4}{d}(1-\frac{s}{2})} \right] \asymp \kappa^{2s-D} h^{(2s-2s)\wedge 4}.
\end{aligned} \tag{5.28}$$

For the last term we have by Lemma 5.3.4

$$\begin{aligned}
\mathbb{E}\|\tilde{u}_{\mathbf{h}} - u_{\mathbf{h}}\|_{\infty}^2 &\lesssim \kappa^{2s-D} \left[ \sum_{i=1}^{n_{\mathbf{h}}} \left| (\kappa^2 + \Lambda_i)^{-\frac{s}{2}} - (\kappa^2 + \Lambda_{\mathbf{h},i})^{-\frac{s}{2}} \right| \right]^2 \\
&\lesssim \kappa^{2s-D} \left[ \sum_{i=1}^{n_{\mathbf{h}}} \Lambda_i^{-\frac{s}{2}-1} |\Lambda_i - \Lambda_{\mathbf{h},i}| \right]^2 \\
&\lesssim \kappa^{2s-D} h^4 \left( \sum_{i=1}^{n_{\mathbf{h}}} \Lambda_i^{1-\frac{s}{2}} \right)^2 \lesssim \kappa^{2s-D} h^{(2s-2d)\wedge 4}.
\end{aligned} \tag{5.29}$$

The result follows by combining (5.27), (5.28), (5.29).  $\square$

*Proof of Proposition 5.3.11.* Again we shall abuse the notation and write

$$u = \kappa^{s-D/2} \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-s/2} \xi_i \Psi_i.$$

$L^{\infty}$  case. Recall that

$$\phi_{f_0}(\varepsilon; u, \|\cdot\|_{\infty}) = \inf_{g \in \mathbb{H}: \|g-f_0\|_{\infty} < \varepsilon} \|g\|_{\mathbb{H}}^2 - \log \mathbb{P}(\|u\|_{\infty} < \varepsilon).$$

By [Li and Linde, 1999, Theorem 1.2], the second term can be bounded by analyzing the  $L^\infty(\mathcal{D})$  metric entropy of  $\mathbb{H}_1$ , the unit ball of  $\mathbb{H}$ . Notice that  $\mathbb{H}_1$  takes the form

$$\mathbb{H}_1 = \left\{ \sum_{i=1}^{\infty} g_i \Psi_i : \sum_{i=1}^{\infty} g_i^2 (\kappa^2 + \Lambda_i)^s \leq 1 \right\}$$

and is contained in a Sobolev ball of order  $s$ , whose  $L^\infty(\mathcal{D})$  metric entropy is bounded by a constant times  $\varepsilon^{-\frac{D}{s}}$  (see e.g. [Edmunds and Triebel, 1996, Theorem 3.3.2]). Then [Li and Linde, 1999, Theorem 1.2] implies that

$$-\log \mathbb{P}(\|u\|_\infty < \varepsilon) \lesssim \varepsilon^{-\frac{2D}{2s-D}} = \varepsilon^{-\frac{D}{\beta}}, \quad (5.30)$$

where we used the assumption that  $s = \beta + \frac{D}{2}$ . For the first term, let  $C_0 = \|f_0\|_{B_{\infty,\infty}^\beta}$  and consider  $g = S_J(\sqrt{\Delta})f_0$  with  $J$  the smallest integer such that  $C_0 2^{-\beta J} < \varepsilon$ . Since  $f_0 \in B_{\infty,\infty}^\beta$  we have

$$\|S_j(\sqrt{\Delta})f_0 - f_0\|_\infty \leq C_0 2^{-\beta j} \quad (5.31)$$

for all  $j$ . In particular,  $\|g - f_0\|_\infty \leq C_0 2^{-\beta J} \leq \varepsilon$ . Moreover we have  $g = \sum_{i=1}^{\infty} S_J(\sqrt{\Lambda_i})f_i \Psi_i$  as a finite series since  $S_J(\sqrt{\Lambda_i}) = 0$  if  $\sqrt{\Lambda_i} > 2^J$  and hence  $g \in \mathbb{H}$ . Now since  $S_j \leq 1$ ,

$$\|g\|_{\mathbb{H}}^2 \leq \sum_{\sqrt{\Lambda_i} \leq 2^J} f_i^2 (\kappa^2 + \Lambda_i)^s = \sum_{\sqrt{\Lambda_i} \leq 1} f_i^2 (\kappa^2 + \Lambda_i)^s + \sum_{j=1}^J \sum_{2^{j-1} < \sqrt{\Lambda_i} \leq 2^j} f_i^2 (\kappa^2 + \Lambda_i)^s.$$

By (5.31) we have

$$\begin{aligned} \sum_{2^{j-1} < \sqrt{\Lambda_i} \leq 2^j} f_i^2 (\kappa^2 + \Lambda_i)^s &\lesssim 2^{2js} \sum_{2^{j-1} < \sqrt{\Lambda_i} \leq 2^j} f_i^2 \lesssim 2^{2js} \|S_j(\sqrt{\Delta})f_0 - f_0\|_2^2 \\ &\lesssim 2^{2js} \|S_j(\sqrt{\Delta})f_0 - f_0\|_\infty^2 \lesssim 2^{2(s-\beta)j}. \end{aligned}$$



Since  $J$  is the smallest integer such that  $C_0 2^{-\beta J} < \varepsilon$ , we have  $2^{-\beta J} \gtrsim \varepsilon$  and hence

$$\|g\|_{\mathbb{H}}^2 \lesssim \sum_{\sqrt{\Lambda_i} \leq 1} f_i^2 + \sum_{j=1}^J 2^{2(s-\beta)j} \lesssim 2^{2(s-\beta)J} \lesssim \varepsilon^{-\frac{2(s-\beta)}{\beta}} = \varepsilon^{-\frac{D}{\beta}}. \quad (5.32)$$

Combining (5.30) and (5.32) we deduce that

$$\phi_{f_0}(\varepsilon) \lesssim \varepsilon^{-\frac{D}{\beta}}$$

and setting  $\varepsilon_n = C n^{-\frac{\beta}{2\beta+D}}$  for a large enough constant  $C$  gives the result.

**$L^2$  case.** We have

$$\phi_{f_0}(\varepsilon; u, \|\cdot\|_2) = \inf_{g \in \mathbb{H}: \|g-f_0\|_2 < \varepsilon} \|g\|_{\mathbb{H}}^2 - \log \mathbb{P}(\|u\|_2 < \varepsilon). \quad (5.33)$$

For the second term, recall that  $u = \kappa^{s-D/2} \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-s/2} \xi_i \Psi_i$ . We then have

$$\begin{aligned} \log \mathbb{P}(\|u\|_2 < \varepsilon) &= \log \mathbb{P} \left( \kappa^{2s-D} \sum_{i=1}^{\infty} (\kappa^2 + \Lambda_i)^{-s} \xi_i^2 < \varepsilon \right) \\ &\geq \log \mathbb{P} \left( \sum_{i=1}^{\infty} i^{-\frac{2s}{D}} \xi_i^2 < C\varepsilon \right) \gtrsim \varepsilon^{-\frac{2}{\frac{2s}{D}-1}} = \varepsilon^{-\frac{D}{\beta}}, \end{aligned} \quad (5.34)$$

where the last step follows from [Dunker et al., 1998, Corollary 6] and the assumption that  $s = \beta + \frac{D}{2}$ . For the first term in (5.33), let  $C_0 = \|f_0\|_{B_{\infty, \infty}^{\beta}}$  and consider  $g = S_J(\sqrt{\Delta})f_0$  with  $J$  the smallest integer so that  $C_0 \sqrt{|\mathcal{D}|} 2^{-\beta J} < \varepsilon$ , where  $|\mathcal{D}|$  is the Lebesgue measure of  $\mathcal{D}$ . Since  $f_0 \in B_{\infty, \infty}^{\beta}$  we have

$$\|S_j(\sqrt{\Delta})f_0 - f_0\|_{\infty} \leq C_0 2^{-\beta j}$$

for all  $j$ . In particular,  $\|g - S_J(\sqrt{\Delta})f_0\|_2 \leq \sqrt{|\mathcal{D}|}\|g - S_J(\sqrt{\Delta})f_0\|_\infty \leq C_0\sqrt{|\mathcal{D}|}2^{-\beta J} \leq \varepsilon$ .

Now proceeding in the same way as the argument in the  $L^\infty$  case we obtain

$$\|g\|_{\mathbb{H}}^2 \lesssim \sum_{\sqrt{\lambda_i} \leq 1} f_i^2 + \sum_{j=1}^J 2^{2(s-\beta)j} \lesssim 2^{2(s-\beta)J} \lesssim \varepsilon^{-\frac{2(s-\beta)}{\beta}} = \varepsilon^{-\frac{D}{\beta}}, \quad (5.35)$$

where we have used the fact that  $2^{-\beta J} \gtrsim \varepsilon$  since  $J$  is the smallest integer such that  $C_0\sqrt{|\mathcal{D}|}2^{-\beta J} < \varepsilon$ . Combining (5.34) and (5.35) we deduce that

$$\phi_{f_0}(\varepsilon) \lesssim \varepsilon^{-\frac{D}{\beta}},$$

and setting  $\varepsilon_n = Cn^{-\frac{\beta}{2\beta+D}}$  for a large enough constant  $C$  gives the result.  $\square$

## CHAPTER 6

# LOCAL REGULARIZATION OF NOISY POINT CLOUDS: IMPROVED GLOBAL GEOMETRIC ESTIMATES AND DATA ANALYSIS

### 6.1 Introduction

Several techniques for the analysis of high dimensional data build on the observation that data-generating mechanisms can often be described by few degrees of freedom. In this chapter we study graph-based methods that employ similarity relationships between data points to uncover the low intrinsic dimension and geometric structure of datasets. Graph-based learning provides a well-balanced compromise between accuracy and interpretability [Coifman and Lafon, 2006], and is popular in a variety of unsupervised and semi-supervised tasks [Zhu, 2005, Von Luxburg, 2007]. These methods have been extensively analyzed in the idealized setting where the data is sampled from a low-dimensional manifold and similarities are computed using the ambient Euclidean distance or the geodesic distance, see e.g. Coifman and Lafon [2006], Singer [2006], Burago et al. [2015], García Trillos et al. [2019a]. The manifold setting is truthful in spirit to the presupposition that data arising from structured systems may be described by few degrees of freedom, but it is not so in that the data are typically noisy. The aim of this chapter is to provide new mathematical theory under the more general and realistic model assumption that the data consist of random perturbations of low-dimensional features lying on a manifold.

By relaxing the manifold assumption we bring forward two fundamental questions that are at the heart of graph-based learning but have not been accounted for by previous theory. First, how to define the inter-point similarities between noisy data points in order to approximate the Euclidean distances between unperturbed data-points? Second, is it possible to recover global geometric features of the manifold from suitably-defined similarities between

noisy data points? We will show by rigorous mathematical reasoning that:

- (i) Denoising inter-point distances leads to an improved approximation of the hidden Euclidean distance between unperturbed points. We illustrate this general idea by analyzing a simple, easily-computable similarity defined in terms of a local-regularization of the noisy dataset.
- (ii) Graph-based objects defined via locally regularized similarities can be guaranteed to satisfy improved error bounds in the recovery of global geometric properties. We illustrate this general idea by showing the spectral approximation of an unnormalized  $\varepsilon$ -graph Laplacian to a Laplace operator defined on the underlying manifold.

In addition to giving theoretical support for the denoising of point clouds, we study the practical use of local regularization in classification problems. Our analytically tractable local-regularization depends on a parameter that modulates the amount of localization, and our analysis suggests the appropriate scaling of said parameter with the level noise level. In our numerical experiments we show that in semi-supervised classification problems this parameter may be chosen by cross-validation, ultimately producing classification rules with improved accuracy. Finally, we propose two alternative denoising methods with similar empirical performance that are sometimes easier to implement. In short, the improved recovery of the geometric structure of the underlying point cloud facilitated by (local) regularization translates into improved graph-based data analysis, and the results seem to be robust to the choice of methodology.

### 6.1.1 Framework

We assume a data model

$$y_i = x_i + z_i, \tag{6.1}$$

where the *unobserved* points  $x_i$  are sampled from an unknown  $m$ -dimensional manifold  $\mathcal{M}$ , the vectors  $z_i \in \mathbb{R}^d$  represent noise, and  $\mathcal{Y}_n = \{y_1, \dots, y_n\} \subseteq \mathbb{R}^d$  is the observed data. Further geometric and probabilistic structure will be imposed to prove our main results—see Assumptions 6.1.1 and 6.1.2 below. Our analysis is motivated by the case, often found in applications, where the number  $n$  of data points and the ambient space dimension  $d$  are large, but the underlying intrinsic dimension  $m$  is small or moderate. Thus, the data-generating mechanism is described (up to a noisy perturbation) by  $m \ll d$  degrees of freedom. We aim to uncover geometric properties of the underlying manifold  $\mathcal{M}$  from the observed data  $\mathcal{Y}_n$  by using *similarity* graphs. The set of vertices of these graphs will be identified with the set  $[n] := \{1, \dots, n\}$ —so that the  $i$ -th node corresponds to the  $i$ -th data point—and the weight  $W(i, j)$  between the  $i$ -th and  $j$ -th data-point will be defined in terms of a similarity function  $\delta : [n] \times [n] \rightarrow [0, \infty)$ .

The first question that we consider is how to choose the similarity function so that  $\delta(i, j)$  approximates the hidden Euclidean distance  $\delta_{\mathcal{X}_n}(i, j) := |x_i - x_j|$ . Full knowledge of the Euclidean distance between the latent variables  $x_i$  would allow to recover, in the large  $n$  limit, global geometric features of the underlying manifold. This motivates the idea of denoising the observed point cloud  $\mathcal{Y}_n$  to approximate the hidden similarity function  $\delta_{\mathcal{X}_n}$ . Here we will study a family of similarity functions based on the Euclidean distance between local averages of points in  $\mathcal{Y}_n$ , i.e. averages of the local measures. We define a denoised dataset  $\bar{\mathcal{Y}}_n = \{\bar{y}_1, \dots, \bar{y}_n\}$  by locally averaging the original dataset, and we then define an associated similarity function

$$\delta_{\bar{\mathcal{Y}}_n}(i, j) := |\bar{y}_i - \bar{y}_j|.$$

In its simplest form,  $\bar{y}_i$  is defined by averaging all points in  $\mathcal{Y}_n$  that are inside the ball of radius  $r > 0$  centered around  $y_i$ , that is,

$$\bar{y}_i := \frac{1}{N_i} \sum_{j \in \mathcal{A}_i} y_j, \tag{6.2}$$

where  $N_i$  is the cardinality of  $\mathcal{A}_i := \{j \in [n] : y_j \in B(y_i, r)\}$ . As discussed in Section 6.1.3, this corresponds to one step of the mean-shift algorithm [Fukunaga and Hostetler, 1975]. Note that  $\bar{\mathcal{Y}}_n$  (and the associated similarity function  $\delta_{\bar{\mathcal{Y}}_n}$ ) depends on  $r$ , but we do not include said dependence in our notation for simplicity. Other possible local and non-local averaging approaches may be considered. We will only analyze the choice made in (6.2) and we will explore other constructions numerically. Introducing the notation

$$\delta_{\mathcal{X}_n}(i, j) = |x_i - x_j|, \quad \delta_{\mathcal{Y}_n}(i, j) = |y_i - y_j|,$$

the first question that we study may be formalized as understanding when, and to what extent, the similarity function  $\delta_{\bar{\mathcal{Y}}_n}$  is a better approximation than  $\delta_{\mathcal{Y}_n}$  (the standard choice) to the hidden similarity function  $\delta_{\mathcal{X}_n}$ . An answer is given in Theorem 6.1.4 below.

The second question that we investigate is how an improvement in the approximation of the hidden similarity function affects the approximation of the Laplace Beltrami operator on the underlying manifold  $\mathcal{M}$ . Specifically, we study how the spectral convergence of graph-Laplacians constructed with noisy data may be improved by local regularization of the point cloud. For concreteness, our theoretical analysis is focused on  $\varepsilon$ -graphs and unnormalized graph-Laplacians, but we expect our results to generalize to other graphs and graph-Laplacians—evidence to support this claim will be given through numerical experiments. We now summarize the necessary background to formalize this question. For a given similarity  $\delta : [n] \times [n] \rightarrow [0, \infty)$  and a parameter  $\varepsilon > 0$ , we define a weighted graph  $\Gamma_{\delta, \varepsilon} = ([n], W)$  by setting the weight between the  $i$ -th and  $j$ -th node to be

$$W(i, j) := \frac{2(m+2)}{\alpha_m \varepsilon^{m+2n}} \mathbf{1}\{\delta(i, j) < \varepsilon\}, \quad (6.3)$$

where  $\alpha_m$  is the volume of the  $m$ -dimensional Euclidean unit ball. Associated to the graph

$\Gamma_{\delta,\varepsilon}$  we define the unnormalized graph Laplacian matrix

$$\Delta_{\delta,\varepsilon} := D - W \in \mathbb{R}^{n \times n}, \quad (6.4)$$

where  $D$  is a diagonal matrix with diagonal entries

$$D(i, i) := \sum_{j=1}^n W(i, j).$$

The motivation for the scaling in (6.3) is so that  $\Delta_{\delta,\varepsilon}$  matches the scale of the Laplace-Beltrami operator (see for example Burago et al. [2015]). For the rest of the chapter we shall denote  $\Gamma_{\mathcal{X}_n,\varepsilon} := \Gamma_{\delta_{\mathcal{X}_n,\varepsilon}}$  and  $\Delta_{\mathcal{X}_n,\varepsilon} := \Delta_{\delta_{\mathcal{X}_n,\varepsilon}}$ . We use analogous notation for  $\mathcal{Y}_n$  and  $\bar{\mathcal{Y}}_n$ . The second question that we consider may be formalized as understanding when, and to what extent,  $\Delta_{\bar{\mathcal{Y}}_n}$  provides a better approximation (in the spectral sense) than  $\Delta_{\mathcal{Y}_n}$  to a Laplace operator on the manifold  $\mathcal{M}$ . An answer is given in Theorem 6.1.6 below.

### 6.1.2 Main results

In this subsection we state our main theoretical results. We first impose some geometric conditions on the underlying manifold  $\mathcal{M}$ .

**Assumption 6.1.1.**  *$\mathcal{M}$  is a smooth, oriented, compact manifold with no boundary and intrinsic dimension  $m$ , embedded in  $\mathbb{R}^d$ . Moreover,  $\mathcal{M}$  has injectivity radius  $\geq i_0$ , maximum of the absolute value of sectional curvature  $\leq K$ , and reach  $\geq R$ . Finally, we assume that  $\mathcal{M}$ 's total volume is normalized and equal to one.*

Loosely speaking, the injectivity radius determines the range of the exponential map (which will be an important tool in our analysis and will be reviewed in the next section) and the sectional curvature controls the metric distortion induced by the exponential map, and thereby its Jacobian. The reach  $R$  can be thought of as an (inverse) conditioning number

of the manifold and controls its second fundamental form; it can also be interpreted as a measure of extrinsic curvature—see, e.g. Aamari et al. [2019], Federer [1959] for technical background. The significance of these geometric quantities and their role in our analysis will be further discussed in Section 6.2.

Next we impose further probabilistic structure into the data model (6.1). We assume that the pairs  $(x_i, z_i)$  are i.i.d. samples of the random vector  $(X, Z) \sim \boldsymbol{\mu} \in \mathcal{P}(\mathcal{M} \times \mathbb{R}^d)$ . Let  $\mu$  and  $\mu_x$  be, respectively, the marginal distribution of  $X$  and the conditional distribution of  $Z$  given  $X = x$ . We assume that  $\mu$  is absolutely continuous with respect to the Riemannian volume form of  $\mathcal{M}$  with density  $p(x)$ , i.e.,

$$d\mu(x) = p(x)d\text{vol}_{\mathcal{M}}(x). \quad (6.5)$$

Furthermore, we assume that  $\mu_x$  is supported on  $T_x\mathcal{M}^\perp$  (the orthogonal complement of the tangent space  $T_x\mathcal{M}$ ) and that it is absolutely continuous with respect to the  $(d - m)$ -dimensional Hausdorff measure  $\mathcal{H}^{d-m}$  restricted to  $T_x\mathcal{M}^\perp$  with density  $p(z|x)$ , i.e.,

$$d\mu_x(z) = p(z|x)d\mathcal{H}^{d-m}(z).$$

To ease the notation we will write  $dz$  instead of  $\mathcal{H}^{d-m}(dz)$ . We make the following assumptions on these densities.

**Assumption 6.1.2.** *It holds that:*

(i) *The density  $p(x)$  is of class  $C^2(\mathcal{M})$  and is bounded above and below by positive constants:*

$$0 < p_{min} \leq p(x) \leq p_{max}, \quad \forall x \in \mathcal{M}.$$

(ii) *For all  $x \in \mathcal{M}$ ,*

$$\int zp(z|x)dz = 0.$$



Moreover, there is  $\sigma < R$  such that  $p(z|x) = 0$  for all  $z$  with  $|z| \geq \sigma$ .

Note that the assumption on  $p(z|x)$  ensures that the noise is centered and bounded by a constant  $\sigma$ . While the assumption that the noise is bounded and orthogonal to the manifold can be relaxed, we choose not to do so here to streamline our results and proofs.

In our first main theorem we study the approximation of the similarity function  $\delta_{\mathcal{X}_n}$  by  $\delta_{\bar{\mathcal{Y}}_n}$ . We consider points  $x_i$  and  $x_j$  that are close with respect to the geodesic distance  $d_{\mathcal{M}}$  on the manifold, and show that local regularization improves the approximation of the hidden similarity provided that  $n$  is large and the noise level  $\sigma$  is small. The local regularity parameter  $r$  needs to be suitably scaled with  $\sigma$ . We make the following standing assumption linking both parameters; we refer to Remark 6.1.5 below for a discussion on the optimal scaling of  $r$  with  $\sigma$ , and to our numerical experiments for practical guidelines.

**Assumption 6.1.3.** *The localization parameter  $r$  and the noise level  $\sigma$  satisfy*

$$\sigma \leq \frac{R}{16m}, \quad r \leq \min \left\{ i_0, \frac{1}{\sqrt{K}}, \sqrt{\frac{\alpha_m}{2CmK}}, \sqrt{\frac{R}{32}} \right\}, \quad \text{and } \sigma \leq \frac{1}{3}r, \quad (6.6)$$

where  $C$  is a universal constant,  $\alpha_m$  denotes the volume of the Euclidean unit ball in  $\mathbb{R}^m$ , and  $i_0$ ,  $R$ , and  $K$  are as in Assumption 6.1.1.

In words, Assumption 6.1.3 requires both  $r$  and  $\sigma$  to be sufficiently small, and  $r$  to be larger than  $\sigma$ .

Now we are ready to state the first main result.

**Theorem 6.1.4.** *Under Assumptions 6.1.1, 6.1.2 and 6.1.3, with probability at least  $1 - 4ne^{-cnr^{\max\{2m, m+4\}}}$ , for all  $x_i$  and  $x_j$  with  $d_{\mathcal{M}}(x_i, x_j) \leq r$  we have*

$$|\delta_{\mathcal{X}_n}(i, j) - \delta_{\bar{\mathcal{Y}}_n}(i, j)| \leq C_{\mathcal{M}} \left( r^3 + r\sigma + \frac{\sigma^2}{r} \right), \quad (6.7)$$

where  $c = \min \left\{ \frac{\alpha_m^2 p_{\min}^2}{4^{m+2}}, \frac{1}{16} \right\}$  and  $C_{\mathcal{M}}$  is a constant depending on  $m, K, R$ , a uniform bound on the change in second fundamental form of  $\mathcal{M}$ , and on the regularity of the density  $p$ .

**Remark 6.1.5.** *Theorem 6.1.4 gives concrete evidence of the importance of the choice of similarity function. For the usual Euclidean distance between observed data,  $\delta_{\mathcal{Y}_n}$ , one can only guarantee that*

$$|\delta_{\mathcal{X}_n}(i, j) - \delta_{\mathcal{Y}_n}(i, j)| \leq 2\sigma,$$

which follows from

$$||x_i - x_j| - |y_i - y_j|| \leq |z_i - z_j| \leq 2\sigma.$$

However, if we choose  $r \propto \sigma^{1/2}$ , then the error in (6.7) is of order  $\sigma^{3/2}$ , which is a considerably smaller quantity in the small noise limit.

Our second main result translates the local similarity bound from Theorem 6.1.4 into a global geometric result concerning the spectral convergence of the graph Laplacian to the Laplace operator formally defined by

$$\Delta_{\mathcal{M}} f = -\frac{1}{p} \operatorname{div}(p^2 \nabla f), \tag{6.8}$$

where  $\operatorname{div}$  and  $\nabla$  denote the divergence and gradient operators on the manifold and  $p$  is the sampling density of the hidden point cloud  $\mathcal{X}_n$ , as introduced in Equation (6.5). It is intuitively clear that the spectral approximation of the discrete graph-Laplacian to the continuum operator  $\Delta_{\mathcal{M}}$  necessarily rests upon having a sufficient number of samples from  $\mu$  (defined in (6.5)). In other words, the empirical measure  $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$  needs to be close to  $\mu$ , the sampling density of the hidden dataset. We characterize the closeness between  $\mu_n$  and  $\mu$  by the  $\infty$ -OT transport distance, defined as

$$d_{\infty}(\mu_n, \mu) := \min_{T: T_{\#}\mu = \mu_n} \operatorname{esssup}_{x \in \mathcal{M}} d_{\mathcal{M}}(x, T(x)),$$

where  $T_{\#}\mu$  denotes the push-forward of  $\mu$  by  $T$ , that is,  $T_{\#}\mu = \mu(T^{-1}(U))$  for any Borel subset  $U$  of  $\mathcal{M}$ . Theorem 2 in García Trillos et al. [2019a] shows that for every  $\beta > 1$ , with probability at least  $1 - C_{\beta, \mathcal{M}}n^{-\beta}$ ,

$$d_{\infty}(\mu_n, \mu) \leq C_{\mathcal{M}} \frac{\log(n)^{p_m}}{n^{1/m}},$$

where  $p_m = 3/4$  if  $m = 2$  and  $p_m = 1/m$  for  $m \geq 3$ . This is the high probability scaling of  $d_{\infty}(\mu_n, \mu)$  in terms of  $n$ .

We introduce some notation before stating our second main result. Let  $\lambda_{\ell}(\Gamma_{\delta, \varepsilon})$  be the  $\ell$ -th smallest eigenvalue of the unnormalized graph-Laplacian  $\Delta_{\delta, \varepsilon}$  defined in Equation (6.4), and let  $\lambda_{\ell}(\mathcal{M})$  be the  $\ell$ -th smallest eigenvalue of the continuum Laplace operator defined in Equation (6.8).

**Theorem 6.1.6.** *Suppose that Assumptions 6.1.1, 6.1.2, and 6.1.3 hold. Suppose further that  $\varepsilon$  is small enough (but not too small) so that*

$$\begin{aligned} \max \left\{ (m+5)d_{\infty}(\mu_n, \mu), 2Cm\eta \right\} < \varepsilon < \min \left\{ 1, \frac{i_0}{10}, \frac{1}{\sqrt{mK}}, \frac{R}{\sqrt{27m}} \right\}, \\ (\sqrt{\lambda_{\ell}(\mathcal{M})} + 1)\varepsilon + \frac{d_{\infty}(\mu_n, \mu)}{\varepsilon} < \tilde{c}_p, \end{aligned} \quad (6.9)$$

where  $\tilde{c}_p$  is a constant that only depends on  $m$  and the regularity of the density  $p$ ,  $C$  is a universal constant, and

$$\eta = C_{\mathcal{M}} \left( r^3 + r\sigma + \frac{\sigma^2}{r} \right)$$

is the bound in (6.7). Then, with probability at least  $1 - 4ne^{-cnr^{\max\{2m, m+4\}}}$ , for all  $\ell = 1, 2, 3, \dots$ ,

$$\frac{|\lambda_{\ell}(\Gamma_{\bar{\mathcal{Y}}_n, \varepsilon}) - \lambda_{\ell}(\mathcal{M})|}{\lambda_{\ell}(\mathcal{M})} \leq \tilde{C} \left( \frac{\eta}{\varepsilon} + \frac{d_{\infty}(\mu_n, \mu)}{\varepsilon} + (1 + \sqrt{\lambda_{\ell}(\mathcal{M})})\varepsilon + \left( K + \frac{1}{R^2} \right) \varepsilon^2 \right),$$

where  $\tilde{C}$  only depends on  $m$  and the regularity of  $p$ , and  $c = \min \left\{ \frac{\alpha_m^2 p_{\min}^2}{4^{m+2}}, \frac{1}{16} \right\}$ .

**Remark 6.1.7.** We will see in Section 6.3 that Theorem 6.1.6 follows by plugging the probabilistic estimate (6.7) into a modification of a deterministic result from [García Trillos et al., 2019a, Corollary 2], which we present for the convenience of the reader in Theorem 6.3.3. We remark that any improvement of Theorem 6.3.3 would immediately translate into an improvement of our Theorem 6.1.6. As discussed in Remark 6.1.5, local regularization enables a smaller  $\eta$  than if no regularization is performed. This in turn allows one to choose, for a given error tolerance, a smaller connectivity  $\varepsilon$ , leading to a sparser graph that is computationally more efficient. Note also that the bound in Theorem 6.1.6 does not depend on the ambient space dimension  $d$ , but only on the intrinsic dimension  $m$  of the data.

**Remark 6.1.8.** Theorem 6.1.6 concretely shows how an improvement in metric approximation translates into an improved estimation of global geometric quantities. We have restricted our attention to analyzing eigenvalues of a Laplacian operator, but we remark that the idea goes beyond this particular choice. For example, one can conduct an asymptotic analysis illustrating the effect of changing the similarity function in the approximation of other geometric quantities of interest like Cheeger cuts. Such analysis could be carried out using the variational convergence approach from García Trillos and Slepčev [2016a]. Finally, we remark that it is possible to study convergence of eigenvectors of graph Laplacians following the results in García Trillos et al. [2019a].

### 6.1.3 Related and Future Work

Graph-based learning algorithms include spectral clustering, total variation clustering, graph-Laplacian regularization for semi-supervised learning, graph based Bayesian semi-supervised learning. A brief and incomplete summary of methodological and review papers is Shi and Malik [2000], Ng et al. [2002], Belkin and Niyogi [2004], Zhou and Schölkopf [2005], Spielman

and Teng [2007], Von Luxburg [2007], Zhu [2005], Bertozzi et al. [2018]. These algorithms involve either a graph Laplacian, the graph total variation, or Sobolev norms involving the graph structure. The large sample  $n \rightarrow \infty$  theory studying the behavior of some of the above methodologies has been analyzed without reference to the intrinsic dimension of the data [Von Luxburg et al., 2008] and in the case of points laying *on* a low dimensional manifold, see e.g. Belkin et al. [2006], García Trillos and Sanz-Alonso [2018], García Trillos et al. [2020b] and references therein. Some papers that account for *both* the noisy and low intrinsic dimensional structure of data are Niyogi et al. [2008], Little and Maggioni [2017], Agapiou et al. [2017], Weed and Bach [2019], Genovese et al. [2012], Aamari and Levrard [2019]. For example, Niyogi et al. [2008] studies the recovery of the homology groups of submanifolds from noisy samples. We use the techniques for the analysis of spectral convergence of graph-Laplacians introduced in Burago et al. [2015] and further developed in García Trillos et al. [2019a]. The results in the latter reference would allow to extend our analysis to other graph Laplacians, but we do not pursue this here for conciseness.

We highlight that the denoising by local regularization occurs at the level of the dataset. That is, rather than denoising each of the observed features individually, we analyze denoising by averaging different data points. In practice combining both forms of denoising may be advantageous. For instance, when each of the data points corresponds to an image, one can first denoise each image at the pixel level and then do regularization at the level of the dataset as proposed here. In this regard, our regularization at the level of the data-set is similar to applying a filter at the level of individual pixels [Tukey and Tukey, 1988]. The success of non-local filter image denoising algorithms suggests that non-local methods may be also of interest at the level of the dataset, but we expect this to be application-dependent. Finally, while in this chapter we only consider first-order regularization based on averages, a topic for further research is the analysis of local PCA regularization [Little and Maggioni, 2017], incorporating covariance information.

With the same motivation for our work, in Memoli et al. [2019] a general construction of metrics on noisy datasets was proposed. The so called *Wasserstein transform* associates to each of the data points a “local” probability distribution, and defines a new metric on the data by computing the Wasserstein distance between the corresponding local measures. A particular construction of local measures closely related to the metric we study here assigns to each observation the empirical measure of the observations restricted to a ball of certain radius around the given data point. The authors of Memoli et al. [2019] propose the Wasserstein transform as a way to generalize the *mean-shift algorithm* and they study how it alleviates the so called *chaining effect* in single linkage clustering. The aim of our work is to provide quantitative evidence of the effect that changing the metric on noisy datasets has on graph-based spectral clustering algorithms. The success of these algorithms hinges on their ability to capture the geometry of the underlying data generating model.

It is worth noting the parallel between the local regularization that we study here and mean-shift and mode seeking methods [Chen et al., 2016, Fukunaga and Hostetler, 1975]. As a matter of fact the points  $\bar{x}_i$  that we construct here correspond to one step in the standard mean shift algorithm. However, we notice that our goal is not to run mean shift for mode seeking, but rather, as a way to construct a metric that better captures the underlying “true” geometric structure of the data that was blurred by noise. This parallelism with mean-shift techniques (or the more general Wasserstein transform in Memoli et al. [2019]) suggests the idea of doing local averaging iteratively. Of course, it is important to notice that unless one prevents points to move tangentially to  $\mathcal{M}$  (as discussed in Wang and Carreira-Perpinán [2010]), a large number of iterations would result in points collapsing to a finite number of local modes.

Local regularization may be also interpreted as a form of dictionary learning, where each data-point is represented in terms of its neighbors. For specific applications it may be of interest to restrict (or extend) the dictionary used to represent each data point [Haddad

et al., 2014]. Finally we refer to Hein and Maier [2006] for alternative techniques on manifold denoising.

### 6.1.4 Outline

The chapter is organized as follows. In Section 6.2 we formalize the geometric setup and prove Theorem 6.1.4. Section 6.3 contains the proof of Theorem 6.1.6 and a lemma that may be of independent interest. Finally, Section 6.4 includes several numerical experiments. In the Appendix we prove a technical lemma that serves as a key ingredient in proving Theorem 6.1.4.

## 6.2 Distance Approximation

In this section we prove Theorem 6.1.4. We start with Section 6.2.1 by giving some intuition on the geometric conditions imposed in Assumption 6.1.1 and introducing the main geometric tools in our analysis. In Section 6.2.2 we decompose the approximation error between the similarity functions  $\delta_{\bar{y}_n}$  and  $\delta_{\mathcal{X}_n}$  into three terms, which are bounded in Sections 6.2.3, 6.2.4, and 6.2.5.

### 6.2.1 Geometric Preliminaries

In this subsection we set our notation and provide some background on geometric concepts used in the remainder of this chapter.

#### 6.2.1.1 Basic Notation

For each  $x \in \mathcal{M}$  we let  $T_x\mathcal{M}$  be the tangent plane of  $\mathcal{M}$  at  $x$  centered at the origin. In particular,  $T_x\mathcal{M}$  is a  $m$ -dimensional subspace of  $\mathbb{R}^d$ , and we denote by  $T_x\mathcal{M}^\perp$  its orthogonal complement. We will use  $\text{vol}_{\mathcal{M}}$  to denote the Riemannian volume form of  $\mathcal{M}$ . We will denote

by  $|x - \tilde{x}|$  the Euclidean distance between arbitrary points in  $\mathbb{R}^d$  and denote by  $d_{\mathcal{M}}(x, \tilde{x})$  the geodesic distance between points in  $\mathcal{M}$ . We denote by  $B_x$  balls in  $T_x\mathcal{M}$  and by  $B_{\mathcal{M}}$  balls in the manifold  $\mathcal{M}$  (with respect to the geodesic distance). Also, unless otherwise specified  $B$ , without subscripts will be used to denote balls in  $\mathbb{R}^d$ . We denote by  $\alpha_m$  the volume of the unit Euclidean ball in  $\mathbb{R}^m$ . Throughout the rest of the chapter we use  $R, i_0$  and  $K$  to denote the reach, injectivity radius, and maximum absolute curvature of  $\mathcal{M}$ , as in Assumption 6.1.1. We now describe at an intuitive level the role that these quantities play in our analysis.

### 6.2.1.2 The Reach

The reach of a closed submanifold  $\mathcal{M}$  is the largest value  $t \in [0, \infty]$  such that the projection map onto  $\mathcal{M}$  is well defined on  $\{x \in \mathbb{R}^d : \inf_{\tilde{x} \in \mathcal{M}} |\tilde{x} - x| < t\}$ , i.e., every point in the tubular neighborhood around  $\mathcal{M}$  of width  $t$  has a unique closest point in  $\mathcal{M}$ . Our assumption that the noise level satisfies  $\sigma < R$  guarantees that  $x_i$  is the (well-defined) projection of  $y_i$  onto the manifold. The reach can be thought of as an inverse conditioning number for the manifold [Niyogi et al., 2008]. We will use that the inverse of the reach provides a uniform upper bound on the second fundamental form (see Lemma 6.2.7).

### 6.2.1.3 Exponential Map, Injectivity Radius and Sectional Curvature

We will make use of the exponential map  $\exp$ , which for every  $x \in \mathcal{M}$  is a map

$$\exp_x : B_x(0, i_0) \rightarrow B_{\mathcal{M}}(x, i_0)$$

where  $i_0$  is the injectivity radius for the manifold  $\mathcal{M}$ . We recall that the exponential map  $\exp_x$  takes a vector  $v \in T_x\mathcal{M}$  and maps it to the point  $\exp_x(v) \in \mathcal{M}$  that is at geodesic distance  $|v|$  from  $x$  along the unit speed geodesic that at time  $t = 0$  passes through  $x$  with velocity  $v/|v|$ . The injectivity radius  $i_0$  is precisely the maximum radius of a ball in  $T_x\mathcal{M}$



centered at the origin for which the exponential map is a well defined diffeomorphism for every  $x$ . We denote by  $J_x$  the Jacobian of the exponential map  $\exp_x$ . Integrals with respect to  $d\text{vol}_{\mathcal{M}}$  can then be written in terms of integrals on  $T_x\mathcal{M}$  weighted by the function  $J_x$ . More precisely, for an arbitrary test function  $\varphi : \mathcal{M} \rightarrow \mathbb{R}$ ,

$$\int_{B_{\mathcal{M}}(x, i_0)} \varphi(\tilde{x}) d\text{vol}_{\mathcal{M}}(\tilde{x}) = \int_{B_x(0, i_0)} \varphi(\exp_x(v)) J_x(v) dv.$$

For fixed  $0 < r \leq \min\{i_0, 1/\sqrt{K}\}$  one can obtain bounds on the metric distortion by the exponential map  $\exp_x : B_x(0, r) \subseteq T_x\mathcal{M} \rightarrow \mathcal{M}$  ([Do Carmo and Flaherty Francis, 1992, Chapter 10] and [Burago et al., 2015, Section 2.2]), and thereby guarantee the existence of a universal constant  $C$  such that, for  $|v| \leq r$ ,

$$(1 + CmK|v|^2)^{-1} \leq J_x(v) \leq (1 + CmK|v|^2). \quad (6.10)$$

An immediate consequence of the previous inequalities is

$$|\text{vol}(B_{\mathcal{M}}(x, r)) - \alpha_m r^m| \leq CmK r^{m+2}, \quad (6.11)$$

where we recall  $\alpha_m$  is the volume of the unit ball in  $\mathbb{R}^m$ . Equations (6.10) and (6.11) will be used in our geometric and probabilistic arguments and motivate our assumptions on the choice of local regularization parameter  $r$  in terms of the injectivity radius and the sectional curvature.

### 6.2.2 Local Distributions

Next we study the local behavior of  $(X, Z)$ . To characterize its local distribution, it will be convenient to introduce the following family of probability measures.

**Definition 6.2.1.** *Let  $y$  be a vector in  $\mathbb{R}^d$  whose distance to  $\mathcal{M}$  is less than  $R$ . Let  $x$  be*

the projection of  $y$  onto  $\mathcal{M}$ . We say that the random variable  $(\tilde{X}, \tilde{Z})$  has the distribution  $\boldsymbol{\mu}_y$  provided that

$$\mathbb{P}((\tilde{X}, \tilde{Z}) \in A_1 \times A_2) := \mathbb{P}((X, Z) \in A_1 \times A_2 | X + Z \in B(y, r)),$$

for all Borel sets  $A_1 \subseteq \mathcal{M}$   $A_2 \subseteq \mathbb{R}^d$ , where in the above  $(X, Z)$  is distributed according to  $\boldsymbol{\mu}$ .

In the remainder we use  $\boldsymbol{\mu}_i$  as shorthand notation for  $\boldsymbol{\mu}_{y_i}$ . As for the original measure  $\boldsymbol{\mu}$ , we characterize  $\boldsymbol{\mu}_i$  in terms of a marginal and conditional distribution. We introduce the density  $\tilde{p}_i : \mathcal{M} \rightarrow \mathbb{R}$  given by

$$\tilde{p}_i(x) := \frac{\mathbb{P}_i(X + Z \in B(y_i, r) | X = x)}{\mathbb{P}_i(X + Z \in B(y_i, r))} \cdot p(x), \quad (6.12)$$

and define

$$\tilde{p}_i(z|x) = \frac{\mathbf{1}_{x+z \in B(y_i, r)}}{\mathbb{P}_i(X + Z \in B(y_i, r) | X = x)} \cdot p(z|x), \quad (6.13)$$

where in the above and in the remainder we use  $\mathbb{E}_i$  and  $\mathbb{P}_i$  to denote conditional expectation and conditional probability given  $(x_i, z_i)$ . It can be easily shown that these functions correspond to the marginal density of  $\tilde{X}_i$  and the conditional density of  $\tilde{Z}_i$  given  $\tilde{X}_i = x$ , where  $(\tilde{X}_i, \tilde{Z}_i) \sim \boldsymbol{\mu}_i$ . The distribution  $\boldsymbol{\mu}_i$  is of relevance because by definition of  $\bar{y}_i$  one has

$$\mathbb{E}_i[\bar{y}_i] = \mathbb{E}_i[\tilde{X}_i + \tilde{Z}_i].$$

Now we are ready to introduce the main decomposition of the error between the similarity

functions  $\delta_{\bar{y}_n}$  and  $\delta_{\mathcal{X}_n}$ . Using the triangle inequality we can write

$$|x_i - x_j| - |\bar{y}_i - \bar{y}_j| \leq |\mathbb{E}_i[\tilde{X}_i] - x_i - (\mathbb{E}_j[\tilde{X}_j] - x_j)| \quad (6.14)$$

$$+ |\mathbb{E}_j[\tilde{Z}_j]| + |\mathbb{E}_i[\tilde{Z}_i]| \quad (6.15)$$

$$+ |\mathbb{E}_i[\bar{y}_i] - \bar{y}_i| + |\mathbb{E}_j[\bar{y}_j] - \bar{y}_j|. \quad (6.16)$$

In the next subsections we bound each of the terms (6.15) (expected conditional noise), (6.14) (difference in *geometric* bias), and (6.16) (sampling error). As we will see in Section 6.2.5 we can control both terms in (6.16) with very high probability using standard concentration inequalities. The other three terms are deterministic quantities that can be written in terms of integrals with respect to the distributions  $\tilde{\mu}_i$  and  $\tilde{\mu}_j$ . To study these integrals it will be convenient to introduce two quantities  $r_- < r < r_+$  (independent of  $i = 1, \dots, n$ ) satisfying:

1. For all  $x \in \mathcal{M}$  with  $d_{\mathcal{M}}(x, x_i) > r_+$  we have

$$\mathbb{P}_i(X + Z \in B(y_i, r) | X = x) = 0.$$

Equivalently, the density  $\tilde{p}_i(x)$  is supported in  $\overline{B_{\mathcal{M}}(x_i, r_+)}$ .

2. For all  $x$  with  $d_{\mathcal{M}}(x, x_i) < r_-$  we have

$$\mathbb{P}_i(X + Z \in B(y_i, r) | X = x) = 1.$$

It should be noted that the choice of both  $r_-$  and  $r_+$  depends on  $r$ . In Appendix 6.5 we present the proof of the following lemma giving estimates for  $r_+$  and  $r_-$ .

**Lemma 6.2.2** (Bounds for  $r_+$  and  $r_-$ ). *Under Assumption 6.1.3, the quantities*

$$r_- := r \left( \sqrt{1 + \frac{4\sigma}{R} + \frac{16\sigma^2}{r^2} + \frac{m\sigma}{R}} \right)^{-1},$$

$$r_+ := r \left( \sqrt{1 - \frac{8r^2}{R} - \frac{4\sigma}{R} - \frac{m\sigma}{R}} \right)^{-1},$$

satisfy properties i) and ii). Furthermore,

$$r_+ - r_- \leq C_{m,R} \left( r^3 + r\sigma + \frac{\sigma^2}{r} \right), \quad C_{m,R} := \max \left\{ \frac{8m + 32}{R}, 64 \right\}$$

and

$$\frac{1}{2}r_+ \leq r \leq 2r_-. \quad (6.17)$$

### 6.2.3 Bounding Expected Conditional Noise

**Proposition 6.2.3.** *Suppose that Assumptions 6.1.1 and 6.1.2 hold. Then,*

$$|\mathbb{E}_i[\tilde{Z}_i]| \leq C_{m,p} \frac{\sigma}{r} (r_+ - r_-), \quad C_{m,p} := \frac{4^{m+1} p_{max}}{m p_{min}}.$$

*Proof.* Using the definition of  $r_+$ ,

$$\begin{aligned} \mathbb{E}_i[\tilde{Z}_i] &= \int_{B_{\mathcal{M}}(x_i, r_+)} \int z \tilde{p}_i(z|x) dz \tilde{p}_i(x) d\text{vol}_{\mathcal{M}}(x) \\ &= \int_{B_{\mathcal{M}}(x_i, r_-)} \int z \tilde{p}_i(z|x) dz \tilde{p}_i(x) d\text{vol}_{\mathcal{M}}(x) \\ &\quad + \int_{B_{\mathcal{M}}(x_i, r_+) \setminus B_{\mathcal{M}}(x_i, r_-)} \int z \tilde{p}_i(z|x) dz \tilde{p}_i(x) d\text{vol}_{\mathcal{M}}(x). \end{aligned}$$

The first integral is the zero vector because for  $x \in B_{\mathcal{M}}(x_i, r_-)$ , we have  $\tilde{p}(z|x) \propto p(z|x)$

and  $p(z|x)$  is assumed to be centered. Therefore,

$$\begin{aligned}
|\mathbb{E}_i[\tilde{Z}_i]| &\leq \sigma \int_{B_{\mathcal{M}}(x_i, r_+) \setminus B_{\mathcal{M}}(x_i, r_-)} \tilde{p}_i(x) d\text{vol}_{\mathcal{M}}(x) \\
&= \frac{\sigma}{\mathbb{P}_i(X + Z \in B(y_i, r))} \int_{B_{\mathcal{M}}(x_i, r_+) \setminus B_{\mathcal{M}}(x_i, r_-)} p(x) d\text{vol}_{\mathcal{M}}(x) \\
&\leq \frac{\sigma p_{\max}}{\mathbb{P}_i(X + Z \in B(y_i, r))} \int_{B_{\mathcal{M}}(x_i, r_+) \setminus B_{\mathcal{M}}(x_i, r_-)} d\text{vol}_{\mathcal{M}}(x) \\
&\leq \frac{\sigma p_{\max}}{\mathbb{P}_i(X + Z \in B(y_i, r))} \int_{B_{x_i}(0, r_+) \setminus B_{x_i}(0, r_-)} J_{x_i}(v) dv \\
&\leq \frac{2\alpha_m \sigma p_{\max}}{\mathbb{P}_i(X + Z \in B(y_i, r))} (r_+^m - r_-^m) \\
&\leq \frac{2\alpha_m \sigma p_{\max}}{m \mathbb{P}_i(X + Z \in B(y_i, r))} (r_+ - r_-) r_+^{m-1},
\end{aligned}$$

where we have used (6.10) and the assumptions on  $r$  to say (in particular) that  $J_{x_i}(v) \leq 2$ , and also the fact that, for  $t > s > 0$ ,

$$t^m - s^m = \int_s^t \frac{u^{m-1}}{m} du \leq (t - s) \frac{t^{m-1}}{m}.$$

Finally, notice that

$$\mathbb{P}_i(X + Z \in B(y_i, r)) \geq \mathbb{P}_i(X \in B_{\mathcal{M}}(x_i, r_-)) = \int_{B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) J_{x_i}(v) dv \geq \frac{1}{2} p_{\min} \alpha_m r_-^m,$$

where again we have used (6.10) to conclude (in particular) that  $J_{x_i}(v) \geq 1/2$ . The result now follows by (6.17). □

### 6.2.4 Bounding Difference in Geometric Bias

In terms of  $r_+$  and  $r_-$ , the difference  $\mathbb{E}_i[\tilde{X}_i] - x_i$  (and likewise  $\mathbb{E}_j[\tilde{X}_j] - x_j$ ) can be written as:

$$\begin{aligned}
\mathbb{E}_i[\tilde{X}_i] - x_i &= \int_{B_{\mathcal{M}}(x_i, r_+)} (x - x_i) \tilde{p}_i(x) d\text{vol}_{\mathcal{M}}(x) \\
&= \int_{B_{x_i}(0, r_+)} (\exp_{x_i}(v) - x_i) \tilde{p}_i(\exp_x(v)) J_{x_i}(v) dv \\
&= \int_{B_{x_i}(0, r_+)} (\exp_{x_i}(v) - x_i) \tilde{p}_i(\exp_x(v)) dv + \\
&\quad \int_{B_{x_i}(0, r_+)} (\exp_{x_i}(v) - x_i) \tilde{p}_i(\exp_x(v)) (J_{x_i}(v) - 1) dv \\
&= \frac{1}{\mathbb{P}_i(X + Z \in B(y_i, r))} \int_{B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) p(\exp_x(v)) dv + \\
&\quad \int_{B_{x_i}(0, r_+) \setminus B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) \tilde{p}_i(\exp_x(v)) dv + \\
&\quad \int_{B_{x_i}(0, r_+)} (\exp_{x_i}(v) - x_i) \tilde{p}_i(\exp_x(v)) (J_{x_i}(v) - 1) dv \\
&:= \frac{1}{\mathbb{P}_i(X + Z \in B(y_i, r))} \int_{B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) p(\exp_x(v)) dv + \xi_i,
\end{aligned}$$

where the second to last equality follows from (6.12). To further simplify the expression for  $x_i - \mathbb{E}_i[\tilde{X}_i]$  let us define

$$b_i := \int_{B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) p(\exp_x(v)) dv.$$

It follows that

$$\begin{aligned}
|\mathbb{E}_i[\tilde{X}_i] - x_i - (\mathbb{E}_j[\tilde{X}_j] - x_j)| &\leq \left| \frac{b_i}{P_i} - \frac{b_j}{P_j} \right| + |\xi_i| + |\xi_j| \\
&\leq \left| \frac{1}{P_i} - \frac{1}{P_j} \right| |b_i| + \frac{1}{P_j} |b_i - b_j| + |\xi_i| + |\xi_j|,
\end{aligned} \tag{6.18}$$

where in the above

$$P_i := \mathbb{P}_i(X + Z \in B(y_i, r)), \quad P_j := \mathbb{P}_j(X + Z \in B(y_j, r)).$$

**Lemma 6.2.4.** *The following hold.*

1. *The terms  $P_i$  satisfy*

$$\frac{1}{2} p_{\min} \alpha_m r_-^m \leq P_i.$$

2. *The terms  $\xi_i$  satisfy:*

$$|\xi_i| \leq C_1(r_+ - r_-) + C_2 r^3,$$

where, up to universal multiplicative constants,

$$C_1 = \frac{4^{m+1} p_{\max}}{m p_{\min}}, \quad C_2 = 4^{m+3} m K \frac{p_{\max}}{p_{\min}}.$$

3. *Suppose that  $d_{\mathcal{M}}(x_i, x_j) \leq r$ . Then,*

$$|P_i - P_j| \leq C_3 r^{m+1} + C_4 (r_+ - r_-) r^{m-1} + C_5 r^{m+2},$$

where, up to universal multiplicative constants,

$$C_3 = C_p \alpha_m, \quad C_4 = \frac{2^{m-1} \alpha_m p_{\max}}{m}, \quad C_5 = m K p_{\max} \alpha_m.$$

and  $C_p$  only depends on bounds on the first derivatives of the density  $p$ .

*Proof.* The first inequality was already obtained at the end of the proof of Proposition 6.2.3.

For the second inequality recall that

$$\begin{aligned} \xi_i &= \int_{B_{x_i}(0, r_+) \setminus B_{x_i}(0, r_-)} (x_i - \exp_{x_i}(v)) \tilde{p}_i(\exp_{x_i}(v)) dv + \\ &\quad \int_{B_{x_i}(0, r_+)} (x_i - \exp_{x_i}(v)) \tilde{p}_i(\exp_{x_i}(v)) [J_{x_i}(v) - 1] dv := I_1 + I_2. \end{aligned}$$

For the first term we notice that  $|x_i - \exp_{x_i}(v)| \leq d_{\mathcal{M}}(x_i, \exp_{x_i}(v)) \leq r_+$ . Thus using i) and the definition of  $\tilde{p}_i$  we have

$$|I_1| \leq \frac{r_+ p_{max} \alpha_m}{\mathbb{P}_i(X + Z \in B(y_i, r))} (r_+^m - r_-^m) \leq \frac{4^{m+1} p_{max}}{m p_{min}} (r_+ - r_-).$$

For the second term we use i) and (6.10) to see that

$$|I_2| \leq \frac{C m K p_{max} \alpha_m}{\mathbb{P}_i(X + Z \in B(y_i, r))} r_+^{m+3} \leq C 4^{m+3} m K \frac{p_{max}}{p_{min}} r_+^3.$$

For iii) we notice that by definition of  $r_-$  and  $r_+$  we can write

$$\mathbb{P}_i(X \in B_{x_i}(0, r_-)) - \mathbb{P}_j(X \in B_{x_j}(0, r_+)) \leq P_i - P_j \leq \mathbb{P}_i(X \in B_{x_i}(0, r_+)) - \mathbb{P}_j(X \in B_{x_j}(0, r_-)),$$

and in particular it is enough to bound  $H_{ij} := |\mathbb{P}_i(X \in B_{\mathcal{M}}(x_i, r_+)) - \mathbb{P}_j(X \in B_{\mathcal{M}}(x_j, r_-))|$ .

We can expand  $H_{ij}$  as follows.

$$\begin{aligned} H_{ij} &= \int_{B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) dv - \int_{B_{x_j}(0, r_-)} p(\exp_{x_j}(\tilde{v})) d\tilde{v} \\ &\quad + \int_{B_{x_i}(0, r_+) \setminus B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) dv \\ &\quad + \int_{B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) (J_{x_i}(v) - 1) dv - \int_{B_{x_j}(0, r_-)} p(\exp_{x_j}(\tilde{v})) (J_{x_j}(\tilde{v}) - 1) d\tilde{v} \\ &:= \mathcal{I}_1 + \mathcal{I}_2 + \mathcal{I}_3. \end{aligned}$$



By a similar argument as above, we can bound  $\mathcal{I}_2$  and  $\mathcal{I}_3$  by

$$\begin{aligned} |\mathcal{I}_2| &\leq p_{max}\alpha_m(r_+^m - r_-^m) \leq \frac{2^{m-1}}{m}\alpha_m p_{max}(r_+ - r_-)r^{m-1}, \\ |\mathcal{I}_3| &\leq 2CmKp_{max}\alpha_m r_-^{m+2} \leq 2CmKp_{max}\alpha_m r^{m+2}. \end{aligned}$$

Finally, we notice that we can identify  $B_{x_i}(0, r_-)$  with  $B_{x_j}(0, r_-)$ . From the assumed smoothness on  $p$  (which in particular is  $C^1$ ) we see that for any  $v \in B_{x_i}(0, r_-)$  we have

$$|p(\exp_{x_i}(v)) - p(\exp_{x_j}(v))| \leq C_p d_{\mathcal{M}}(\exp_{x_i}(v), \exp_{x_j}(v)) \leq 3C_p r.$$

Then it follows that  $|\mathcal{I}_1| \leq 3C_p \alpha_m r^{m+1}$  and we get the desired result.  $\square$

We now bound the difference  $|b_i - b_j|$  for nearby points  $x_i, x_j$ , where we recall that

$$b_i := \int_{B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) p(\exp_{x_i}(v)) dv.$$

**Proposition 6.2.5.** *Suppose that  $x_i$  and  $x_j$  are such that  $d_{\mathcal{M}}(x_i, x_j) \leq r$ . Then,*

$$|b_i - b_j| \leq Cr^{m+3},$$

where the constant  $C$  can be written as

$$C = p_{max}\alpha_m \left( \frac{6\sqrt{m}}{R^2} + \left(1 + \frac{4}{R}\right) C_{\mathcal{M}} \right) + \frac{C_p}{R} \alpha_m,$$

where  $C_p$  is a constant that depends on bounds on first and second derivatives of the density  $p$ , and  $C_{\mathcal{M}}$  is a constant that depends only on the change in second fundamental form along  $\mathcal{M}$  (a third order term).

As we will see Proposition 6.2.5 can be proved combining several ideas from differen-

tial geometry. We present the required auxiliary results as we develop the proof of the proposition.

We start by conveniently writing  $b_i$  and  $b_j$  in a way that facilitates their direct comparison. Indeed, for any given  $v \in B_{x_i}(0, r_-)$  let us consider the curves

$$\gamma_{v,i}(t) := \exp_{x_i} \left( t \frac{v}{|v|} \right), \quad t \in [0, |v|],$$

and

$$t \in [0, |v|] \mapsto x_i + t \frac{v}{|v|}.$$

Thus,  $\gamma_{v,i}$  is an arc-length parameterized geodesic on  $\mathcal{M}$  that starts at the point  $x_i$  and at time  $|v|$  passes through the point  $\exp_{x_i}(v)$ . Its initial velocity  $\dot{\gamma}_{v,i}(0)$  is the vector  $v/|v|$ . On the other hand, while the second curve does not stay in  $\mathcal{M}$  for  $t > 0$ , it does have the same starting point and velocity as  $\gamma_{v,i}$ . We can use the fundamental theorem of calculus to write:

$$\exp_{x_i}(v) - (x_i + v) = \int_0^{|v|} \left( \dot{\gamma}_{v,i} - \frac{v}{|v|} \right) dt,$$

as well as

$$\dot{\gamma}_{v,i}(t) - \frac{v}{|v|} = \int_0^t \ddot{\gamma}_{v,i}(s) ds, \quad \forall t \in [0, |v|]. \quad (6.19)$$

In particular, we have the second order representation

$$\exp_{x_i}(v) - x_i = v + \int_0^{|v|} \int_0^t \ddot{\gamma}_{v,i}(s) ds dt. \quad (6.20)$$

As a consequence of the previous formula we can rewrite  $b_i$  as

$$\begin{aligned} b_i &= \int_{B_{x_i}(0, r_-)} (\exp_{x_i}(v) - x_i) p(\exp_{x_i}(v)) dv \\ &= \int_{B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) \int_0^{|v|} \int_0^t \ddot{\gamma}_{v,i}(s) ds dt dv + \int_{B_{x_i}(0, r_-)} v p(\exp_{x_i}(v)) dv. \end{aligned} \quad (6.21)$$

Completely analogous definitions and statements can be introduced to represent  $b_j$ .

With the objective of using the formula (6.21) to compare  $b_i$  and  $b_j$  we relate vectors in  $T_{x_i}\mathcal{M}$  with vectors in  $T_{x_j}\mathcal{M}$  by a convenient linear isometry  $F_{ij} : T_{x_i}\mathcal{M} \mapsto T_{x_j}\mathcal{M}$  constructed using parallel transport.

**Lemma 6.2.6.** *Suppose that  $x_i$  and  $x_j$  are such that  $d_{\mathcal{M}}(x_i, x_j) \leq r$ . Let  $\phi : t \in [0, d_{\mathcal{M}}(x_i, x_j)] \mapsto \phi(t) \in \mathcal{M}$ , be the arc-length parameterized geodesic starting at  $x_i$  at time zero and passing through  $x_j$  at time  $t = d_{\mathcal{M}}(x_i, x_j)$ . For an arbitrary vector  $v \in T_{x_i}\mathcal{M}$  let  $V_v$  be the (unique) vector field along  $\phi$  that solves the ODE*

$$\begin{cases} \frac{D}{dt}V_v(t) = 0, & t \in (0, d_{\mathcal{M}}(x_i, x_j)), \\ V_v(0) = v, \end{cases},$$

where  $\frac{D}{dt}$  denotes the covariant derivative (on  $\mathcal{M}$ ) along the curve  $\phi$ . Then, the map  $F_{ij}$  defined by

$$F_{ij} : v \mapsto \tilde{v} := V_v(d_{\mathcal{M}}(x_i, x_j))$$

is a linear isometry. Moreover,

$$|v - \tilde{v}| \leq \frac{1}{R}|v|d_{\mathcal{M}}(x_i, x_j), \quad \forall v \in T_{x_i}\mathcal{M}. \quad (6.22)$$

*Proof.* First note that  $F_{ij}$  is a linear isometry since the ODE defining  $V_v$  is linear and the vector fields  $V_v$  are parallel to the curve  $\phi$  by definition. To get the estimate (6.22) we can use the fundamental theorem of calculus and write

$$\tilde{v} = v + \int_0^t \dot{V}_v(s)ds,$$

where  $t := d_{\mathcal{M}}(x_i, x_j)$ . The fact that  $V_v$  is parallel along the curve  $\phi$  implies that  $\dot{V}_v(s) \in$

$T_{\phi(s)}\mathcal{M}^\perp$  and furthermore that for arbitrary unit norm  $\eta$  with  $\eta \in T_{\phi(s)}\mathcal{M}^\perp$  we have

$$|\langle \dot{V}_v(s), \eta \rangle| = |\langle S_\eta(V_v(s)), \dot{\phi}(s) \rangle| \leq \|S_\eta\| \|V_v(s)\| |\dot{\phi}(s)| = \|S_\eta\| |v|,$$

where  $S_\eta$  is the so called *shape operator* representing the second fundamental form (see Proposition 2.3 Chapter 6 in Do Carmo and Flaherty Francis [1992]). The relevance of the previous inequality is that when combined with Proposition 6.1 in Niyogi et al. [2008] (which shows that the operator norm of the second fundamental form is bounded by  $1/R$ ) it implies that

$$|\dot{V}_v(s)| \leq \frac{|v|}{R}, \quad \forall s \in [0, t].$$

Therefore,

$$|\tilde{v} - v| \leq \int_0^t |\dot{V}_v(s)| ds \leq \frac{|v|}{R} d_{\mathcal{M}}(x_i, x_j),$$

establishing in this way the desired bound.  $\square$

From now on, for a given  $v \in B_{x_i}(0, r_-)$  we let  $\tilde{v} \in B_{x_j}(0, r_-)$  be its image under  $F_{ij}$ .

We consider the curve

$$\gamma_{\tilde{v},j}(t) := \exp_{x_j} \left( t \frac{\tilde{v}}{|\tilde{v}|} \right), \quad t \in [0, |\tilde{v}|],$$

where we recall that  $|v| = |\tilde{v}|$  because  $F_{ij}$  is a linear isometry. We can then make a change of variables and write  $b_j$  as

$$b_j = \int_{B_{x_i}(0, r_-)} p(\exp_{x_j}(\tilde{v})) \int_0^{|v|} \int_0^t \ddot{\gamma}_{\tilde{v},j}(s) ds dt dv + \int_{B_{x_i}(0, r_-)} \tilde{v} p(\exp_{x_j}(\tilde{v})) dv. \quad (6.23)$$

In the next lemma we find bounds for the norms of accelerations.

**Lemma 6.2.7.** *Let  $v \in B_{x_i}(0, r_-)$  and let  $\tilde{v}$  be as in Lemma 6.2.6. Then, for all  $t \in [0, |v|]$  we have*

$$|\ddot{\gamma}_{v,i}(t)| \leq \frac{1}{R},$$

and

$$|\dot{\gamma}_{v,i}(t) - \dot{\gamma}_{\tilde{v},j}(t)| \leq 2\frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R}.$$

*Proof.* The first inequality appears in the proof of Proposition 2 in Niyogi et al. [2008] and is obtained in a completely analogous way as we obtained the bound for  $\dot{V}_v$  in the proof of Lemma 6.2.6 (given that unit speed geodesics are auto parallel).

To prove the second estimate, we notice that from the first bound and (6.19) it follows that

$$\left| \dot{\gamma}_{v,i}(t) - \frac{v}{|v|} \right| \leq \frac{|v|}{R}, \quad \forall t \in [0, |v|].$$

Naturally, a similar inequality holds for  $\gamma_{\tilde{v},j}$ . Using Lemma 6.2.6 we conclude that for all  $t \in [0, |v|]$  (recall that  $|v| = |\tilde{v}|$ )

$$\begin{aligned} |\dot{\gamma}_{v,i}(t) - \dot{\gamma}_{\tilde{v},j}(t)| &\leq \left| \frac{v}{|v|} - \frac{\tilde{v}}{|\tilde{v}|} \right| + \left| \dot{\gamma}_{v,i}(t) - \frac{v}{|v|} \right| + \left| \dot{\gamma}_{\tilde{v},j}(t) - \frac{\tilde{v}}{|\tilde{v}|} \right| \\ &\leq \frac{1}{|v|} |v - \tilde{v}| + 2\frac{|v|}{R} \\ &\leq \frac{d_{\mathcal{M}}(x_i, x_j)}{R} + 2\frac{|v|}{R}. \end{aligned}$$

□

From our assumption that the density  $p$  was in  $C^2(\mathcal{M})$  it follows that

$$p(\exp_{x_i}(v)) = p(x_i) + \langle \nabla p(x_i), v \rangle + R_i(v),$$

$$p(\exp_{x_j}(\tilde{v})) = p(x_j) + \langle \nabla p(x_j), \tilde{v} \rangle + R_j(\tilde{v}),$$

where  $\nabla p(x_i)$  is the gradient (in  $\mathcal{M}$ ) of  $p$  at the point  $x_i$ , and the remainder terms satisfy

$$\max\{|R_i(v)|, |R_j(\tilde{v})|\} \leq C_p |v|^2,$$

for a constant  $C_p$  that depends on a uniform bound on second derivatives of  $p$ . Likewise,

$$\max\{|p(x_i) - p(x_j)|, |\nabla p(x_i) - \nabla p(x_j)|\} \leq C_p d_{\mathcal{M}}(x_i, x_j).$$

Plugging the previous identities in the expressions (6.21) and (6.23), using

$$\int_{B_{x_i}(0, r_-)} p(x_i) v dv = 0, \quad \int_{B_{x_i}(0, r_-)} p(x_j) \tilde{v} d\tilde{v} = 0,$$

inequality (6.22), the bound on accelerations from Lemma 6.2.7, and finally Assumption 6.1.2, we can conclude that

$$\begin{aligned} |b_i - b_j| &\leq \int_{B_{x_i}(0, r_-)} \int_0^{|v|} \int_0^t |p(x_j) \ddot{\gamma}_{\tilde{v}, j}(s) - p(x_i) \ddot{\gamma}_{v, i}(s)| ds dt dv + \frac{C_p}{R} \alpha_m r^{m+2} (r + d_{\mathcal{M}}(x_i, x_j)) \\ &\leq p_{max} \cdot \int_{B_{x_i}(0, r_-)} \int_0^{|v|} \int_0^t |\ddot{\gamma}_{\tilde{v}, j}(s) - \ddot{\gamma}_{v, i}(s)| ds dt dv \\ &\quad + \int_{B_{x_i}(0, r_-)} \int_0^{|v|} \int_0^t |p(x_j) - p(x_i)| |\ddot{\gamma}_{v, i}(s)| ds dt dv + \frac{C_p}{R} \alpha_m r^{m+2} (r + d_{\mathcal{M}}(x_i, x_j)) \\ &\leq p_{max} \cdot \int_{B_{x_i}(0, r_-)} \int_0^{|v|} \int_0^t |\ddot{\gamma}_{\tilde{v}, j}(s) - \ddot{\gamma}_{v, i}(s)| ds dt dv + \frac{C_p}{R} \alpha_m r^{m+2} (r + d_{\mathcal{M}}(x_i, x_j)). \end{aligned} \tag{6.24}$$

In the above,  $C_p$  is a constant that depends on derivatives of  $p$  of order 1 and order 2 (and in particular is equal to zero when  $p$  is constant) and  $\alpha_m$  is the volume of the  $m$ -dimensional unit ball.

Proposition 6.2.5 now follows from the next lemma where we bound the difference of accelerations.

**Lemma 6.2.8.** *Let  $v \in B_{x_i}(0, r_-)$  and let  $\tilde{v}$  be as in Lemma 6.2.6. Then, for all  $t \in [0, |v|]$*

we have

$$|\ddot{\gamma}_{v,i}(t) - \ddot{\gamma}_{\tilde{v},j}(t)| \leq \left(2\frac{\sqrt{m}}{R^2} + C_{\mathcal{M}}\right) (2|v| + d_{\mathcal{M}}(x_i, x_j)) + 2C_{\mathcal{M}} \left(\frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R}\right),$$

where  $C_{\mathcal{M}}$  is a constant that depends on  $\mathcal{M}$  (a third order term).

*Proof.* For a fixed  $t \in [0, |v|]$  we let

$$x := \gamma_{v,i}(t), \quad \tilde{x} := \gamma_{\tilde{v},j}(t).$$

We start by constructing a convenient linear map

$$\eta \in T_x \mathcal{M}^\perp \mapsto \tilde{\eta} \in T_{\tilde{x}} \mathcal{M}^\perp.$$

For this purpose we use an *orthonormal frame*  $E_1, \dots, E_m$  on a neighborhood (in  $\mathcal{M}$ ) of  $x$  containing the geodesic connecting  $x$  and  $\tilde{x}$ . The frame is constructed by parallel transporting an orthonormal basis  $E_1(x), \dots, E_m(x)$  of  $T_x \mathcal{M}$  along geodesics emanating from  $x$ . Now, associated to  $\eta \in T_x \mathcal{M}^\perp$  we define the (normal) vector field  $N_\eta$  by

$$N_\eta := \eta - \sum_{l=1}^m \langle E_l, \eta \rangle E_l.$$

Equivalently,  $N_\eta$  can be written as

$$N_\eta(z) = \Pi_z(\eta),$$

where for a point  $z \in \mathcal{M}$ ,  $\Pi_z$  denotes the projection onto  $T_z \mathcal{M}^\perp$  (the orthogonal complement of the tangent plane at  $z$ ).

Let  $\phi_{x\tilde{x}}$  be the arc-length parameterized geodesic with  $\phi_{x\tilde{x}}(0) = x$  and  $\phi_{x\tilde{x}}(\tilde{t}) = \tilde{x}$ . We restrict the vector field  $N_\eta$  to the curve  $\phi_{x\tilde{x}}$  and abuse notation slightly to write  $N_\eta(s)$  and

$E_j(s)$  for the value of the vector fields at the point  $\phi_{x\tilde{x}}(s)$ . We let  $\tilde{\eta} := N_\eta(\tilde{t})$  and notice that

$$|\eta - \tilde{\eta}| = \left( \sum_{l=1}^m \langle E_l(\tilde{t}), \eta \rangle^2 \right)^{1/2} = \left( \sum_{l=1}^m \langle E_l(\tilde{t}) - E_l(0), \eta \rangle^2 \right)^{1/2} \leq \frac{\sqrt{m}d_{\mathcal{M}}(x, \tilde{x})|\eta|}{R}, \quad (6.25)$$

where in the last line we have used that  $|E_l(\tilde{t}) - E_l(0)| \leq \frac{\tilde{t}}{R}$  (proved in the exact same way as (6.22)).

Let  $\eta \in T_x\mathcal{M}^\perp$  be a unit norm vector and let  $\tilde{\eta}$  be as constructed before. Since  $N_\eta$  is a normal vector field which locally extends  $\eta$  we can follow the characterization for the shape operator in Proposition 2.3 Chapter 6 in Do Carmo and Flaherty Francis [1992] and deduce that:

$$\langle \ddot{\gamma}_{v,i}(t), \eta \rangle = \langle S_\eta(\dot{\gamma}_{v,i}), \dot{\gamma}_{v,i}(t) \rangle = \left\langle \frac{d}{dt} N_\eta(\gamma_{v,i}(t)), \dot{\gamma}_{v,i}(t) \right\rangle.$$

Moreover, the smoothness of the manifold  $\mathcal{M}$  allows us to extend  $N_\eta$  smoothly to a neighborhood in  $\mathbb{R}^d$  of  $x$  and  $\tilde{x}$  (we also use  $N_\eta$  to represent the extension). Indeed, for any point  $z$  in a tubular neighborhood of  $\mathcal{M}$  of width smaller than  $R$  we can define

$$N_\eta(z) := N_\eta(Proj_{\mathcal{M}}(z)),$$

where  $Proj_{\mathcal{M}}$  is the projection onto  $\mathcal{M}$  (which is well defined for points within distance  $R$  from  $\mathcal{M}$ ). The smoothness of  $N_\eta$  in particular implies that

$$\|DN_\eta(x) - DN_\eta(\tilde{x})\| \leq C_{\mathcal{M}}|x - \tilde{x}| \leq C_{\mathcal{M}}d_{\mathcal{M}}(x, \tilde{x}),$$

$$\|DN_\eta(\tilde{x})\| \leq C_{\mathcal{M}},$$

where  $DN_\eta$  is the matrix of derivatives of the vector field  $N_\eta$ , and where  $C_{\mathcal{M}}$  is some constant that only depends on  $\mathcal{M}$ . With this extension at hand, we can then use the chain rule and



write:

$$\langle \ddot{\gamma}_{v,i}(t), \eta \rangle = \langle DN(x)\dot{\gamma}_{v,i}(t), \dot{\gamma}_{v,i}(t) \rangle,$$

and in a similar fashion

$$\langle \ddot{\gamma}_{\tilde{v},j}(t), \eta \rangle = \langle \ddot{\gamma}_{\tilde{v},j}(t), \eta - \tilde{\eta} \rangle + \langle \ddot{\gamma}_{\tilde{v},j}(t), \tilde{\eta} \rangle = \langle \ddot{\gamma}_{\tilde{v},j}(t), \eta - \tilde{\eta} \rangle + \langle -DN(\tilde{x})\dot{\gamma}_{\tilde{v},j}(t), \dot{\gamma}_{\tilde{v},j}(t) \rangle.$$

Using the triangle and Cauchy-Schwarz inequalities we obtain

$$\begin{aligned} |\langle \ddot{\gamma}_{v,i}(t) - \ddot{\gamma}_{\tilde{v},j}(t), \eta \rangle| &\leq |\ddot{\gamma}_{\tilde{v},j}| |\eta - \tilde{\eta}| + \|DN(x) - DN(\tilde{x})\| |\dot{\gamma}_{v,i}|^2 + \\ &\quad \|DN(\tilde{x})\| |\dot{\gamma}_{v,i} - \dot{\gamma}_{\tilde{v},j}| (|\dot{\gamma}_{\tilde{v},j}| + |\dot{\gamma}_{v,i}|) \\ &\leq \frac{\sqrt{m}}{R^2} d_{\mathcal{M}}(x, \tilde{x}) + C_{\mathcal{M}} d_{\mathcal{M}}(x, \tilde{x}) + 2C_{\mathcal{M}} \left( \frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R} \right). \end{aligned}$$

Since the above inequality holds for all  $\eta \in T_x \mathcal{M}^\perp$  with norm one, we conclude that

$$|\Pi_x(\ddot{\gamma}_{v,i}(t)) - \Pi_x(\ddot{\gamma}_{\tilde{v},j}(t))| \leq \frac{\sqrt{m}}{R^2} d_{\mathcal{M}}(x, \tilde{x}) + C_{\mathcal{M}} d_{\mathcal{M}}(x, \tilde{x}) + 2C_{\mathcal{M}} \left( \frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R} \right),$$

where we recall  $\Pi_x$  represents the projection onto  $T_x \mathcal{M}^\perp$ . Moreover, since  $\ddot{\gamma}_{v,i}(t)$  is the acceleration of a unit speed geodesic passing through  $x$ , we know that  $\ddot{\gamma}_{v,i}(t) \in T_x \mathcal{M}^\perp$ , so that  $\Pi_x(\ddot{\gamma}_{v,i}) = \ddot{\gamma}_{v,i}$ . Similarly we have  $\Pi_{\tilde{x}}(\ddot{\gamma}_{\tilde{v},j}) = \ddot{\gamma}_{\tilde{v},j}$  (where  $\Pi_{\tilde{x}}$  represents projection onto  $T_{\tilde{x}} \mathcal{M}^\perp$ ). Hence

$$|\ddot{\gamma}_{v,i}(t) - \ddot{\gamma}_{\tilde{v},j}(t)| \leq |\Pi_x \ddot{\gamma}_{v,i}(t) - \Pi_x \ddot{\gamma}_{\tilde{v},j}(t)| + |\Pi_x \ddot{\gamma}_{\tilde{v},j}(t) - \ddot{\gamma}_{\tilde{v},j}(t)|, \quad (6.26)$$

and so it remains to find a bound for  $|\Pi_x \ddot{\gamma}_{\tilde{v},j}(t) - \ddot{\gamma}_{\tilde{v},j}(t)|$ . We can write

$$\Pi_x \ddot{\gamma}_{\tilde{v},j} = \ddot{\gamma}_{\tilde{v},j} - \sum_{l=1}^m \langle \ddot{\gamma}_{\tilde{v},j}, E_l(0) \rangle E_l(0).$$

Therefore,

$$|\ddot{\gamma}_{\tilde{v},j} - \Pi_x \ddot{\gamma}_{\tilde{v},j}| = \left( \sum_{l=1}^m \langle \ddot{\gamma}_{\tilde{v},j}, E_l(0) \rangle^2 \right)^{1/2} = \left( \sum_{l=1}^m \langle \ddot{\gamma}_{\tilde{v},j}, E_l(0) - E_l(\tilde{t}) \rangle^2 \right)^{1/2} \leq \sqrt{m} \frac{d_{\mathcal{M}}(x, \tilde{x})}{R^2}.$$

Putting everything together we deduce that

$$\begin{aligned} |\ddot{\gamma}_{v,i} - \ddot{\gamma}_{\tilde{v},j}| &\leq \left(2 \frac{\sqrt{m}}{R^2} + C_{\mathcal{M}}\right) d_{\mathcal{M}}(x, \tilde{x}) + 2C_{\mathcal{M}} \left( \frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R} \right) \\ &\leq \left(2 \frac{\sqrt{m}}{R^2} + C_{\mathcal{M}}\right) (2|v| + d_{\mathcal{M}}(x_i, x_j)) + 2C_{\mathcal{M}} \left( \frac{|v|}{R} + \frac{d_{\mathcal{M}}(x_i, x_j)}{R} \right), \end{aligned}$$

where in the last step we have used the triangle inequality

$$d_{\mathcal{M}}(x, \tilde{x}) \leq d_{\mathcal{M}}(x, x_i) + d_{\mathcal{M}}(x_i, x_j) + d_{\mathcal{M}}(x_j, \tilde{x}) \leq 2|v| + d_{\mathcal{M}}(x_i, x_j).$$

□

**Remark 6.2.9.** Notice that the computations in the proof of Proposition 6.2.5 also show that

$$|b_i| \leq Cr^{m+2}, \quad i = 1, \dots, n.$$

Indeed, this can be seen directly from (6.21), Lemma 6.2.7 (which bounds the acceleration term), and the fact that the first term on the right-hand side of the following expression drops by symmetry:

$$\int_{B_{x_i}(0, r_-)} p(\exp_{x_i}(v)) v dv = p(x_i) \int_{B_{x_i}(0, r_-)} v dv + \int_{B_{x_i}(0, r_-)} (\langle \nabla p(x_i), v \rangle + R_i(v)) v dv.$$

### 6.2.5 Bounding Sampling Error

We will make use of two concentration inequalities to bound the sampling error. We first recall Hoeffding's inequality.

**Lemma 6.2.10** (Hoeffding's inequality). *Let  $w_1, \dots, w_n$  be i.i.d samples from a random variable  $w$  taking values in the interval  $[0,1]$  and let  $\bar{w}$  be the sample average. Then,*

$$\mathbb{P}(|\bar{w} - \mathbb{E}[\bar{w}]| > t) \leq 2e^{-2nt^2}.$$

The next is a generalization for random vectors that follows directly from the simple and elegant work Pinelis [1992] (more precisely, Theorem 3).

**Lemma 6.2.11.** *Let  $W_1, \dots, W_n$  be i.i.d samples from a random vector  $W$  such that  $|W| \leq M$  for some constant  $M$ , and  $\mathbb{E}[W] = 0$ . Let  $\bar{W}$  be the sample average. Then,*

$$\mathbb{P}\left(|\bar{W} - \mathbb{E}[\bar{W}]| > \sqrt{\frac{M^2}{n}}t\right) \leq 2e^{-t^2/16}.$$

**Proposition 6.2.12.** *Suppose Assumption 6.1.3 holds. Then,*

$$\mathbb{P}\left(|\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{2^{m+4}}{\alpha_m p_{min}}}r^3\right) \leq 4e^{-c nr^{\max\{2m, m+4\}}}, \text{ where } c = \min\left\{\frac{\alpha_m^2 p_{min}^2}{4^{m+2}}, \frac{1}{16}\right\}.$$

*In particular, if  $nr^{\max\{2m, m+4\}} \gg 1$ , then  $|\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| \leq \sqrt{\frac{2^{m+4}}{\alpha_m p_{min}}}r^3$  with high probability.*

*Proof.* Let  $N_i$  be the number of points in  $B(y_i, r)$ . Notice that  $\tilde{x}_i + \tilde{z}_i - \mathbb{E}_i[\tilde{X}_i + \tilde{Z}_i]$  is centered and bounded by  $2r$  in norm, and  $\bar{y}_i = \overline{\tilde{x}_i + \tilde{z}_i}$ . Then Lemma 6.2.11 implies

$$\mathbb{P}_i\left(|\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{N_i}}t \mid N_i\right) \leq 2e^{-t^2/16}.$$

By the law of iterated expectations it follows that

$$\mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{N_i}t} \right) \leq 2e^{-t^2/16}.$$

Next note that  $N_i$ , the number of points  $y_j$  in  $B(y_i, r)$ , can be bounded below by  $\tilde{N}_i$ , the number of points  $x_j$  that lie in the ball  $B_{\mathcal{M}}(x_i, r_-)$ . Thus,

$$\mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{\tilde{N}_i}t} \right) \leq 2e^{-t^2/16}. \quad (6.27)$$

Now we find probabilistic bound for  $\tilde{N}_i$ . Let  $w_j = \mathbf{1}\{x_j \in B(x_i, r_-)\}$ . Then given  $x_i$ , the  $w_j$  are i.i.d samples from Bernoulli( $q_i$ ), where  $q_i = \mu(B_{\mathcal{M}}(x_i, r_-))$ . Lemma 6.2.10 implies

$$\mathbb{P}_i \left( |\tilde{N}_i - nq_i| > nt | x_i \right) \leq 2e^{-2nt^2}.$$

Again by the law of iterated expectation and rearranging terms, we have

$$\mathbb{P} \left( \tilde{N}_i < n(q_i - t) \right) \leq 2e^{-2nt^2}. \quad (6.28)$$

Combining (6.27) and (6.28), we obtain

$$\begin{aligned} \mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{n(q_i - s)}t} \right) &= \mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{n(q_i - s)}t}, \tilde{N}_i < n(q_i - s) \right) \\ &\quad + \mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{n(q_i - s)}t}, \tilde{N}_i \geq n(q_i - s) \right) \\ &\leq \mathbb{P} \left( \tilde{N}_i < n(q_i - s) \right) + \mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{4r^2}{\tilde{N}_i}t} \right) \\ &\leq 2e^{-2ns^2} + 2e^{-t^2/16}. \end{aligned}$$

Under Assumption 6.1.3, (6.11) implies  $q_i \geq \frac{\alpha_m p_{\min}}{2^{m+1}} r^m$ . Taking  $s = \frac{\alpha_m p_{\min}}{2^{m+2}} r^m$  and  $t = \sqrt{nr^{m+4}}$ , we see that

$$\mathbb{P} \left( |\bar{y}_i - \mathbb{E}_i[\bar{y}_i]| > \sqrt{\frac{2^{m+4}}{\alpha_m p_{\min}}} r^3 \right) \leq 2e^{-\frac{\alpha_m^2 p_{\min}^2}{4^{m+2}} nr^{2m}} + 2e^{-nr^{m+4}/16} \leq 4e^{-c nr^{\max\{2m, m+4\}}},$$

where  $c = \min \left\{ \frac{\alpha_m^2 p_{\min}^2}{4^{m+2}}, \frac{1}{16} \right\}$ . The result then follows.  $\square$

Theorem 6.1.4 now follows by combining Lemma 6.2.4, Propositions 6.2.3, 6.2.5, and Proposition 6.2.12 together with a union bound.

### 6.3 From Local Regularization to Global Estimates

In this section we use the local estimates (6.7) to show spectral convergence of  $\Delta_{\bar{y}_{n,\varepsilon}}$  towards the continuum Laplace-Beltrami operator. We first make some definitions. Recall that the graph  $\Gamma_{\delta,\varepsilon} = ([n], W)$  has weights

$$W(i, j) = \frac{2(m+2)}{\alpha_m \varepsilon^{m+2} n} \mathbf{1}\{\delta(i, j) < \varepsilon\},$$

where  $m$  is the dimension of  $\mathcal{M}$  and  $\alpha_m$  is the volume of the  $m$ -dimensional Euclidean unit ball. For a function  $u : [n] \rightarrow \mathbb{R}$ , we denote its value on the  $i$ -th node as  $u(i)$ . We then define the discrete Dirichlet energy of  $u$  as

$$E_{\delta,\varepsilon}[u] = \frac{m+2}{\alpha_m \varepsilon^{m+2} n} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{\delta(i, j) < \varepsilon\} |u(i) - u(j)|^2$$

and the  $L^2$  norm of  $u$  as

$$\|u\|^2 = \frac{1}{n} \sum_{i=1}^n |u(i)|^2.$$

Given that  $\Delta_{\delta,\varepsilon}$  is a positive semi-definite operator, we can use the minimax principle (see for example Lieb and Loss [2001]) to write

$$\lambda_\ell(\Gamma_{\delta,\varepsilon}) = \min_L \max_{u \in L \setminus \{0\}} \frac{E_{\delta,\varepsilon}[u]}{\|u\|^2},$$

where  $\lambda_\ell(\Gamma_{\delta,\varepsilon})$  is the  $\ell$ -th smallest eigenvalue of  $\Delta_{\Gamma_{\delta,\varepsilon}}$  and the minimum is taken over all subspaces  $L$  of dimension  $\ell$ . The following lemma compares the eigenvalues of the discrete graphs constructed using  $\delta_{\mathcal{X}_n}$  and  $\delta_{\bar{\mathcal{Y}}_n}$ .

**Lemma 6.3.1.** *Let  $\eta$  be the bound in (6.7) so that for all  $i, j$  with  $d_{\mathcal{M}}(x_i, x_j) \leq r$  we have*

$$|\delta_{\mathcal{X}_n}(i, j) - \delta_{\bar{\mathcal{Y}}_n}(i, j)| \leq \eta.$$

*Suppose that  $\varepsilon$  is chosen so that  $\varepsilon \geq 2Cm\eta$ , for some universal constant  $C$ . Then,*

$$\left(1 - Cm\frac{\eta}{\varepsilon}\right)\lambda_\ell(\Gamma_{\mathcal{X}_n, \varepsilon - \eta}) \leq \lambda_\ell(\Gamma_{\bar{\mathcal{Y}}_n, \varepsilon}) \leq \left(1 + Cm\frac{\eta}{\varepsilon}\right)\lambda_\ell(\Gamma_{\mathcal{X}_n, \varepsilon + \eta}). \quad (6.29)$$

*Proof.* We first compare the Dirichlet energies. Since  $\delta_{\mathcal{X}_n}(i, j) < \delta_{\bar{\mathcal{Y}}_n}(i, j) + \eta$ , we have

$$\begin{aligned} E_{\bar{\mathcal{Y}}_n, \varepsilon}[u] &= \frac{m+2}{\alpha_m \varepsilon^{m+2} n} \sum_i \sum_j \mathbf{1}\{\delta_{\bar{\mathcal{Y}}_n}(i, j) < \varepsilon\} |u_i - u_j|^2 \\ &\leq \frac{m+2}{\alpha_m \varepsilon^{m+2} n} \sum_i \sum_j \mathbf{1}\{\delta_{\mathcal{X}_n}(i, j) < \varepsilon + \eta\} |u_i - u_j|^2 \\ &= \left(\frac{\varepsilon + \eta}{\varepsilon}\right)^{m+2} E_{\mathcal{X}_n, \varepsilon + \eta}[u] \\ &\leq \left(1 + Cm\frac{\eta}{\varepsilon}\right) E_{\mathcal{X}_n, \varepsilon + \eta}[u]. \end{aligned} \quad (6.30)$$

Now we use the minimax principle to show the upper-bound on (6.29). Let  $u_1, \dots, u_\ell$  be the first  $\ell$  eigenvectors of  $\Delta_{\mathcal{X}_n, \varepsilon + \eta}$  and let  $L = \text{span}\{u_1, \dots, u_\ell\}$ . Then  $\dim L = \ell$  and for any

$u \in L$ ,  $E_{\mathcal{X}_{n,\varepsilon+\eta}}[u] \leq \lambda_\ell(\Gamma_{\mathcal{X}_{n,\varepsilon+\eta}})\|u\|^2$ . Then by (6.30), we have

$$\lambda_\ell(\Gamma_{\bar{\mathcal{Y}}_{n,\varepsilon}}) \leq \max_{L \setminus \{0\}} \frac{E_{\bar{\mathcal{Y}}_{n,\varepsilon}}[u]}{\|u\|^2} \leq \left(1 + Cm \frac{\eta}{\varepsilon}\right) \max_{L \setminus \{0\}} \frac{E_{\mathcal{X}_{n,\varepsilon+\eta}}[u]}{\|u\|^2} = \left(1 + Cm \frac{\eta}{\varepsilon}\right) \lambda_\ell(\Gamma_{\mathcal{X}_{n,\varepsilon+\eta}}).$$

By a similar argument applied to  $\Gamma_{\mathcal{X}_{n,\varepsilon-\eta}}$  and  $\Gamma_{\bar{\mathcal{Y}}_{n,\varepsilon}}$ , we get the lower-bound in (6.29).  $\square$

**Remark 6.3.2.** *With the convergence of eigenvalues and the relationship between the Dirichlet energies it is also possible to make statements about convergence of eigenvectors (or better yet, spectral projections).*

The spectral convergence towards the continuum (Theorem 6.1.6) is a consequence of the following theorem, proved in [García Trillos et al., 2019a, Corollary 2].

**Theorem 6.3.3.** *Let  $d_\infty$  be the  $\infty$ -OT distance between  $\mu_n$  and  $\mu$ . Suppose  $\varepsilon$  satisfies the conditions in Equation (6.9) and that Assumptions 6.1.1 and 6.1.2 hold. Then*

$$\frac{|\lambda_\ell(\Gamma_{\mathcal{X}_{n,\varepsilon}}) - \lambda_\ell(\mathcal{M})|}{\lambda_\ell(\mathcal{M})} \leq \tilde{C} \left( \frac{d_\infty}{\varepsilon} + (1 + \sqrt{\lambda_\ell(\mathcal{M})})\varepsilon + \left(K + \frac{1}{R^2}\right)\varepsilon^2 \right),$$

where  $\tilde{C}$  only depends on  $m$  and the regularity of  $p$ .

Combining Lemma 6.3.1 and Theorem 6.3.3 gives Theorem 6.1.6.

## 6.4 Numerical Experiments

In this section we present a series of numerical experiments where we conduct local regularization on three different datasets. In Section 6.4.1 we consider a toy example with artificial data generated by perturbing points sampled uniformly from the unit, two-dimensional sphere embedded in  $\mathbb{R}^d$  with  $d = 100$ . We show that the approximation of the hidden Euclidean distances between unperturbed points is significantly improved by locally regularizing the data, and that this improvement translates into better spectral approximation

of the spherical Laplacian. Our numerical findings corroborate the theory developed in the previous two sections. In Section 6.4.2 we consider the two-moon and MNIST datasets and show that graphs constructed with locally regularized data can be used to improve the performance of a simple graph-based optimization method for semi-supervised classification.

### 6.4.1 Distance $\mathcal{E}$ Spectrum

Here we study the effect of local regularization on distance approximation and spectral convergence, as an illustration of the results from Sections 6.2 and 6.3. In our toy model we consider uniform samples from the unit two-dimensional sphere  $\mathcal{M} = \mathcal{S}$  embedded in  $\mathbb{R}^d$ , with  $d = 100$ . The motivation for such a choice is that the eigenvalues of the associated Laplace-Beltrami operator on  $\mathcal{S}$  are known explicitly (see for example Olver). Indeed, after appropriate normalization,  $\Delta_{\mathcal{S}}$  admits eigenvalues  $\ell(\ell + 1), \ell \in \mathbb{N}$ , with corresponding multiplicity  $2\ell + 1$ .

The dataset is generated by sampling  $n = 3000$  points  $x_i$  uniformly from the sphere and adding uniform noise  $z_i$  normal to the tangent plane, and bounded by  $\sigma$  in norm. To be more precise, the noise is normal to the sphere for the first three dimensions and uniform in all directions for the rest dimensions. Local regularization is performed by taking  $r \propto \sqrt{\sigma}$  and the graph is constructed with  $\varepsilon = 2n^{-1/4}$ . The optimal proportion constant in  $r$  is not obvious from our theory and in the experiments below we choose  $r = \sqrt{\sigma}/3$  for  $\sigma = 0.1$  and  $r = \sqrt{\sigma}$  for the rest of the  $\sigma$ 's. We first show that the  $\bar{y}_i$  give a better approximation of the pairwise distances of the  $x_i$  than the  $y_i$  do. We only consider those nodes  $i, j$  such that  $\delta_{\mathcal{X}_n}(i, j) < \varepsilon$  (i.e. the nodes that are relevant for the construction of the graph Laplacians). More precisely, let  $D_{\mathcal{X}_n}$  be the matrix whose  $ij$ th entry is  $\delta_{\mathcal{X}_n}(i, j)\mathbf{1}\{\delta_{\mathcal{X}_n}(i, j) < \varepsilon\}$ . Similarly, we define  $[D_{\mathcal{Y}_n}]_{ij} = \delta_{\mathcal{Y}_n}(i, j)\mathbf{1}\{\delta_{\mathcal{X}_n}(i, j) < \varepsilon\}$  and  $[D_{\bar{\mathcal{Y}}_n}]_{ij} = \delta_{\bar{\mathcal{Y}}_n}(i, j)\mathbf{1}\{\delta_{\mathcal{X}_n}(i, j) < \varepsilon\}$ . In Table 6.1 we compare the entrywise  $\infty$ -norm of the  $D_{\mathcal{X}_n} - D_{\mathcal{Y}_n}$  and  $D_{\mathcal{X}_n} - D_{\bar{\mathcal{Y}}_n}$  for different values of  $\sigma$ . We see that the improvement is substantial.



	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	$\sigma = 0.7$	$\sigma = 0.9$
$\ D\mathcal{X}_n - D\mathcal{Y}_n\ _\infty$	0.095	0.298	0.505	0.723	0.937
$\ D\mathcal{X}_n - D\bar{\mathcal{Y}}_n\ _\infty$	0.084	0.090	0.087	0.093	0.144

Table 6.1: Entrywise  $\infty$ -norm of  $D\mathcal{X}_n - D\mathcal{Y}_n$  and  $D\mathcal{X}_n - D\bar{\mathcal{Y}}_n$  on  $\mathcal{S}$  for several  $\sigma$ 's.

Next we study the spectral approximation of Laplacians by comparing the spectra of  $\Delta\mathcal{X}_{n,\varepsilon}$ ,  $\Delta\mathcal{Y}_{n,\varepsilon}$  with that of  $\Delta\bar{\mathcal{Y}}_{n,\varepsilon}$ . Note that since the  $x_i$  are uniformly distributed, the density  $p$  on  $\mathcal{S}$  that they are sampled from is constant and equal to  $\frac{1}{\text{vol}\mathcal{M}}$ . So for the spectra of the graph Laplacians to match in scale with that of  $\Delta_{\mathcal{S}}$ , the weights should be rescaled according to

$$W(i, j) = \frac{2(m+2)\text{vol}(\mathcal{M})}{\alpha_m \varepsilon^{m+2n}},$$

where  $\text{vol}(\mathcal{M})$  is the volume of the manifold and equals  $4\pi$  in this case. In Figure 6.1 we compare the first 100 eigenvalues of  $\Delta\mathcal{X}_{n,\varepsilon}$ ,  $\Delta\mathcal{Y}_{n,\varepsilon}$ , and  $\Delta\bar{\mathcal{Y}}_{n,\varepsilon}$  with the continuum spectrum. We see that when the noise size is large, the Euclidean graph Laplacian  $\Delta\mathcal{Y}_{n,\varepsilon}$  does not give a meaningful approximation of the continuum spectrum, while the locally regularized version  $\Delta\bar{\mathcal{Y}}_{n,\varepsilon}$  still performs well.

**Remark 6.4.1.** *While our theory in Section 6.2 suggests the choice  $r \propto \sqrt{\sigma}$  in the small  $r$  and large  $n$  limit, for practical purposes some other scalings may give better results. Indeed for the above  $\sigma$ 's, choosing  $r = \sigma$  seems to give better spectral approximation. The choice of the local-regularization parameter will be further investigated in Section 6.4.2.2 in the context of a semi-supervised classification task, where a data-driven (cross-validation) approach can be used.*

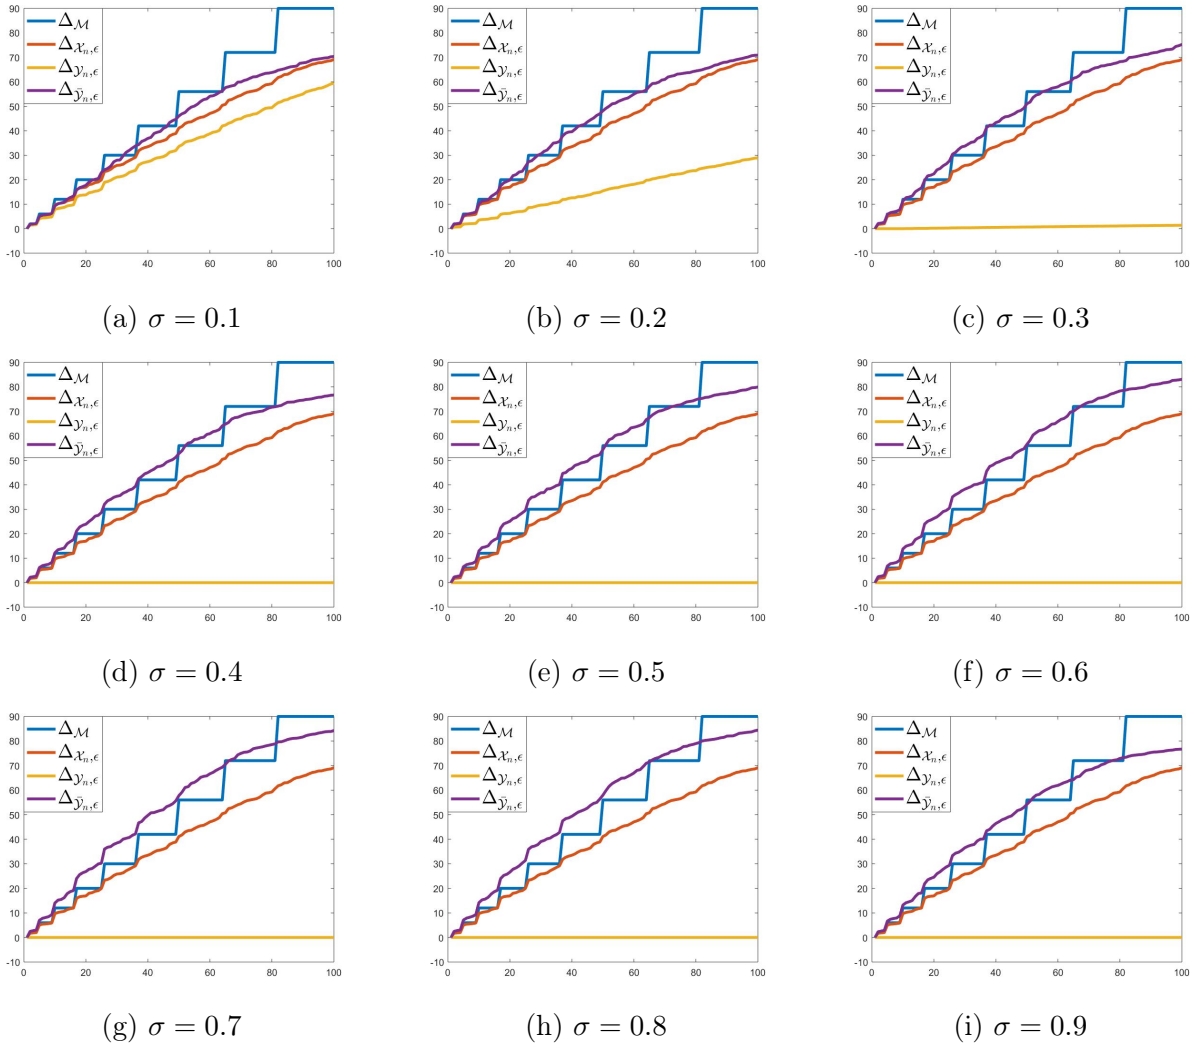


Figure 6.1: Comparison of spectra of continuum Laplacian,  $\Delta_{\mathcal{X}_{n,\varepsilon}}$ ,  $\Delta_{\mathcal{Y}_{n,\varepsilon}}$  and  $\Delta_{\bar{\mathcal{Y}}_{n,\varepsilon}}$  for different values of  $\sigma$ .

### 6.4.2 Classification

In this subsection we demonstrate the practical use of local regularization by applying it to classification problems. To show the potential benefits, we consider synthetic and real datasets, namely the two moons and MNIST datasets. Since in one of our experiments we study a real dataset, where in general the connectivity parameter in an  $\varepsilon$  graph is hard to tune, we instead consider fully connected graphs with self-tuning weights. Precisely, given a similarity  $\delta : [n] \times [n] \rightarrow [0, \infty)$  we define, following Zelnik-Manor and Perona [2005], the

weights by

$$W(i, j) = \exp\left(-\frac{\delta(i, j)^2}{2\tau(i)\tau(j)}\right), \quad (6.31)$$

where  $\tau(i)$  is the similarity between the  $i$ -th data point and its  $K$ -th nearest neighbor with respect to the distance  $\delta$ . As before, we denote by  $\Gamma_{\mathcal{X}_n}$ ,  $\Gamma_{\mathcal{Y}_n}$  and  $\Gamma_{\bar{\mathcal{Y}}_n}$  the graphs constructed with similarities  $\delta_{\mathcal{X}_n}$ ,  $\delta_{\mathcal{Y}_n}$ , and  $\delta_{\bar{\mathcal{Y}}_n}$ . Instead of specifying a universal  $\varepsilon$  representing the connectivity length-scale, the neighborhood for each point is selected from using the local geometry which varies in space. It amounts to choosing different values of  $\varepsilon$  adaptively depending on the local scale, as proposed in Zelnik-Manor and Perona [2005]. Since the  $\tau(i)$  are defined by considering  $K$ -nearest neighbors, a natural variant of the above fully connected graph is to set the weights to be 0 whenever  $x_i$  and  $x_j$  are not among the  $K$ -nearest neighbors of each other. In other words, we can construct a (symmetrized)  $K$ -NN graph with the same  $K$  as in the definition of  $\tau(i)$  and the nonzero weights are the same as above. It turns out that empirically this  $K$ -NN version can improve the classification performance substantially, but to illustrate the local regularization idea, we will present results for both graph constructions. We shall denote these two types of graphs as fully-connected and  $K$ -NN variants for brevity, or fully and  $K$ -NN for short.

In the following, we focus on the semi-supervised learning setting where we are given  $n$  data points with the first  $J$  being labeled. The classification is done by minimizing a probit functional as explained below. Let  $\Delta_\delta$  be a normalized graph Laplacian constructed on the dataset, which will be constructed using  $\mathcal{X}_n$ ,  $\mathcal{Y}_n$  and  $\bar{\mathcal{Y}}_n$  and  $\Delta_\delta = I - D^{-1/2}WD^{-1/2}$  as compared with (6.4). Let  $(\lambda_i, q_i)$ ,  $i = 1, \dots, n$  be the associated eigenvalue-eigenvector pairs, and let  $U = \text{span}\{q_2, \dots, q_n\}$ . The classifier is set to be the sign of the minimizer  $u$  of the

functional

$$\mathcal{J}(u) := \frac{1}{2c} \langle u, \Delta_\delta u \rangle - \sum_{j=1}^J \log \left( \Phi(y(j)u(j); \gamma) \right), \text{ with } c := n \left( \sum_{i=2}^n \lambda_i^{-1} \right)^{-1},$$

where  $\{y(j)\}_{j=1}^J$  is the vector of labels and  $\Phi$  is the cdf of  $\mathcal{N}(0, \gamma^2)$ . The functional  $\mathcal{J}$  can be interpreted as the negative log posterior in a Bayesian setting, as discussed in Bertozzi et al. [2018]. Throughout our experiments we set  $\gamma = 0.1$ .

#### 6.4.2.1 Two Moons

We first study the two moons dataset (Bühler and Hein [2009]), which is generated by sampling points uniformly from two semi-circles of unit radius centered at  $(0, 0)$  and  $(1, 0.5)$  and then embedding the dataset in  $\mathbb{R}^d$ , with  $d = 100$ . We then perturb the data by adding uniform noise with norm bounded by  $\sigma$ . As before, the noise in the first two dimensions are normal to the semicircles; the noise is taken to be uniform in the ambient space in the remaining dimensions. In addition to the semi-supervised setting, we also examine the unsupervised case.

We consider  $n = 1000$  points 1% of which have labels and we set  $K = 10$ . As pointed out in Remark 6.4.1, we choose the regularization parameter  $r$  to be equal to  $\sigma$ . We compare the approximation of distance matrix and classification performance on  $\mathcal{X}_n, \mathcal{Y}_n$ , and  $\bar{\mathcal{Y}}_n$ 's, as in Table 6.2 and Figure 6.4. Instead of comparing nodes that are within  $\delta_{\mathcal{X}_n}$ -distance  $\varepsilon$ , we consider nodes that are  $K$ -nearest neighbors of each other with respect to  $\delta_{\mathcal{X}_n}$ . As before, the regularized points  $\bar{\mathcal{Y}}_n$  approximate the pairwise distances better and moreover, they improve the classification performance. Especially for the fully-connected case, we see that  $\bar{\mathcal{Y}}_n$  is able to capture the exact correct labeling as the clean data does for moderate  $\sigma$ 's, while the noisy data  $\mathcal{Y}_n$  is making mistakes even when  $\sigma$  is as small as 0.3.

For further understanding, in Figure 6.3 we plot the first two coordinates of the points

	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	$\sigma = 0.7$	$\sigma = 0.9$
$\ D\mathcal{X}_n - D\mathcal{Y}_n\ _\infty$	0.109	0.344	0.589	0.795	0.996
$\ D\mathcal{X}_n - D\bar{\mathcal{Y}}_n\ _\infty$	0.064	0.164	0.240	0.372	0.431

Table 6.2: Entrywise  $\infty$ -norm of  $D\mathcal{X}_n - D\mathcal{Y}_n$  and  $D\mathcal{X}_n - D\bar{\mathcal{Y}}_n$  on two moons for different values of  $\sigma$ .

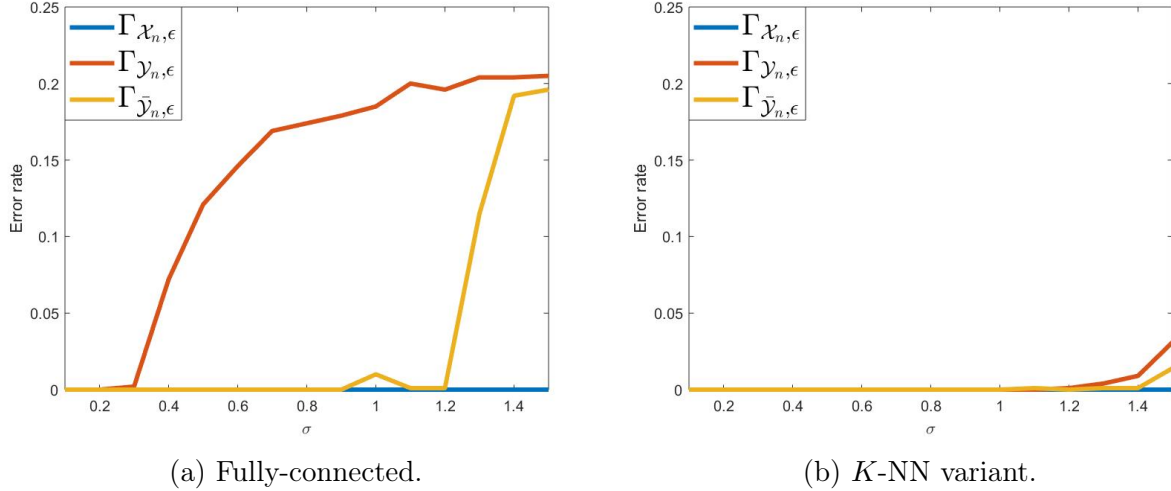


Figure 6.2: Classification error rates for  $\Gamma_{\mathcal{X}_n}$ ,  $\Gamma_{\mathcal{Y}_n}$  and  $\Gamma_{\bar{\mathcal{Y}}_n}$  on two moons for different values of  $\sigma$ .

in  $\mathcal{X}_n$ ,  $\mathcal{Y}_n$  and  $\bar{\mathcal{Y}}_n$  for large values of  $\sigma$ . We see that after local regularization, the first two coordinates of  $\bar{\mathcal{Y}}_n$  lie almost on the underlying manifold. The denoising effect of local regularization is apparent. Furthermore, we observe that the semicircles for  $\bar{\mathcal{Y}}_n$  are “shorter” than those of  $\mathcal{X}_n$ . In other words, points near the ends are pulled away from the boundaries. Moreover, if one looks carefully at the plots for  $\bar{\mathcal{Y}}_n$ , points are denser near the top and bottom. This illustrates that local regularization not only reduces noise, but also moves points to regions of high probability. We refer to Chen et al. [2016], Fukunaga and Hostetler [1975] and the references therein for some discussion on mean-shift and mode-seeking type algorithms.

**Remark 6.4.2.** *The two moons dataset is sampled from a manifold with boundaries, and so our theory does not directly apply. However, the numerical results seem to suggest that our theory continues to hold in this setting.*

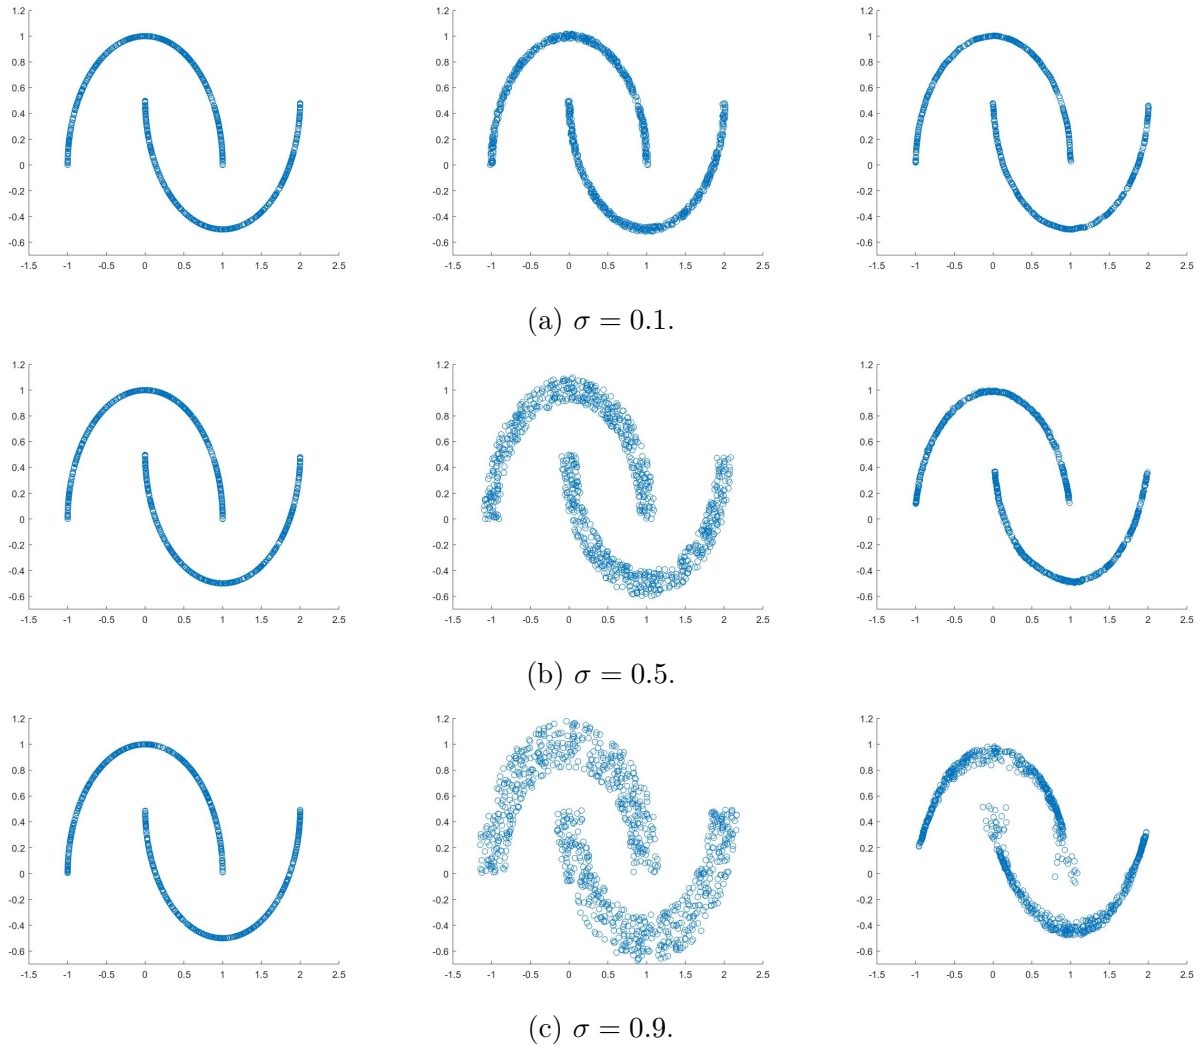


Figure 6.3: Visualization of the point clouds  $\mathcal{X}_n$ ,  $\mathcal{Y}_n$ , and  $\bar{\mathcal{Y}}_n$ . Each row contains scatter plots of the first two coordinates of the points in the datasets  $\mathcal{X}_n$ ,  $\mathcal{Y}_n$ , and  $\bar{\mathcal{Y}}_n$ .

**Remark 6.4.3.** *Additional numerical experiments not shown here suggest that applying local regularization within unsupervised spectral clustering gives qualitatively similar results to those shown in Figure 6.2 for a semi-supervised setting.*

### 6.4.2.2 MNIST

In this subsection we apply local regularization on the MNIST data-set of hand-written digits (LeCun [1998]). Each digit is described by a 784-dimensional vector, but the number

of degrees of freedom of the data-generating mechanism is much smaller. For instance, in Hein and Audibert [2005] the authors estimate the intrinsic dimension of the digits 1 from MNIST to be 8. However, unlike in the previous examples, here there is no explicitly available underlying manifold from which the digits are sampled. Instead of adding additional noise to the dataset, we directly apply local regularization to the digits and show that doing so improves their binary classification. Since the level of noise is unknown choosing the localization parameter  $r$  cannot be guided by the theory and in our experiments, we tune it by performing 2-fold cross validation on the label sets. When there are few labels, we repeatedly generate holdout sets and compare the overall error. Due to this practical difficulty of tuning  $r$ , we propose two variants of  $\Gamma_{\bar{\mathcal{Y}}_n}$  that can serve as alternatives in practice.

We study the classification performance of  $\Gamma_{\bar{\mathcal{Y}}_n}$  for different pairs of digits. We consider a semi-supervised learning problem with  $n = 1000$  images and  $K = 20$  for the  $K$ -nearest neighbor variant. Table 6.3 shows the classification error percentage for four different pairs of digits when 4% of the digits are labeled. Table 6.4 shows the decrease on the classification error on the pair 4&9 as the percentage of labeled digits is increased.

Fully	3&8	5&8	4&9	7&9	$K$ -NN	3&8	5&8	4&9	7&9
$\Gamma_{\mathcal{Y}_n}$	27.7%	48.0%	48.0%	48.0%	$\Gamma_{\mathcal{Y}_n}$	7.6%	5.5%	13.3%	7.3%
$\Gamma_{\bar{\mathcal{Y}}_n}$	13.4%	17.4%	30.0%	15.3%	$\Gamma_{\bar{\mathcal{Y}}_n}$	6.0%	3.6%	9.6%	5.4%

Table 6.3: Classification error for different pairs of digits 3&8, 5&8, 4&9, and 7&9.

Fully	4%	8%	12%	16%	$K$ -NN	4%	8%	12%	16%
$\Gamma_{\mathcal{Y}_n}$	48.0%	42.7%	38.8%	29.4%	$\Gamma_{\mathcal{Y}_n}$	13.3%	10.9%	7.6%	5.1%
$\Gamma_{\bar{\mathcal{Y}}_n}$	30.0%	26.1%	21.9%	18.2%	$\Gamma_{\bar{\mathcal{Y}}_n}$	9.6%	6.4%	6.0%	4.5%

Table 6.4: Classification error for 4&9 with different number of labels.

Again the  $K$ -NN variant performs much better than the fully-connected graph. As in Table 6.3, we see that except for the pair 3&8, the classification error for the other three pairs with  $\Gamma_{\mathcal{Y}_n}$  is 48.0: after respecting the 40 labels, the other 960 images are classified as part of the same group. However, after regularization, the classification error is greatly reduced




Figure 6.4: Visualization of the regularization effects. The second row is the regularized version of the corresponding image in the first row. While arguably more blurred, the digits in the second row are more homogeneous within each group, making classification easier.

with  $\Gamma_{\bar{y}_n}$ . The same is true when we use the  $K$ -NN variant, but the improvement is smaller. Similarly as in Table 6.4, the improvement for local regularization becomes less dramatic as we go from the fully-connect graph to its  $K$ -NN variant and as the number of labels increases. This implies that there is certainly a limit for the improvement that local regularization can provide. Moreover, such improvement is most effective when label information is limited and one has to extract information from the geometry. Our theory and our experiments show that local regularization improves the recovery of geometric information and thereby boosts the classification performance in that scenario. We present a visualization of the effect of local regularization in Figure 6.3. The two rows represent the image before and after local regularization respectively. We can see that especially for the eights, many of the images get “fixed” after regularization. Moreover, at a high level, images within each group in the second row look more similar among themselves than those in the first row. Because of this we expect the classification to be better.

**Remark 6.4.4.** *The four chosen pairs of digits are the hardest pairs to classify but local regularization can improve the performance for other pairs too. For unsupervised spectral clustering, local regularization still gives improvement, but using cross validation to choose  $r$  is no longer possible.*

### 6.4.3 Future Directions

As mentioned above, the practical choice of  $r$  can be challenging. We propose two alternatives that may be easier to work with and investigate their competence on the MNIST dataset.



### 6.4.3.1 $k$ -NN regularization

This is a natural variant of  $\Gamma_{\bar{y}_n}$  based on  $k$ -nearest neighbor regularization. Instead of specifying a neighborhood of  $y_i$  of radius  $r$ , we simply regress the data by averaging over its  $k$  nearest neighbors. Here  $k$  is not necessarily the same as  $K$  (the number of neighbors used to construct a similarity graph). Conceptually, choosing  $k$  amounts to setting different values of  $r$  at different points in such a way that the resulting neighborhoods contain roughly the same number of points. This construction is easier to work with since  $k$  is in general easier to tune than  $r$ .

### 6.4.3.2 Self-tuning regularization

This is a global regularization variant that does not require hyper parameters. Instead of averaging over a neighborhood of radius  $r$ , we take a global weighted average of the whole point cloud, where the weights are proportional to the similarities between the  $y_i$ . More specifically, we define a new distance in terms of the points  $\hat{y}_i$ , where

$$\hat{y}_i = \sum_{j=1}^n W(i, j) y_j,$$

and  $W(i, j)$  is defined as in (6.31). We see that points far from  $y_i$  have small contribution in the definition of  $\hat{y}_i$  and so essentially one ends up summing over points in a neighborhood that is implicitly specified by the similarities. For points close to  $y_i$ , the weights are roughly on the same order. Hence  $\hat{y}_i$  can be seen approximately as  $\bar{y}_i$  plus a small contribution from points that are far from  $y_i$ . We expect this construction to behave a little worse than the  $\Gamma_{\bar{y}_n}$  with optimal  $r$ . However, the fact that this construction does not require the tuning of any hyper-parameter makes it an appealing choice. Table 6 compares the classification performance of all graphs mentioned above (with the four different choices of distance function, and the two alternatives to build similarity graphs).

Fully	3&8	5&8	4&9	7&9	$K$ -NN	3&8	5&8	4&9	7&9
$\Gamma_{\mathcal{Y}_n}$	27.7%	48.0%	48.0%	48.0%	$\Gamma_{\mathcal{Y}_n}$	7.6%	5.5%	12.8%	7.3%
$\Gamma_{\bar{\mathcal{Y}}_n}$	13.4%	17.4%	36.9%	15.3%	$\Gamma_{\bar{\mathcal{Y}}_n}$	6.9%	3.6%	9.7%	5.4%
$k$ -NN	11.5%	7.4%	43.1%	18.3%	$k$ -NN	5.3%	5.9%	9.6%	6.1%
self-tuning	16.1%	13.9%	33.4%	26.3%	self-tuning	7.6%	3.1%	8.8%	5.6%

Table 6.5: Comparison of classification errors with 4% labeled data.

**Remark 6.4.5.** *The idea of using labels to learn  $r$  (or  $k$ ) can be understood as a specific instance of a more general idea: to use labels to better inform the learning of the underlying geometry of a dataset. What is more, one can try to simultaneously learn the geometry of the input space with the learning of the labeling function, instead of looking at these two problems in sequential form. This will be the topic of future research.*

## 6.5 Appendix: Estimating $r_-$ and $r_+$

### 6.5.1 Estimating $r_-$

We want to find values of  $t > \frac{r}{2}$  for which for all  $v \in T_{x_i}\mathcal{M}$  with  $|v| \leq t$ , and for all  $\eta \in T_{\exp_{x_i}(v)}\mathcal{M}^\perp$  with  $|\eta| \leq \sigma$  we have

$$|\exp_{x_i}(v) + \eta - y_i| < r.$$

We will later take the maximum value of  $t$  for which this holds and set  $r_-$  to be this maximum value.

Let  $x = \exp_{x_i}(v)$ . First, with the parallel transport map used in the proof of the geometric bias estimates (as in (6.25)) we can associate a vector  $\tilde{\eta} \in T_x\mathcal{M}^\perp$  to a vector  $\eta \in T_{\exp_{x_i}(v)}\mathcal{M}^\perp$  with norm less than  $\sigma$ , for which

$$|\eta - \tilde{\eta}| \leq \frac{m}{R}\sigma t.$$

Now,

$$\begin{aligned}
|x + \eta - y_i| &\leq |x - x_i + \hat{\eta} - z_i| + |\eta - \hat{\eta}| \\
&= \left( |x - x_i|^2 + 2\langle x - x_i, \hat{\eta} - z_i \rangle + |\hat{\eta} - z_i|^2 \right)^{1/2} + |\eta - \hat{\eta}| \\
&\leq \left( |x - x_i|^2 + 2\langle x - x_i, \hat{\eta} - z_i \rangle + |\hat{\eta} - z_i|^2 \right)^{1/2} + \frac{m}{R}\sigma t.
\end{aligned}$$

We have

$$|x - x_i| \leq d_{\mathcal{M}}(x, x_i) = |v| \leq t,$$

and also

$$\langle x - x_i, \hat{\eta} - z_i \rangle = \langle x - (x_i + v), \hat{\eta} - z_i \rangle,$$

as it follows from the fact that  $\eta, z_i \in T_{x_i}\mathcal{M}^\perp$  and  $v \in T_{x_i}\mathcal{M}$ . Using this, Cauchy-Schwartz, and (6.20) we conclude that

$$|\langle x - (x_i + v), \hat{\eta} - z_i \rangle| \leq 2\sigma|x - (x_i + v)| \leq 2\sigma\frac{|v|^2}{R},$$

and hence

$$\begin{aligned}
|x + \eta - y_i| &\leq \left( t^2 + \frac{4}{R}\sigma t^2 + 4\sigma^2 \right)^{1/2} + \frac{m\sigma t}{R} \\
&= t \left( \sqrt{1 + \frac{4\sigma}{R} + \frac{4\sigma^2}{t^2}} + \frac{m\sigma}{R} \right) \\
&\leq t \left( \sqrt{1 + \frac{4\sigma}{R} + \frac{16\sigma^2}{r^2}} + \frac{m\sigma}{R} \right).
\end{aligned}$$

From the above it follows that  $r_-$  defined as

$$r_- := r \left( \sqrt{1 + \frac{4\sigma}{R} + \frac{16\sigma^2}{r^2}} + \frac{m\sigma}{R} \right)^{-1}, \quad (6.32)$$

satisfies the desired properties and moreover

$$r - r_- \leq r \left( 1 - \left( \sqrt{1 + \frac{4}{R}\sigma + \frac{16\sigma^2}{r^2} + \frac{m\sigma}{R}} \right)^{-1} \right). \quad (6.33)$$

### 6.5.2 Estimating $r_+$

To estimate  $r_+$ , we need the following lemma proved in García Trillos et al. [2019a].

**Lemma 6.5.1.** *Suppose  $x, \tilde{x} \in \mathcal{M}$  are such that  $|x - \tilde{x}| \leq R/2$ . Then*

$$|x - \tilde{x}| \leq d_{\mathcal{M}}(x, \tilde{x}) \leq |x - \tilde{x}| + \frac{8}{R}|x - \tilde{x}|^3. \quad (6.34)$$

To construct  $r_+$  we find values of  $t$  with  $2r \geq r + \sigma > t > 0$  such that if  $|v| > t$  then

$$|\exp_{x_i}(v) + \sigma\eta - y_i| \geq r$$

for all  $\eta \in T_{\exp_{x_i}(v)}\mathcal{M}^\perp$  of norm no larger than  $\sigma$ . As in the construction of  $r_-$  we let

$x := \exp_{x_i}(v)$ . Similar computations give

$$\begin{aligned}
|x + \eta - y_i| &\geq |x - x_i + \hat{\eta} - z_i| - |\eta - \hat{\eta}| \\
&= \left( |x - x_i|^2 + 2\langle x - x_i, \hat{\eta} - z_i \rangle + |\hat{\eta} - z_i|^2 \right)^{1/2} - |\eta - \hat{\eta}| \\
&\geq \left( |x - x_i|^2 + 2\langle x - x_i, \hat{\eta} - z_i \rangle + |\hat{\eta} - z_i|^2 \right)^{1/2} - \frac{m}{R}\sigma|v| \\
&\geq \left( \left( |v| - \frac{1}{R}|v|^3 \right)^2 - \frac{4}{R}\sigma|v|^2 \right)^{1/2} - \frac{m}{R}\sigma|v| \\
&\geq |v| \left( \sqrt{1 - \frac{2|v|^2}{R} - \frac{4\sigma}{R} - \frac{m\sigma}{R}} \right) \\
&\geq |v| \left( \sqrt{1 - \frac{8r^2}{R} - \frac{4\sigma}{R} - \frac{m\sigma}{R}} \right) \\
&\geq t \left( \sqrt{1 - \frac{8r^2}{R} - \frac{4\sigma}{R} - \frac{m\sigma}{R}} \right),
\end{aligned}$$

where in the third inequality we have used (6.34) to conclude that

$$|\exp_{x_i}(v) - x_i| \geq d_{\mathcal{M}}(x, x_i) - C(d_{\mathcal{M}}(x, x_i))^3 = |v| - C|v|^3.$$

We can then take  $t$  to be such that the right hand side of (6.33) is equal to  $r$ . That is, we can take

$$r_+ := r \left( \sqrt{1 - \frac{8r^2}{R} - \frac{4\sigma}{R} - \frac{m\sigma}{R}} \right)^{-1}.$$

From these estimates we see that

$$r_+ - r_- \leq c \left( r^3 + r\sigma + \frac{\sigma^2}{r} \right).$$

## REFERENCES

- E. Aamari and C. Levrard. Nonasymptotic rates for manifold, tangent space and curvature estimation. *The Annals of Statistics*, 47(1):177–204, 2019.
- E. Aamari, J. Kim, F. Chazal, B. Michel, A. Rinaldo, and L. Wasserman. Estimating the reach of a manifold. *Electronic Journal of Statistics*, 13(1):1359–1399, 2019.
- K. Abraham and R. Nickl. On statistical Calderón problems. *Mathematical Statistics and Learning*, 2(2):165–216, 2020.
- P. Abrahamsen. A review of Gaussian random fields and correlation functions. *Norsk Regnesentral/Norwegian Computing Center Oslo*, 1997.
- R. J. Adler. *The Geometry of Random Fields*, volume 62. SIAM, 2010.
- S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A. M. Stuart. Importance sampling: Intrinsic dimension and computational cost. *Statistical Science*, 32(3):405–431, 2017.
- E. B. Anderes and M. L. Stein. Estimating deformations of isotropic Gaussian random fields on the plane. *The Annals of Statistics*, 36(2):719–741, 2008.
- I. Babuska, R. Tempone, and G. E. Zouraris. Galerkin finite element approximations of stochastic elliptic partial differential equations. *SIAM Journal on Numerical Analysis*, 42(2):800–825, 2004.
- H. Bakka, H. Rue, A. Fuglstad, G.-A. and Riebler, D. Bolin, J. Illian, E. Krainski, D. Simpson, and F. Lindgren. Spatial modeling with R-INLA: A review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(6):e1443, 2018.
- S. Banerjee. On geodesic distance computations in spatial modeling. *Biometrics*, 61(2):617–625, 2005.
- S. Banerjee, A. E. Gelfand, A. O. Finley, and H. Sang. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):825–848, 2008.
- S. Banerjee, B. P. Carlin, and A. E. Gelfand. *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC, 2014.
- J. M. Bardeen, J. Bond, N. Kaiser, and A. Szalay. The statistics of peaks of Gaussian random fields. *The Astrophysical Journal*, 304:15–61, 1986.
- J. M. Bardsley and J. Kaipio. Gaussian Markov random field priors for inverse problems. *Inverse Problems & Imaging*, 7(2), 2013.
- M. Belkin and P. Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56(1-3):209–239, 2004.

- M. Belkin and P. Niyogi. Convergence of Laplacian eigenmaps. *Advances in Neural Information Processing Systems (NIPS)*, 19:129, 2007.
- M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008.
- M. Belkin, I. Matveeva, and P. Niyogi. Regularization and semi-supervised learning on large graphs. In *International Conference on Computational Learning Theory*, pages 624–638. Springer, 2004.
- M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7(Nov):2399–2434, 2006.
- T. Berry and J. Harlim. Variable bandwidth diffusion kernels. *Applied and Computational Harmonic Analysis*, 40(1):68–96, 2016.
- T. Berry and T. Sauer. Local kernels and the geometric structure of data. *Applied and Computational Harmonic Analysis*, 40(3):439–469, 2016.
- M. Bertalmío, L.-T. Cheng, S. Osher, and G. Sapiro. Variational problems and partial differential equations on implicit surfaces. *Journal of Computational Physics*, 174(2):759–780, 2001.
- A. L. Bertozzi, X. Luo, A. M. Stuart, and K. C. Zygalakis. Uncertainty quantification in graph-based classification of high dimensional data. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):568–595, 2018.
- A. L. Bertozzi, B. Hosseini, H. Li, K. Miller, and A. M. Stuart. Posterior consistency of semi-supervised regression on graphs. *Inverse Problems*, 37(10):105011, 2021.
- J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 36(2):192–225, 1974.
- J. Besag. Statistical analysis of non-lattice data. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 24(3):179–195, 1975.
- P. J. Bickel and B. Li. Local polynomial regression on unknown manifolds. In *Complex Datasets and Inverse Problems*, pages 177–186. Institute of Mathematical Statistics, 2007.
- D. Bigoni, Y. Chen, N. García Trillos, Y. Marzouk, and D. Sanz-Alonso. Data-driven forward discretizations for Bayesian inversion. *Inverse Problems*, 36(10):105008, 2020.
- V. I. Bogachev. *Gaussian Measures*. Number 62. American Mathematical Soc., 1998.
- D. Bolin. Spatial Matérn fields driven by non-Gaussian noise. *Scandinavian Journal of Statistics*, 41(3):557–579, 2014.

- D. Bolin and K. Kirchner. The rational SPDE approach for Gaussian random fields with general smoothness. *Journal of Computational and Graphical Statistics*, 29(2):274–285, 2020.
- D. Bolin and F. Lindgren. Spatial models generated by nested stochastic partial differential equations, with an application to global ozone mapping. *The Annals of Applied Statistics*, 5(1):523–550, 2011.
- D. Bolin and F. Lindgren. A comparison between Markov approximations and other methods for large spatial data sets. *Computational Statistics & Data Analysis*, 61:7–21, 2013.
- D. Bolin, K. Kirchner, and M. Kovács. Weak convergence of Galerkin approximations for fractional elliptic stochastic PDEs with spatial white noise. *BIT Numerical Mathematics*, 58(4):881–906, 2018.
- D. Bolin, K. Kirchner, and M. Kovács. Numerical solution of fractional elliptic stochastic PDEs with spatial white noise. *IMA Journal of Numerical Analysis*, 40(2):1051–1073, 2020.
- A. Bonito, J. M. Cascón, K. Mekchay, P. Morin, and R. H. Nochetto. High-order AFEM for the Laplace–Beltrami operator: Convergence rates. *Foundations of Computational Mathematics*, 16(6):1473–1539, 2016.
- M. Brennan, D. Bigoni, O. Zahm, A. Spantini, and Y. Marzouk. Greedy inference with structure-exploiting lazy maps. *Advances in Neural Information Processing Systems*, 33: 8330–8342, 2020.
- E. Brochu, V. M. Cora, and N. De Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- T. Bühler and M. Hein. Spectral clustering based on the graph p-Laplacian. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 81–88. ACM, 2009.
- T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
- L. Bungert, J. Calder, and T. Roith. Uniform convergence rates for Lipschitz learning on graphs. *arXiv preprint arXiv:2111.12370*, 2021.
- D. Burago, S. Ivanov, and Y. Kurylev. A graph discretization of the Laplace–Beltrami operator. *Journal of Spectral Theory*, 4(4):675–714, 2015.
- J. Calder and N. García Trillos. Improved spectral convergence rates for graph Laplacians on  $\varepsilon$ -graphs and  $k$ -nn graphs. *Applied and Computational Harmonic Analysis*, 2022.



- J. Calder and D. Slepčev. Properly-weighted graph Laplacian for semi-supervised learning. *Applied Mathematics & Optimization*, 82(3):1111–1159, 2020.
- J. Calder, D. Slepčev, and M. Thorpe. Rates of convergence for Laplacian semi-supervised learning with low labeling rates. *arXiv preprint arXiv:2006.02765*, 2020.
- J. Calder, N. García Trillos, and M. Lewicka. Lipschitz regularity of graph Laplacians on random data clouds. *SIAM Journal on Mathematical Analysis*, 54(1):1169–1222, 2022.
- D. Calvetti and E. Somersalo. *An Introduction to Bayesian Scientific Computing: Ten Lectures on Subjective Computing*, volume 2. Springer Science & Business Media, 2007.
- F. Camacho and A. Demlow.  $L^2$  and pointwise a posteriori error estimates for FEM for elliptic PDEs on surfaces. *IMA Journal of Numerical Analysis*, 35(3):1199–1227, 2015.
- M. Cameletti, F. Lindgren, D. Simpson, and H. Rue. Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *AStA Advances in Statistical Analysis*, 97(2):109–131, 2013.
- Y. Canzani. Analysis on manifolds via the Laplacian. *Lecture Notes available at: <http://www.math.harvard.edu/canzani/docs/Laplacian.pdf>*, 2013.
- I. Castillo, G. Kerkycharian, and D. Picard. Thomas Bayes’ walk on manifolds. *Probability Theory and Related Fields*, 158(3):665–710, 2014.
- O. Chapelle, B. Scholkopf, and A. Zien. Semi-supervised Learning. *Cambridge, Massachusetts: The MIT Press View Article*, 2, 2006.
- M. A. Chaplain, M. Ganesh, I. G. Graham, and G. Lolas. Mathematical modelling of solid tumour growth: applications of pre-pattern formation. In *Morphogenesis and Pattern Formation in Biological Systems*, pages 283–293. Springer, 2003.
- Y.-C. Chen, C. R. Genovese, and L. Wasserman. A comprehensive approach to mode clustering. *Electronic Journal of Statistics*, 10(1):210–241, 2016.
- F. R. K. Chung. *Spectral Graph Theory*. Number 92. American Mathematical Soc., 1997.
- F. S. Cohen, Z. Fan, and M. A. Patel. Classification of rotated and scaled textured images using Gaussian Markov random field models. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (2):192–202, 1991.
- R. Cohen, A. Devore and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDEs. *Analysis and Applications*, 9(01):11–47, 2011.
- R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.

- R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005.
- R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer. Graph Laplacian tomography from unknown random projections. *IEEE Transactions on Image Processing*, 17(10):1891–1899, 2008.
- J. A. Costa and A. O. Hero. Determining intrinsic dimension and entropy of high-dimensional shape spaces. In *Statistics and Analysis of Shapes*, pages 231–252. Springer, 2006.
- S. Cotter, M. Dashti, and A. M. Stuart. Approximation of Bayesian inverse problems for PDEs. *SIAM Journal on Numerical Analysis*, 48(1):322–345, 2010a.
- S. Cotter, M. Dashti, and A. M. Stuart. MCMC methods for functions: modifying old algorithms to make them faster. *SIAM Journal on Numerical Analysis*, 48(1):322–345, 2010b.
- T. Coulhon, G. Kerkycharian, and P. Petrushev. Heat kernel generated frames in the setting of Dirichlet spaces. *Journal of Fourier Analysis and Applications*, 18(5):995–1066, 2012.
- S. G. Cox and K. Kirchner. Regularity and convergence analysis in Sobolev and Hölder spaces for generalized Whittle–Matérn fields. *Numerische Mathematik*, 146(4):819–873, 2020.
- F. G. Cozman and I. Cohen. Risks of semi-supervised learning: How unlabeled data can degrade performance of generative classifiers. *Semi-supervised Learning*, 4:57–72, 2006.
- K. Crane. Keenan’s 3d model repository. URL <http://www.cs.cmu.edu/~kmcrane/Projects/ModelRepository>.
- T. Cui, K. J. Law, and Y. M. Marzouk. Dimension-independent likelihood-informed MCMC. *Journal of Computational Physics*, 304:109–137, 2016.
- M. Dashti and A. M. Stuart. Bayesian approach to inverse problems. *Handbook of Uncertainty Quantification*, pages 311–428, 2017.
- A. Datta, S. Banerjee, A. O. Finley, and A. E. Gelfand. Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812, 2016.
- E. B. Davies. *Spectral Theory and Differential Operators*. Number 42. Cambridge University Press, 1996.
- M. P. Do Carmo and J. Flaherty Francis. *Riemannian Geometry*, volume 6. Springer, 1992.

- H. Donnelly. Eigenfunctions of the Laplacian on compact Riemannian manifolds. *Asian Journal of Mathematics*, 10(1):115–126, 2006.
- J. Du, H. Zhang, and V. Mandrekar. Fixed-domain asymptotic properties of tapered maximum likelihood estimators. *The Annals of Statistics*, 37(6A):3330–3361, 2009.
- T. Dunker, M. Lifshits, and W. Linde. Small deviation probabilities of sums of independent random variables. In *High Dimensional Probability*, pages 59–74. Springer, 1998.
- M. M. Dunlop and A. M. Stuart. The Bayesian formulation of EIT: Analysis and algorithms. *Inverse Problems and Imaging*, 10(4):1007–1036, 2016.
- M. M. Dunlop, M. A. Iglesias, and A. M. Stuart. Hierarchical Bayesian level set inversion. *Statistics and Computing*, 27(6):1555–1584, 2017.
- D. B. Dunson, H.-T. Wu, and N. Wu. Spectral convergence of graph Laplacian and heat kernel reconstruction in  $L^\infty$  from random samples. *Applied and Computational Harmonic Analysis*, 55:282–336, 2021.
- G. Dziuk and C. M. Elliott. Finite element methods for surface PDEs. *Acta Numerica*, 22:289–396, 2013.
- D. E. Edmunds and H. Triebel. *Function Spaces, Entropy Numbers, Differential Operators*, volume 120. Cambridge University Press Cambridge, 1996.
- C. Eilks and C. M. Elliott. Numerical simulation of dealloying by surface dissolution via the evolving surface finite element method. *Journal of Computational Physics*, 227(23):9727–9741, 2008.
- A. El Alaoui, X. Cheng, A. Ramdas, M. J. Wainwright, and M. I. Jordan. Asymptotic behavior of  $L_p$ -based Laplacian regularization in semi-supervised learning. In *Conference on Learning Theory*, pages 879–906, 2016.
- C. M. Elliott and B. Stinner. A surface phase field model for two-phase biological membranes. *SIAM Journal on Applied Mathematics*, 70(8):2904–2928, 2010.
- H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.
- X. Fernique. Intégrabilité des vecteurs Gaussiens. *CR Acad. Sci. Paris Serie A*, 270:1698–1699, 1970.
- P. Frauenfelder, C. Schwab, and R. A. Todor. Finite elements for elliptic problems with stochastic coefficients. *Computer Methods in Applied Mechanics and Engineering*, 194(2-5):205–228, 2005.
- P. I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

- G.-A. Fuglstad, F. Lindgren, D. Simpson, and H. Rue. Exploring a new class of non-stationary spatial Gaussian random fields with varying local anisotropy. *Statistica Sinica*, pages 115–133, 2015a.
- G.-A. Fuglstad, D. Simpson, F. Lindgren, and H. Rue. Does non-stationary spatial data always require non-stationary random fields? *Spatial Statistics*, 14:505–531, 2015b.
- G.-A. Fuglstad, D. Simpson, F. Lindgren, and H. Rue. Constructing priors that penalize the complexity of Gaussian random fields. *Journal of the American Statistical Association*, 114(525):445–452, 2019.
- K. Fukunaga and L. Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on Information Theory*, 21(1):32–40, 1975.
- R. Furrer, M. G. Genton, and D. Nychka. Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523, 2006.
- T. Gao, S. Z. Kovalsky, and I. Daubechies. Gaussian process landmarking on manifolds. *SIAM Journal on Mathematics of Data Science*, 1(1):208–236, 2019.
- N. García Trillos and D. Sanz-Alonso. The Bayesian formulation and well-posedness of fractional elliptic inverse problems. *Inverse Problems*, 33(6):065006, 2017.
- N. García Trillos and D. Sanz-Alonso. Continuum limits of posteriors in graph Bayesian inverse problems. *SIAM Journal on Mathematical Analysis*, 50(4):4020–4040, 2018.
- N. García Trillos and D. Slepčev. Continuum limit of total variation on point clouds. *Archive for rational mechanics and analysis*, 220(1):193–241, 2016a.
- N. García Trillos and D. Slepčev. Continuum limit of total variation on point clouds. *Archive for Rational Mechanics and Analysis*, 220(1):193–241, 2016b.
- N. García Trillos, M. Gerlach, M. Hein, and D. Slepčev. Error estimates for spectral convergence of the graph Laplacian on random geometric graphs toward the Laplace–Beltrami operator. *Foundations of Computational Mathematics*, pages 1–61, 2019a.
- N. García Trillos, D. Sanz-Alonso, and R. Yang. Local regularization of noisy point clouds: Improved global geometric estimates and data analysis. *Journal of Machine Learning Research*, 20(136):1–37, 2019b.
- N. García Trillos, M. Gerlach, M. Hein, and D. Slepčev. Error estimates for spectral convergence of the graph Laplacian on random geometric graphs toward the Laplace–Beltrami operator. *Foundations of Computational Mathematics*, 20(4):827–887, 2020a.
- N. García Trillos, Z. Kaplan, T. Samakhoana, and D. Sanz-Alonso. On the consistency of graph-based Bayesian semi-supervised learning and the scalability of sampling algorithms. *Journal of Machine Learning Research*, 21(28):1–47, 2020b.

- G. Gaspari and S. E. Cohn. Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757, 1999.
- A. E. Gelfand, P. Diggle, P. Guttorp, and M. Fuentes. *Handbook of Spatial Statistics*. CRC press, 2010.
- C. Genovese, M. Perone-Pacifco, I. Verdinelli, and L. Wasserman. Minimax manifold estimation. *Journal of Machine Learning Research*, 13(May):1263–1291, 2012.
- C. J. Geoga, M. Anitescu, and M. L. Stein. Scalable Gaussian process computations using hierarchical matrices. *Journal of Computational and Graphical Statistics*, pages 1–11, 2019.
- S. Ghosal and A. Van Der Vaart. Convergence rates of posterior distributions for noniid observations. *The Annals of Statistics*, 35(1):192–223, 2007.
- S. Ghosal and A. Van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*, volume 44. Cambridge University Press, 2017.
- S. Ghosal, J. K. Ghosh, and A. Van Der Vaart. Convergence rates of posterior distributions. *The Annals of Statistics*, pages 500–531, 2000.
- F. Gilani and J. Harlim. Approximating solutions of linear elliptic PDE’s on a smooth manifold using local kernel. *Journal of Computational Physics*, 2019.
- D. Gilbarg and N. S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Springer, 2015.
- E. Giné and V. Koltchinskii. Empirical graph Laplacian approximation of Laplace–Beltrami operators: Large sample results. In *High Dimensional Probability*, pages 238–259. Institute of Mathematical Statistics, 2006.
- M. Giordano and R. Nickl. Consistency of Bayesian inference with Gaussian process priors in an elliptic inverse problem. *Inverse Problems*, 36(8):085001, 2020.
- T. Gneiting, A. E. Raftery, A. H. Westveld III, and T. Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133(5):1098–1118, 2005.
- R. B. Gramacy and H. K. H. Lee. Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103(483):1119–1130, 2008.
- P. Greengard and M. O’Neil. Efficient reduced-rank methods for Gaussian processes with eigenfunction expansions. *arXiv preprint arXiv:2108.05924*, 2021.
- P. Guttorp and T. Gneiting. Studies in the history of probability and statistics XLIX: On the Matérn correlation family. *Biometrika*, 93(4):989–995, 2006.

- A. Haddad, D. Kushnir, and R. R. Coifman. Texture separation via a reference set. *Applied and Computational Harmonic Analysis*, 36(2):335–347, 2014.
- S. Harizanov, R. Lazarov, S. Margenov, P. Marinov, and Y. Vutov. Optimal solvers for linear systems with fractional powers of sparse spd matrices. *Numerical Linear Algebra with Applications*, 25(5):e2167, 2018.
- J. Harlim, D. Sanz-Alonso, and R. Yang. Kernel methods for Bayesian elliptic inverse problems on manifolds. *SIAM/ASA Journal on Uncertainty Quantification*, 8(4):1414–1445, 2020.
- M. J. Heaton, A. Datta, A. O. Finley, R. Furrer, J. Guinness, F. Guhaniyogi, R. and Gerber, R. B. Gramacy, D. Hammerling, M. Katzfuss, et al. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24(3):398–425, 2019.
- M. Hein. Uniform convergence of adaptive graph-based regularization. In *International Conference on Computational Learning Theory*, pages 50–64. Springer, 2006.
- M. Hein and J.-Y. Audibert. Intrinsic dimensionality estimation of submanifolds in  $\mathbb{R}^d$ . In *Proceedings of the 22nd International Conference on Machine Learning*, pages 289–296, 2005.
- M. Hein and M. Maier. Manifold denoising. *Advances in Neural Information Processing Systems*, 19, 2006.
- M. Hein, J.-Y. Audibert, and U. Von Luxburg. From graphs to manifolds—weak and strong pointwise consistency of graph Laplacians. In *International Conference on Computational Learning Theory*, pages 470–485. Springer, 2005.
- M. Hein, J.-Y. Audibert, and U. v. Luxburg. Graph Laplacians and their convergence on random neighborhood graphs. *Journal of Machine Learning Research*, 8(6), 2007.
- P. Hennig, M. A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2179):20150142, 2015.
- F. Hoffmann, B. Hosseini, Z. Ren, and A. M. Stuart. Consistency of semi-supervised learning algorithms on graphs: Probit and one-hot methods. *Journal of Machine Learning Research*, 21:1–55, 2020.
- P. L. Houtekamer and H. L. Mitchell. A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129(1):123–137, 2001.
- M. A. Iglesias and C. Dawson. The representer method for state and parameter estimation in single-phase darcy flow. *Computer Methods in Applied Mechanics and Engineering*, 196(45-48):4577–4596, 2007.

- T. Isaac, N. Petra, G. Stadler, and O. Ghattas. Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet. *Journal of Computational Physics*, 296:348–368, 2015.
- J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*. Springer Science & Business Media, 160, 2006.
- M. Kang and M. Katzfuss. Correlation-based sparse inverse Cholesky factorization for fast Gaussian-process inference. *arXiv preprint arXiv:2112.14591*, 2021.
- M. Katzfuss and J. Guinness. A general framework for Vecchia approximations of Gaussian processes. *Statistical Science*, 36(1):124–141, 2021.
- M. C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- H. Kersting and P. Hennig. Active uncertainty calibration in Bayesian ode solvers. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pages 309–318, 2016.
- U. Khristenko, L. Scarabosio, P. Swierczynski, E. Ullmann, and B. Wohlmuth. Analysis of boundary effects on PDE-based sampling of Whittle–Matérn random fields. *SIAM/ASA Journal on Uncertainty Quantification*, 7(3):948–974, 2019.
- H.-M. Kim, B. K. Mallick, and C. C. Holmes. Analyzing nonstationary spatial data using piecewise Gaussian processes. *Journal of the American Statistical Association*, 100(470):653–668, 2005.
- A. Kirichenko and H. van Zanten. Estimating a smooth function on a large graph by Bayesian Laplacian regularisation. *Electronic Journal of Statistics*, 11(1):891–915, 2017.
- R. I. Kondor and J. Lafferty. Diffusion kernels on graphs and other discrete structures. In *Proceedings of the 19th International Conference on Machine Learning*, volume 2002, pages 315–322, 2002.
- A. Kushpel and J. Levesley. Entropy of Sobolev’s classes on compact homogeneous Riemannian manifolds. *arXiv preprint arXiv:1504.06508*, 2015.
- A. Lang, J. Potthoff, M. Schlather, and D. Schwab. Continuity of random fields on Riemannian manifolds. *Communications on Stochastic Analysis*, 10(2):4, 2016.
- Y. LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- W. V. Li and W. Linde. Approximation, metric entropy and small ball estimates for Gaussian measures. *The Annals of Probability*, 27(3):1556–1578, 1999.

- Y. Li, B. Mark, G. Raskutti, and R. Willett. Graph-based regularization for regression problems with highly-correlated designs. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 740–742. IEEE, 2018.
- Z. Li and Z. Shi. A convergent point integral method for isotropic elliptic equations on a point cloud. *Multiscale Modeling & Simulation*, 14(2):874–905, 2016.
- Z. Li, Z. Shi, and J. Sun. Point integral method for solving Poisson-type equations on manifolds from point clouds with convergence guarantees. *Communications in Computational Physics*, 22(1):228–258, 2017.
- F. Liang, S. Mukherjee, and M. West. The use of unlabeled data in predictive modeling. *Statistical Science*, pages 189–205, 2007.
- E. H. Lieb and M. Loss. *Analysis*, volume 14. American Mathematical Soc., 2001.
- F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.
- F. Lindgren, D. Bolin, and H. Rue. The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running. *Spatial Statistics*, page 100599, 2022.
- A. Lischke, G. Pang, M. Gulian, F. Song, C. Glusa, X. Zheng, Z. Mao, W. Cai, M. M. Meerschaert, M. Ainsworth, et al. What is the fractional Laplacian? A comparative review with new results. *Journal of Computational Physics*, 404:109009, 2020.
- A. V. Little and L. Maggioni, M. and Rosasco. Multiscale geometric methods for data sets I: Multiscale SVD, noise and curvature. *Applied and Computational Harmonic Analysis*, 43(3):504–567, 2017.
- F. Liu, S. Chakraborty, F. Li, Y. Liu, and A. C. Lozano. Bayesian regularization via graph Laplacian. *Bayesian Analysis*, 9(2):449–474, 2014.
- R. J. Lorentzen, K. M. Flornes, G. Nævdal, et al. History matching channelized reservoirs using the ensemble Kalman filter. *SPE Journal*, 17(01):137–151, 2012.
- D. J. C. MacKay. Gaussian processes—a replacement for supervised neural networks? *NIPS tutorial*, 1997.
- J. D. Martin and T. W. Simpson. Use of kriging models to approximate deterministic computer models. *AIAA journal*, 43(4):853–863, 2005.
- Y. Marzouk and D. Xiu. A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics*, 6(4):826–847, 2009.
- B. Matérn. *Spatial Variation*, volume 36. Second Edition, Springer Science & Business Media, 2013.



- D. McLaughlin and L. R. Townley. A reassessment of the groundwater inverse problem. *Water Resources Research*, 32(5):1131–1161, 1996.
- F. Mémoli, G. Sapiro, and P. Thompson. Implicit brain imaging. *NeuroImage*, 23:S179–S188, 2004.
- F. Memoli, Z. Smith, and Z. Wan. The Wasserstein transform. In *International Conference on Machine Learning*, pages 4496–4504. PMLR, 2019.
- F. Monard, R. Nickl, and G. P. Paternain. Consistent inversion of noisy non-Abelian X-ray transforms. *Communications on Pure and Applied Mathematics*, 74(5):1045–1099, 2021.
- S. Montagna and S. T. Tokdar. Computer emulation with nonstationary Gaussian processes. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):26–47, 2016.
- K. Monterrubio-Gómez, L. Roininen, S. Wade, T. Damoulas, and M. Girolami. Posterior inference for sparse hierarchical non-stationary models. *Computational Statistics & Data Analysis*, 148:106954, 2020.
- J. L. Mueller and S. Siltanen. *Linear and nonlinear inverse problems with practical applications*, volume 10. SIAM, 2012.
- B. Nadler, N. Srebro, and X. Zhou. Semi-supervised learning with the graph Laplacian: The limit of infinite unlabelled data. *Advances in Neural Information Processing Systems*, 22:1330–1338, 2009.
- R. Neal. Regression and classification using Gaussian process priors. *Bayesian Statistics*, 6, 1998.
- A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems*, pages 849–856, 2002.
- Y. C. Ng, N. Colombo, and R. Silva. Bayesian semi-supervised learning with graph Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 1683–1694, 2018.
- R. Nickl, S. van de Geer, and S. Wang. Convergence rates for penalized least squares estimators in PDE constrained regression problems. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):374–413, 2020.
- L. I. Nicolaescu. *Lectures on the Geometry of Manifolds*. World Scientific, 2020.
- P. Niyogi. Manifold regularization and semi-supervised learning: Some theoretical analyses. *Journal of Machine Learning Research*, 14(5), 2013.
- P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.

- P. J. Olver. *Introduction to Partial Differential Equations*. Springer.
- H. Owhadi and C. Scovel. *Operator-Adapted Wavelets, Fast Solvers, and Numerical Homogenization: From a Game Theoretic Approach to Numerical Approximation and Algorithm Design*, volume 35. Cambridge University Press, 2019.
- G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- I. Pinelis. An approach to inequalities for the distributions of infinite-dimensional martingales. In *Probability in Banach Spaces, 8: Proceedings of the Eighth International Conference*, pages 128–134. Springer, 1992.
- J. Ping and D. Zhang. History matching of channelized reservoirs with vector-based level-set parameterization. *SPE Journal*, 19(03):514–529, 2014.
- C. Piret. The orthogonal gradients method: A radial basis functions method for solving partial differential equations on arbitrary surfaces. *Journal of Computational Physics*, 231(14):4662–4675, 2012.
- Y. Qiao, C. Shi, C. Wang, H. Li, M. Haberland, X. Luo, A. M. Stuart, and A. L. Bertozzi. Uncertainty quantification for semi-supervised multi-class classification in image processing and ego-motion analysis of body-worn videos. *Electronic Imaging*, 2019(11):264–1, 2019.
- J. Quinero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Machine learning of linear differential equations using Gaussian processes. *Journal of Computational Physics*, 348:683–693, 2017.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Numerical Gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 40(1):A172–A198, 2018.
- L. Roininen, J. M. Huttunen, and S. Lasanen. Whittle-Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography. *Inverse Problems & Imaging*, 8(2):561, 2014.
- L. Roininen, M. Girolami, S. Lasanen, and M. Markkanen. Hyperpriors for Matérn fields with applications in Bayesian inversion. *Inverse Problems & Imaging*, 13(1):1–29, 2019.
- H. Rue and L. Held. *Gaussian Markov Random Fields: Theory and Applications*. Chapman and Hall/CRC, 2005.
- H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):319–392, 2009.

- S. J. Ruuth and B. Merriman. A simple embedding method for solving partial differential equations on surfaces. *Journal of Computational Physics*, 227(3):1943–1961, 2008.
- P. D. Sampson, D. Damian, and P. Guttorp. Advances in modeling and inference for environmental processes with nonstationary spatial covariance. In *GeoENV III—Geostatistics for Environmental Applications*, pages 17–32. Springer, 2001.
- X. Sanchez-Vila, A. Guadagnini, and J. Carrera. Representative hydraulic conductivities in saturated groundwater flow. *Reviews of Geophysics*, 44(3), 2006.
- D. Sanz-Alonso. Importance sampling and necessary sample size: An information theory approach. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):867–879, 2018.
- D. Sanz-Alonso and R. Yang. The SPDE approach to Matérn fields: Graph representations. *To appear in Statistical Science*, 2022a.
- D. Sanz-Alonso and R. Yang. Unlabeled data help in graph-based semi-supervised learning: a Bayesian nonparametrics perspective. *To appear in Journal of Machine Learning Research*, 2022b.
- D. Sanz-Alonso and R. Yang. Finite element representations of Gaussian processes: Balancing numerical and statistical accuracy. *To appear in SIAM/ASA Journal on Uncertainty Quantification*, 2022c.
- D. Sanz-Alonso, A. M. Stuart, and A. Taeb. Inverse Problems and Data Assimilation. *arXiv preprint arXiv:1810.06191*, 2019.
- F. Schäfer, M. Katzfuss, and H. Owhadi. Sparse Cholesky factorization by Kullback–Leibler minimization. *SIAM Journal on Scientific Computing*, 43(3):A2019–A2046, 2021.
- M. Seeger. Learning with labeled and unlabeled data. Technical report, 2000a.
- M. Seeger. Relationships between Gaussian processes, support vector machines and smoothing splines. *Machine Learning*, 2000b.
- X. Shen and L. Wasserman. Rates of convergence of posterior distributions. *The Annals of Statistics*, 29(3):687–714, 2001.
- J. Shi and J. Malik. Normalized cuts and image segmentation. *Departmental Papers (CIS)*, page 107, 2000.
- L. Shi, R. Mihalcea, and M. Tian. Cross language text classification by model translation and semi-supervised learning. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1057–1067, 2010.
- Z. Shi and J. Sun. Convergence of the point integral method for Poisson equation on point cloud. *arXiv preprint arXiv:1403.2141*, 2016.

- D. Simpson, F. Lindgren, and H. Rue. Think continuous: Markovian Gaussian models in spatial statistics. *Spatial Statistics*, 1:16–29, 2012.
- A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.
- A. Singer and H.-T. Wu. Spectral convergence of the connection Laplacian from random samples. *Information and Inference: A Journal of the IMA*, 6(1):58–123, 2017.
- A. Singh, R. Nowak, and J. Zhu. Unlabeled data: Now it helps, now it doesn't. *Advances in Neural Information Processing Systems*, 21, 2008.
- A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing*, 30(2):419–446, 2020.
- P. Sollich. Bayesian methods for support vector machines: Evidence and predictive class probabilities. *Machine Learning*, 46(1-3):21–52, 2002.
- E. Somersalo, M. Cheney, and D. Isaacson. Existence and uniqueness for electrode models for electric current computed tomography. *SIAM Journal on Applied Mathematics*, 52(4):1023–1040, 1992.
- D. A. Spielman and S.-H. Teng. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421(2-3):284–305, 2007.
- V. Stathopoulos, V. Zamora-Gutierrez, K. Jones, and M. Girolami. Bat call identification with Gaussian process multinomial probit regression and a dynamic time warping kernel. In *Artificial Intelligence and Statistics*, pages 913–921, 2014.
- M. L. Stein. Bounds on the efficiency of linear predictions using an incorrect covariance function. *The Annals of Statistics*, pages 1116–1138, 1990a.
- M. L. Stein. Uniform asymptotic optimality of linear predictions of a random field using an incorrect second-order structure. *The Annals of Statistics*, pages 850–872, 1990b.
- M. L. Stein. A simple condition for asymptotic optimality of linear predictions of random fields. *Statistics & Probability Letters*, 17(5):399–404, 1993.
- M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Science & Business Media, 1999a.
- M. L. Stein. Predicting random fields with increasing dense observations. *The Annals of Applied Probability*, 9(1):242–273, 1999b.
- M. L. Stein. The screening effect in kriging. *The Annals of Statistics*, 30(1):298–323, 2002.
- G. Strang and G. J. Fix. *An Analysis of the Finite Element Method*. Englewood Cliffs, N. J., Prentice-Hall, Inc., 1973. 318 p, 1973.

- A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- A. M. Stuart and A. Teckentrup. Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Mathematics of Computation*, 87(310):721–753, 2018.
- T. J. Sullivan. *Introduction to Uncertainty Quantification*, volume 63. Springer, 2015.
- W. Tao and Z. Shi. Convergence of Laplacian spectra from random samples. *Journal of Computational Mathematics*, 38(6):952–984, 2020.
- J. E. Taylor and K. J. Worsley. Detecting sparse signals in random fields, with an application to brain mapping. *Journal of the American Statistical Association*, 102(479):913–928, 2007.
- E. H. Thiede, D. Giannakis, A. R. Dinner, and J. Weare. Galerkin approximation of dynamical quantities using trajectory data. *The Journal of Chemical Physics*, 150(24):244111, 2019.
- D. Ting, L. Huang, and M. I. Jordan. An analysis of the convergence of graph Laplacians. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1079–1086, 2010.
- J. A. Toth and S. Zelditch. Riemannian manifolds with uniformly bounded eigenfunctions. *Duke Mathematical Journal*, 111(1):97–132, 2002.
- H. Triebel. *Theory of Function Spaces II*. Monographs in Mathematics. Springer Basel, 1992.
- A. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York, 2008.
- J. W. Tukey and P. A. Tukey. Computer graphics and exploratory data analysis: An introduction. *The Collected Works of John W. Tukey: Graphics: 1965-1985*, 5:419, 1988.
- A. van der Vaart and H. van Zanten. Rates of contraction of posterior distributions based on Gaussian process priors. *The Annals of Statistics*, 36(3):1435–1463, 2008a.
- A. van der Vaart and H. van Zanten. Reproducing kernel Hilbert spaces of Gaussian priors. In *Pushing the limits of contemporary statistics: contributions in honor of Jayanta K. Ghosh*, pages 200–222. Institute of Mathematical Statistics, 2008b.
- A. van der Vaart and H. van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12(60):2095–2119, 2011.
- J. E. Van Engelen and H. H. Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020.
- A. V. Vecchia. Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 50(2):297–312, 1988.

- U. Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.
- U. Von Luxburg, M. Belkin, and O. Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, pages 555–586, 2008.
- D. Wang and W.-L. Loh. On fixed-domain asymptotics and covariance tapering in Gaussian random field models. *Electronic Journal of Statistics*, 5:238–269, 2011.
- W. Wang and M. A. Carreira-Perpinán. Manifold blurring mean shift algorithms for manifold denoising. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1759–1766. IEEE, 2010.
- L. Wasserman. *All of Nonparametric Statistics*. Springer Science & Business Media, 2006.
- L. Wasserman and J. Lafferty. Statistical analysis of semi-supervised regression. *Advances in Neural Information Processing Systems*, 20, 2007.
- L. Wasserman and J. D. Lafferty. Statistical analysis of semi-supervised regression. In *Advances in Neural Information Processing Systems*, pages 801–808, 2008.
- J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- J. Weston, C. Leslie, E. Ie, D. Zhou, A. Elisseeff, and W. S. Noble. Semi-supervised protein classification using cluster kernels. *Bioinformatics*, 21(15):3241–3247, 2005.
- P. Whittle. On stationary processes in the plane. *Biometrika*, pages 434–449, 1954.
- P. Whittle. Stochastic-processes in several dimensions. *Bulletin of the International Statistical Institute*, 40(2):974–994, 1963.
- A. Wiens, D. Nychka, and W. Kleiber. Modeling spatial data using local likelihood estimation and a Matérn to spatial autoregressive translation. *Environmetrics*, 31(6):e2652, 2020.
- C. K. I. Williams and C. E. Rasmussen. Gaussian processes for regression. In *Advances in Neural Information Processing Systems*, pages 514–520, 1996.
- C. K. I. Williams and C. E. Rasmussen. *Gaussian Processes for Machine Learning*, volume 2. MIT press Cambridge, MA, 2006.
- B. Xu. Asymptotic behavior of  $L^2$ -normalized eigenfunctions of the Laplace-Beltrami operator on a closed Riemannian manifold. *Harmonic Analysis and its Applications*, pages 99–117, 2006.
- J.-J. Xu, Z. Li, J. Lowengrub, and H. Zhao. A level-set method for interfacial flows with surfactant. *Journal of Computational Physics*, 212(2):590–616, 2006.

- Y. Yang and D. B. Dunson. Bayesian manifold regression. *The Annals of Statistics*, 44(2): 876–905, 2016.
- L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Advances in Neural Information Processing Systems*, pages 1601–1608, 2005.
- D. Zhou and B. Schölkopf. Regularization on discrete spaces. In *Joint Pattern Recognition Symposium*, pages 361–368. Springer, 2005.
- X. Zhu. *Semi-supervised learning with graphs*. Carnegie Mellon University, 2005.
- X. Zhu, Z. Ghahramani, and J. D. Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the 20th International Conference on Machine Learning*, pages 912–919, 2003.