THE UNIVERSITY OF CHICAGO

NON-ASYMPTOTIC STATISTICAL ANALYSIS OF ENSEMBLE BASED FILTERING ALGORITHMS

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

DEPARTMENT OF STATISTICS

BY
OMAR AL-GHATTAS

CHICAGO, ILLINOIS
JUNE 2025

ABSTRACT

At its core, this dissertation aims to formalize and explain—through a statistical lens—the empirical success of popular ensemble-based algorithms in the data assimilation literature. A key component of this effort is the derivation of non-asymptotic, dimension-free bounds for the estimation of covariance operators. To achieve this, we leverage existing techniques from high-dimensional probability while also developing new theoretical tools to analyze the behavior of a certain class of covariance estimators under structural assumptions. This dissertation rigorously establishes fundamental guarantees for these estimators, shedding light on the mechanisms that drive their effectiveness and providing a deeper understanding of their practical success.

To Nadereh, Mae, and Leila.

My mother, my wife, and my daughter,

the three people I cherish most.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

To begin with, I would like to thank my supervisor, Daniel Sanz-Alonso. It is difficult to adequately and eloquently capture in a few short sentences the impact that Daniel has had on my development as a researcher. From our first meetings it was clear to me that Daniel is a dedicated and thoughtful mentor who is genuinely invested in the growth of his students. Over the years, I have greatly valued his generosity with his time, his ability to motivate me when my enthusiasm wavered, and his reassurance during the inevitable challenges of the PhD. For these reasons – and many more that I could not hope to list here – I am truly grateful and proud to have been one of his students.

Throughout my PhD, I've been fortunate to collaborate with exceptional researchers – Jiajun Bao, Jiaheng Chen, and Nathan Waniorek. I've learned a great deal (and continue to) from our joint work, and I've deeply valued our partnership.

I am thankful to Rina Foygel Barber and Rebecca Willett for agreeing to be on my thesis committee, and for their advice and feedback about my work.

My first formal introduction to high-dimensional statistics came through a series of masterfully delivered courses by Chao Gao. Beyond laying a rigorous foundation, he expertly surveyed recent advances in the field – insights that profoundly shaped my research direction. I am very glad to have had Chao as one of my teachers. I am also indebted to Yandi Shen for introducing me to the literature on dimension-free covariance estimation.

I am indebted to Max Farrell for his support and generosity throughout my PhD. Max not only introduced me to foundational works in econometrics and statistics—expanding my perspective significantly—but also offered invaluable guidance through countless insightful discussions over the years. His mentorship has been instrumental to my growth, and I sincerely appreciate his time and wisdom.

I am deeply grateful to Youssef Marzouk for hosting me as a visiting student at MIT during the final year of my PhD. Youssef and his Uncertainty Quantification group extended extraordinary hospitality to me and my family, helping us settle comfortably in a new environment. Being part of their intellectual community has been both an honor and a pleasure. I would also like to thank Philippe Rigollet for his valuable research guidance and career advice and for generously including me in his group. I would like to especially thank Aimee Maurais, Tudor Manole and Matthew Li from the MIT community for their friendship.

Finally, I am very lucky to have formed two friendships with YoonHaeng Hur and Subhodh Kotekal. Looking back, many of the seemingly insignificant everyday moments – stressing over problem set deadlines in our first year office, Zoom chats during the COVID years, the countless coffees, lunches, and dinners, and the endless blackboard discussions – have become some of my most treasured memories of the PhD journey. Their friendship made my time in Chicago truly special.

CHAPTER 1

INTRODUCTION

This thesis advances two research areas: (1) the statistical analysis of ensemble-based algorithms and (2) the estimation of covariance operators with sparse structure. This section aims to provide motivation for the thesis by introducing the concept of an ensemble Kalman update and the challenges inherent in its statistical analysis. We then emphasize the significance of structured covariance operator estimation as a crucial element in understanding ensemble-based algorithms. Finally, we offer a high-level overview of key technical tools from empirical process theory that are utilized throughout the thesis.

1.1 Ensemble Kalman Algorithms

Many algorithms for inverse problems and data assimilation rely on ensemble Kalman updates to blend prior predictions with observed data. As a motivating example, consider the inverse problem of recovering $u \in \mathbb{R}^d$ from data $y \in \mathbb{R}^k$, corrupted by noise η, where

$$y = \mathcal{G}(u) + \eta, \tag{1.1}$$

$\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^k$ is the forward model, and $\eta \sim \mathbb{P}_\eta = \mathcal{N}(0, \Gamma)$ is the observation error with positive-definite covariance matrix Γ. An ensemble Kalman update takes as input a *prior ensemble* $\{u_n\}_{n=1}^N$ and observed data y, and returns as output an *updated ensemble* $\{\upsilon_n\}_{n=1}^N$ that blends together the information in the prior ensemble and in the newly observed data.

1.1.1 Single-Step Ensemble Kalman Update

Throughout this work, we explore different notions of recovery, but as an illustrative example in this introduction, we focus on a specific type, which we call *posterior-approximation*.

Specifically, when the forward model is linear, $\mathcal{G}(u) = Au$, with ill-conditioned A or $d \gg k$, naive inversion amplifies small observation errors into large reconstruction errors. Regularization stabilizes the solution, and a Bayesian approach achieves this by placing a Gaussian prior $u \sim \mathcal{N}(m, C)$, where C acts as a probabilistic regularizer. The posterior $\mathbb{P}_{u|y}$ is Gaussian, $\mathcal{N}(\mu, \Sigma)$, with

$$\mu = m + CA^\top (ACA^\top + \Gamma)^{-1}(y - Am), \tag{1.2}$$

$$\Sigma = C - CA^\top (ACA^\top + \Gamma)^{-1} AC. \tag{1.3}$$

These require storing $d \times d$ matrices, making computation infeasible for large d. Instead, an ensemble Kalman update transforms a prior ensemble $\{u_n\}_{n=1}^N$ into an updated ensemble $\{\upsilon_n\}_{n=1}^N$, whose sample mean and covariance approximate those of $\mathbb{P}_{u|y}$. We refer to such methods as posterior-approximation algorithms. Numerous such algorithms exist in the literature, with the Perturbed Observations (PO) and Square-root Filter (SR) updates being among the most popular. For example, the PO update transforms each particle of the prior ensemble according to

$$\upsilon_n = u_n + \widehat{C}A^\top (A\widehat{C}A^\top + \Gamma)^{-1}\big(y - Au_n - \eta_n\big), \quad 1 \le n \le N,$$

where $\widehat{C}$ denotes the empirical covariance matrix of the prior ensemble and $\{\eta_n\}_{n=1}^N$ are i.i.d. copies of the noise variable η. The PO update is therefore a (perturbed) Monte-Carlo estimate of the true posterior update. One then obtains estimates of the posterior parameters (μ, Σ) by using the sample mean and covariance of the updated ensemble, denoted $(\hat{\mu}_{\text{PO}}, \widehat{\Sigma}_{\text{PO}})$. In the SR update, we instead obtain the estimators $(\hat{\mu}_{\text{SR}}, \widehat{\Sigma}_{\text{SR}})$.

The primary motivation for ensemble Kalman methods is their ability to perform well even with a small ensemble size N, which is critical in applications where generating each particle is computationally expensive. Most theoretical studies have focused on large-ensemble asymptotics, examining the limit as $N \to \infty$. While these *mean-field* results are mathematically insightful and have led to practical advancements, they do not fully explain the observed success of ensemble Kalman methods when used with small ensemble sizes. Moreover, asymptotic analyses cannot differentiate between different algorithms. For example, in the limit $N \to \infty$, the SR and PO updates are equivalent, even though they exhibit markedly different behaviors in practice. Finally, in many practical scenarios, the state dimension d is extremely large—often much larger than N or even infinite—making it essential for any theoretical framework to account for the challenges introduced by high-dimensional state spaces.

In this thesis, we propose a novel analysis of the error of ensemble Kalman updates under a variety of instantiations. Concretely, we provide high-probability and in-expectation upper bounds of the form $\|\hat{\mu}-\mu\|_2 \leq \varepsilon_\mu$ whenever $N \gtrsim N_0$ and similarly for the covariance deviation in operator norm. Here, N_0 is a quantity that depends on the problem specific parameters and error level ε_μ, but need not depend on the intrinsic dimension d. We therefore refer to our bounds as being dimension free. The non-asymptotic nature of our results also allow us to distinguish between PO and SR updates.

1.1.2 Multi-Step Ensemble Kalman Update

Ensemble Kalman updates are often employed to solve filtering problems which arise in the data assimilation literature. Here, the goal is to estimate a time-evolving state from partial and noisy observations. To make things concrete, we consider here the following linear version of the hidden Markov model governing the relationship between the state and observation processes:

$$
\begin{aligned}
&\text{(Initialization)} \qquad u^{(0)} \sim \mathcal{N}(\mu^{(0)}, \Sigma^{(0)}), && && (1.4)\\
&\text{(Dynamics)} \qquad u^{(j)} = Au^{(j-1)} + \xi^{(j)}, && \xi^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Xi), && j = 1, 2, \ldots \quad (1.5)\\
&\text{(Observation)} \qquad y^{(j)} = Hu^{(j)} + \eta^{(j)}, && \eta^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Gamma), && j = 1, 2, \ldots \quad (1.6)
\end{aligned}
$$

with $u^{(0)}$ independent of the i.i.d. sequences $\{\xi^{(j)}\}$ and $\{\eta^{(j)}\}$. For a given time index $j \in \mathbb{N}$, the *filtering* goal is to compute the *filtering distribution* $p\big(u^{(j)}|Y^{(j)}\big)$, where $Y^{(j)} := \{y^{(1)}, \ldots, y^{(j)}\}$. The filtering distribution provides a probabilistic summary of the state $u^{(j)}$ conditional on observations up to time j. Given access to the filtering distribution at the preceding time-step $j-1$, $p\big(u^{(j)}|Y^{(j)}\big)$ may be obtained by the following two-step procedure:

$$
\begin{aligned}
&\text{(Forecast)} \qquad p\big(u^{(j)}|Y^{(j-1)}\big) = \int \mathcal{N}(u^{(j)}; Au^{(j-1)}, \Xi) p\big(u^{(j-1)}|Y^{(j-1)}\big)\, du^{(j-1)}, \quad (1.7)\\
&\text{(Analysis)} \qquad p\big(u^{(j)}|Y^{(j)}\big) \propto \mathcal{N}(y^{(j)}; Hu^{(j)}, \Gamma) p\big(u^{(j)}|Y^{(j-1)}\big). \quad (1.8)
\end{aligned}
$$

The *forecast distribution* $p\big(u^{(j)}|Y^{(j-1)}\big)$ represents our knowledge of the state at time j given past observations and is computed using the dynamics model. In the analysis step, the new observation y_j is assimilated via Bayes' formula, with the prior given by the forecast distribution and the likelihood determined by the observation model. When the state dimension d is large or the dynamics are nonlinear, making exact computation infeasible, the Ensemble Kalman Filter (EnKF) is commonly used.

At $j = 0$, an *initial ensemble* of N particles is drawn from $\mathcal{N}(\mu^{(0)}, \Sigma^{(0)})$. The ensemble is then iteratively updated: in the forecast step, it is propagated through the system dynamics, producing the *forecast ensemble*; in the analysis step, each ensemble member is updated using a single-step ensemble Kalman update, yielding the *analysis ensemble*. Thus, the algorithm can be seen as a multi-step ensemble Kalman update, where each step builds on the previous ensemble output.

A key theoretical challenge is the dependence structure of the ensemble. While the initial particles are independent, dependence arises at the first analysis step ($j = 1$) and becomes increasingly complex due to the recursive nature of the algorithm. Each updated particle υ_n depends nonlinearly on $\widehat{C}$, which itself is a function of all prior ensemble members $\{u_n\}_{n=1}^N$. This intricate dependence makes non-asymptotic analysis particularly difficult.

In this thesis, we propose a novel algorithm called the Resampled Ensemble Kalman Filter REnKF, which employs a simple resampling step at each filtering cycle to break the correlations between ensemble members. The algorithm is amenable to theoretical analysis in the linear Gaussian setting, and shows good performance in practice, comparable to the EnKF, even in non-linear settings, making it a promising approach for a wide range of applications.

1.2 Covariance Operator Estimation

The study of covariance operator estimation is motivated by the fact that operational algorithms for numerical weather prediction (for example, the EnKF) rely on an ensemble of forecasts to estimate a background prior covariance. In these applications and many others, the data used to specify the prior covariance represent finely discretized functions. As data resolution continues to improve, we wish to understand the fundamental dimension-free, discretization-independent quantities that determine the difficulty of estimating the prior covariance. Relatedly, operator learning, i.e. the task of recovering an operator from pairs of inputs and outputs or from trajectory data, has also received increased attention motivated by recent machine learning techniques to solve partial differential equations. In this line of work, we have investigated a class of L^q-sparse operators where the kernel need not concentrate around its diagonal, and an even more flexible family of weighted L^q-sparse operators that further allow for extreme heterogeneity of the underlying process across the domain. We further identify a sufficient interpretable condition which we term the *small-lengthscale* setting in which many of the most widely used covariance operators can provably be shown to belong to these structured classes. In this regime, we further show that the sample complexity is determined by the correlation lengthscale of the operator.

1.2.1 Unstructured Case

In the high-dimensional setting where $d \gg N$, estimating the covrariance operator has been thoroughly studied Vershynin [2010], [Wainwright, 2019, Chapter 6]. Given i.i.d. observations $u, u_1, \ldots, u_N$ drawn from a d-dimensional centered Gaussian with covariance operator (matrix) $\Sigma := \mathbb{E}[uu^\top]$, a natural estimator for Σ is the sample average $\hat{\Sigma} := \frac{1}{N}\sum_{n=1}^{N} u_n u_n^\top$. The goal is to control the operator-norm deviation $\|\hat{\Sigma} - \Sigma\|$. For simplicity, we focus on the Gaussian case in this chapter, although many results extend easily to the sub-Gaussian setting. Classical results imply that for a universal constant $c > 0$ and any $t \geq 1$, then with probability at least $1 - e^{-t}$, then

$$\|\widehat{\Sigma} - \Sigma\| \leq c\|\Sigma\| \left(\sqrt{\frac{d}{N}} \vee \frac{d}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

This directly implies that, with probability exceeding $1 - e^{-d}$, it suffices to take $N \geq d$ samples to achieve a small estimation error. However, since the ensemble size N is often much smaller than d, this result fails to explain the empirical success of ensemble Kalman algorithms. While the bound is sharp in certain cases (e.g., when $\Sigma = I_d$), it is overly pessimistic in practical settings where ensemble Kalman updates operate. A key breakthrough in Koltchinskii and Lounici [2017] established a dimension-free alternative: there exists a universal constant $c > 0$ such that for any $t \geq 1$, with probability at least $1 - e^{-t}$,

$$\|\widehat{\Sigma} - \Sigma\| \asymp c\|\Sigma\| \left(\sqrt{\frac{r(\Sigma)}{N}} \vee \frac{r(\Sigma)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right),$$

where $r(\Sigma) := \text{Tr}(\Sigma)/\|\Sigma\|$ is the effective dimension of Σ. The effective dimension remains small whenever the covariance spectrum decays rapidly. Notably, this result extends to infinite-dimensional settings where the random vectors $u, u_1, \ldots, u_N$ reside in a Hilbert space H. In the finite-dimensional case $H = \mathbb{R}^d$, we have $r(\Sigma) \leq d$, implying that accurate estimation is possible even when the ensemble size is significantly smaller than the state dimension, provided the covariance operator has low effective dimension. In this thesis, we demonstrate how such bounds can be used to control the error of ensemble Kalman updates.

1.2.2 Structured Case

In much of the literature on ensemble Kalman updates (e.g. Tong and Morzfeld [2023], Bergemann and Reich [2010a], Petrie [2008]), a modified version of the sample covariance estimator known as the *localized* covariance estimator and denoted $\widehat{\Sigma}_\rho$, where ρ represents the localization radius is utilized. This approach is particularly relevant when the state vector components correspond to spatial locations: the sample covariance is first computed and then adjusted elementwise to down-weight terms associated with distant state coordinates. In much of the literature, this modification is heuristically motivated as a way to "remove spurious correlations", though the optimal choice of its hyper-parameters remains unclear. In this thesis, we identify Localization as a form of regularized covariance estimation, aligning with the foundational works Wu and Pourahmadi [2003], Bickel and Levina [2008a,b] and extensively explored in high-dimensional statistics (see Pourahmadi [2013] for a textbook discussion). More concretely, when the true covariance matrix Σ exhibits an inherent *sparse* structure –meaning many of its elements are exactly or nearly zero– a thresholding-based estimator provides a preferable alternative. These estimators determine whether each entry of the sample covariance falls below a carefully selected threshold ρ, treating such entries as spurious and either setting them to zero or shrinking them accordingly.

For such estimators that take advantage of sparsity in the target, it can be shown that there exists a universal positive constant c such that, with probability at least $1 - e^{-t}$, an appropriately chosen radius ρ ensures:

$$\|\widehat{\Sigma}_\rho - \Sigma\| \leq c\|\Sigma\|_{\max} \left(\sqrt{\frac{\log d}{N}} \vee \frac{\log d}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

This bound demonstrates that accurate covariance estimation necessitates an ensemble size that scales with the logarithm of the dimension, representing a significant improvement over scenarios without structural assumptions. This leads to the following key questions in our study of ensemble Kalman updates:

1. How should sparsity be defined for an infinite dimensional covariance operator?

2. Is there an analogous concept of effective dimension for infinite-dimensional covariance operators with additional sparse structure?

3. Can sparse covariance matrix estimation bounds be extended to an infinite-dimensional setting?

In this thesis, we affirmatively address all three questions. Concretely, we investigate the setting in which u is an infinite-dimensional random field with covariance model that satisfies a novel notion of approximate sparsity. We show that the statistical error of thresholded estimators can be bounded in terms of two dimension-free quantities: the expected supremum of the field and the sparsity level.

1.2.3 Small Lengthscale Analysis

One of the major contributions of this work is to showcase the benefit of thresholding estimators in the challenging regime where the correlation lengthscale of the field is small relative to the size of the physical domain. Mathematically, given a process with covariance function $k = k_\lambda$ where $\lambda > 0$ is the correlation lengthscale, we study the regime in which $\lambda \to 0$. In this setting, our theory characterizes the aforementioned dimension free quantities in terms of λ. While a vast literature in nonparametric statistics and approximation theory highlights the key role of smoothness in determining optimal convergence rates for many nonparametric estimation tasks, our non-asymptotic theory emphasizes that the lengthscale rather than the smoothness of the covariance function drives the difficulty of the estimation problem and the advantage of thresholded estimators. Fields with small correlation lengthscale are ubiquitous in applications. For instance, they arise naturally in climate science and numerical weather forecasting, where global forecasts need to account for the effect of local processes with a small correlation lengthscale, such as cloud formation or propagation of gravitational waves.

1.2.4 Empirical Process Theory

The dimension-free bounds in this work are primarily enabled by techniques from empirical process theory applied to the covariance estimation problem. Specifically, let $u, u_1, \ldots, u_N \overset{\text{i.i.d.}}{\sim} \mathbb{P}$ be a sequence of random variables on a probability space $(\Omega, \mathbb{P})$. The quadratic empirical process associated with a function class $\mathcal{F}$ on $(\Omega, \mathbb{P})$ is given by

$$f \mapsto \frac{1}{N} \sum_{n=1}^{N} f^2(u_n) - \mathbb{E} f^2(u), \quad f \in \mathcal{F}.$$

It can be shown that the variational form of $\|\hat{\Sigma} - \Sigma\|$ corresponds to the supremum of a quadratic empirical process over a suitably chosen function class $\mathcal{F}$. This observation underpins the proof of the fundamental bound in Koltchinskii and Lounici [2017], which relies on the following inequality from Klartag and Mendelson [2005]:

$$\mathbb{E} \sup_{f \in \mathcal{F}} \left| \frac{1}{N} \sum_{n=1}^{N} f^2(u_n) - \mathbb{E} f^2(u) \right| \leq c \left(\sup_{f \in \mathcal{F}} \|f\|_{\psi_1} \frac{\gamma_2(\mathcal{F}; \psi_2)}{\sqrt{N}} \vee \frac{\gamma_2^2(\mathcal{F}; \psi_2)}{N} \right),$$

which holds whenever $\mathcal{F}$ is a symmetric function class satisfying $\mathbb{E}f(u) = 0$. The bound is expressed in terms of the ψ_1-Orlicz norm of the function class and Talagrand's generic chaining functional. Precise control over these quantities gives rise to the concept of an effective dimension in the deviation bound for the sample covariance operator. In this work, we extend these techniques to develop a new notion of effective dimension specifically suited for a newly introduced class of structured covariance operators.

1.2.5 Outline and Main Contributions

We now provide an outline of the upcoming chapters and summarize their key contributions.

1.2.6 Chapter 2 - Non-Asymptotic Analysis of Ensemble Kalman Updates: Effective Dimension and Localization

In Chapter 2, we establish non-asymptotic error bounds in terms of suitable notions of effective dimension of the prior covariance model that account for spectrum decay (which may represent smoothness of a prior random field) and approximate sparsity (which may represent spatial decay of correlations). Our work complements mean-field analyses of ensemble Kalman updates and identifies scenarios where mean-field behavior holds with moderate N. In addition to demystifying the practical success of ensemble Kalman methods with a small ensemble size, our non-asymptotic perspective allows us to tell apart, on accuracy grounds, implementations of ensemble Kalman updates that use perturbed observations and square root filtering. These implementations become equivalent in the large N limit, and therefore their differences in accuracy cannot be captured by asymptotic results. Furthermore, our non-asymptotic perspective provides new understanding on the importance of localization, a procedure widely used by practitioners that involves tapering or "localizing" empirical covariance estimates to avoid spurious correlations. A key contribution of our framework is to obtain *dimension-free* bounds. Removing the dependence on the state dimension d is particularly important since in many applications d represents the discretization of some infinite dimensional field. Our bounds therefore capture the intrinsic geometric complexity of the problem as opposed to being a function of the discretization level. This chapter is adapted from the following publication:

O. Al-Ghattas and D. Sanz-Alonso, *Non-asymptotic analysis of ensemble Kalman updates: effective dimension and localization*, *Information and Inference: A Journal of the IMA*, vol. 13, no. 1, pp. iaad043, 2024.

1.2.7 Chapter 3 - Ensemble Kalman Filters with Resampling

In Chapter 3 we study multi-step settings that are complicated by perturbed observations and stochastic dynamics, which are commonly used in the Ensemble Kalman Filtering (EnKF) literature. The EnKF is particularly difficult to analyse theoretically due to the presence of correlations between ensemble members, since the Kalman gain used to update each particle depends on the entire ensemble. In this work, we investigate a simple modification of EnKF that incorporates a resampling step to break these correlations. The new algorithm is amenable to a theoretical analysis that extends and improves upon those available for filters without resampling, while also maintaining a similar empirical performance. We consider a simple parametric resampling scheme: at the beginning of each filtering step, members of the ensemble are independently sampled from a Gaussian distribution whose mean and covariance match those of the ensemble at the previous time-step. Thereafter, the filtering step can be carried out using any of the numerous existing EnKF variants. For the resulting algorithm, which we term REnKF, we establish theoretical guarantees that extend and improve upon those available for filters without resampling. Our theoretical guarantees hold in the linear-Gaussian setting in which we provide a detailed error analysis of the ensemble mean and covariance as estimators of the mean and covariance of the filtering distributions, given by the Kalman filter. Our theory covers both stochastic and deterministic dynamical systems; in addition, it covers both stochastic implementations based on *perturbed observations* and deterministic implementations based on *square-root* filters. Importantly, our error-bounds are non-asymptotic and dimension-free: they hold for any given ensemble size and are written in terms of the effective-dimension of the covariance of the initial distribution, and of the dynamics and observation models. This chapter is adapted from the following publication:

O. Al-Ghattas, J. Bao, and D. Sanz-Alonso, *Ensemble Kalman filters with resampling*, *SIAM/ASA Journal on Uncertainty Quantification*, vol. 12, no. 2, pp. 411–441, 2024.

1.2.8 Chapter 4 - Covariance Operator Estimation: Sparsity, Lengthscale, and Ensemble Kalman Filters

In Chapter 4, we first lift the theory of covariance estimation from finite to infinite dimension. In the finite-dimensional setting, a rich body of work shows that, exploiting various forms of sparsity, it is possible to consistently estimate the covariance matrix of a vector $u \in \mathbb{R}^{d_u}$ with $N \sim \log(d_u)$ samples as opposed to the non-sparse setting in which $N \sim d_u$ samples are needed. In this work we investigate the setting in which u is an infinite-dimensional random field with an approximately sparse covariance model. Specifically, we generalize notions of approximate sparsity often employed in the finite-dimensional covariance estimation literature. We show that the statistical error of thresholded estimators can be bounded in terms of two dimension-free quantities: the expected supremum of the field and the sparsity level. The second contribution is to showcase the benefit of thresholding in the challenging regime where the correlation lengthscale of the field is small relative to the size of the physical domain. While a vast literature in nonparametric statistics and approximation theory highlights the key role of smoothness in determining optimal convergence rates for many nonparametric estimation tasks, our non-asymptotic theory emphasizes that the lengthscale rather than the smoothness of the covariance function drives the difficulty of the estimation problem and the advantage of thresholded estimators. Fields with small correlation lengthscale are ubiquitous in applications. For instance, they arise naturally in climate science and numerical weather forecasting, where global forecasts need to account for the effect of local processes with a small correlation lengthscale, such as cloud formation or propagation of gravitational waves. The third contribution of this paper is to demonstrate the advantage of using thresholded covariance estimators within ensemble Kalman filters, generalizing the theory in Al-Ghattas and Sanz-Alonso [2024b] to the infinite dimensional setting. Our theory explains when and why localized EnKFs are expected to out-perform non-localized filters. This chapter is adapted from the following publication:

O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek, *Covariance operator estimation: sparsity, lengthscale, and ensemble Kalman filters*, *Bernoulli, 31(3), 2377-2402*, 2025

1.2.9 Chapter 5 - Covariance Operator Estimation via Adaptive Thresholding

In Chapter 5, consider estimating the covariance operator of a highly nonstationary process with marginal variance that is permitted to vary widely in the domain. These operators satisfy a *weighted* L_q-sparsity condition. For covariance operators in this class, we establish a bound on the operator norm error of the *adaptive threshold* estimator in terms of two dimension-free quantities: the sparsity level and the expected supremum of the normalized field. In contrast to Al-Ghattas et al. [2023], our theory allows for covariance models with unbounded marginal variance functions. We then compare our adaptive threshold estimator with other estimators of interest, namely the *universal threshold* estimator in Al-Ghattas et al. [2023] and the *sample covariance* estimator. For universal thresholding, we prove a lower bound that is larger than our upper bound for adaptive thresholding. We generalize the small lengthscale setting of Al-Ghattas et al. [2023] to one in which both the lengthscale and a parameter controlling the range of the marginal variance function is allowed to be arbitrarily small, and further prove an exponential improvement in sample complexity of the adaptive threshold estimator compared to the sample covariance. A key technical contribution of this work is the analysis of the product sub-exponential classes that arise in estimating the variance components of the process, which in turn are needed to adaptively set the threshold radius. The analysis demonstrates how to derive non-asymptotic and dimension free bounds for a large family of product empirical processes and are potentially of independent interest. The results are adapted from the following paper which received a minor revision at *Stochastic Processes and their Applications:*

O. Al-Ghattas and D. Sanz-Alonso, *Covariance Operator Estimation via Adaptive Thresholding*, *arXiv preprint arXiv:2405.18562*, 2024.

1.2.10 Additional Work

The following recent pre-prints carried out during my PhD are not included in this thesis:

- In Al-Ghattas et al. [2024b], we derive the information-theoretic limits of covariance operator estimation in the structured setting through the use of a minimax framework. In addition to the L_q-sparsity first considered in Al-Ghattas et al. [2023], we also consider banded integral operators with kernels that decay rapidly off-the-diagonal. For both classes, we establish minimax optimal lower bounds using a novel and general framework that lifts the theory from high-dimensional matrix estimation to the operator setting. In so doing, we identify the dimension-free quantities that determine the sample complexity. Additionally, we show that tapering and thresholding estimators achieve the minimax optimal rate in the two respective classes.

  O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek, *Optimal estimation of structured covariance operators*, *arXiv preprint arXiv:2408.02109*, 2024.

- In Al-Ghattas et al. [2025], we establish sharp dimension-free concentration inequalities and expectation bounds for the deviation of the sum of simple random tensors from its expectation. As part of our analysis, we use generic chaining techniques to obtain a sharp high-probability upper bound on the suprema of multiproduct empirical processes. In so doing, we generalize classical results for quadratic and product empirical processes to higher-order settings.

  O. Al-Ghattas, J. Chen, and D. Sanz-Alonso, *Sharp Concentration of Simple Random Tensors*, *arXiv preprint arXiv:2502.16916*, 2025.

CHAPTER 2

NON-ASYMPTOTIC ANALYSIS OF ENSEMBLE KALMAN UPDATES: EFFECTIVE DIMENSION AND LOCALIZATION

This chapter is adapted from the publication listed below and is used with permission of the publisher.

O. Al-Ghattas and D. Sanz-Alonso, *Non-asymptotic analysis of ensemble Kalman updates: effective dimension and localization*, *Information and Inference: A Journal of the IMA*, vol. 13, no. 1, pp. iaad043, 2024.

2.1 Introduction

The aim of this chapter is to develop a *non-asymptotic* analysis of ensemble Kalman updates that rigorously explains why, and under what circumstances, a small ensemble size may suffice. To that end, we establish non-asymptotic error bounds in terms of suitable notions of effective dimension of the prior covariance model that account for spectrum decay (which may represent smoothness of a prior random field) and approximate sparsity (which may represent spatial decay of correlations). Our work complements mean-field analyses of ensemble Kalman updates and identifies scenarios where mean-field behavior holds with moderate N.

In addition to demystifying the practical success of ensemble Kalman methods with a small ensemble size, our non-asymptotic perspective allows us to tell apart, on accuracy grounds, implementations of ensemble Kalman updates that use perturbed observations and square root filtering. These implementations become equivalent in the large N limit, and therefore their differences in accuracy cannot be captured by asymptotic results. Furthermore, our non-asymptotic perspective provides new understanding on the importance of localization, a procedure widely used by practitioners that involves tapering or "localizing" empirical covariance estimates to avoid spurious correlations.

Rather than providing a complete, definite analysis of any particular ensemble Kalman method, our goal is to bring to bear a new set of tools from high-dimensional probability and statistics to the study of these algorithms. In particular, our work builds on and contributes to the theory of high-dimensional covariance estimation, which we believe is fundamental to the understanding of ensemble Kalman methods. To make the presentation accessible to a wide audience, we assume no background knowledge on covariance estimation or on ensemble Kalman methods.

2.1.1 Problem Description

Consider the inverse problem of recovering $u \in \mathbb{R}^d$ from data $y \in \mathbb{R}^k$, corrupted by noise η, where

$$y = \mathcal{G}(u) + \eta, \tag{2.1}$$

$\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^k$ is the forward model, and $\eta \sim \mathbb{P}_\eta = \mathcal{N}(0, \Gamma)$ is the observation error with positive-definite covariance matrix Γ. An ensemble Kalman update takes as input a *prior ensemble* $\{u_n\}_{n=1}^N$ and observed data y, and returns as output an *updated ensemble* $\{\upsilon_n\}_{n=1}^N$ that blends together the information in the prior ensemble and in the data. Two main types of problems will be investigated: posterior approximation and sequential optimization. In the former, ensemble Kalman updates are used to approximate a posterior distribution in a Bayesian linear setting; in the latter, they are used within optimization algorithms for nonlinear inverse problems.

Posterior Approximation

If the forward model is linear, i.e. $\mathcal{G}(u) = Au$ for some matrix $A \in \mathbb{R}^{k \times d}$, and A is ill-conditioned or $d \gg k$, naive inversion of the data by means of the (generalized) inverse of A

results in an amplification of small observation error η into large error in the reconstruction of u. In such situations, regularization is needed to stabilize the solution. To this end, one may adopt a Bayesian approach and place a Gaussian prior on the unknown $u \sim \mathbb{P}_u = \mathcal{N}(m, C)$ with positive-definite C; the prior distribution then acts as a probabilistic regularizer. The Bayesian solution to the inverse problem (2.1) is a full characterization of the posterior distribution $\mathbb{P}_{u|y}$, that is, the distribution of u given y. A standard calculation shows that $\mathbb{P}_{u|y} = \mathcal{N}(\mu, \Sigma)$, with

$$
\begin{aligned}
\mu &= m + CA^\top (ACA^\top + \Gamma)^{-1}(y - Am), \\
\Sigma &= C - CA^\top (ACA^\top + \Gamma)^{-1} AC,
\end{aligned}
\tag{2.2}
$$

which require storage of $d \times d$ matrices and consequently are difficult to compute explicitly when the state dimension d is large. A posterior-approximation ensemble Kalman update transforms a prior ensemble $\{u_n\}_{n=1}^N$ drawn from $\mathbb{P}_u$ into an updated ensemble $\{\upsilon_n\}_{n=1}^N$ whose sample mean and sample covariance approximate the mean and covariance of $\mathbb{P}_{u|y}$. Ensemble Kalman updates enjoy a low computational and memory cost when the ensemble size N is smaller than the state dimension d. In Section 2.2 we establish non-asymptotic error bounds that ensure that if N is larger than a suitably defined *effective dimension*, then the sample mean and sample covariance of the updated ensemble approximate well the true posterior mean and covariance in (2.2). We refer to methods that are capable of approximating well the posterior $\mathbb{P}_{u|y}$ in a linear-Gaussian setting as *posterior-approximation* algorithms.

Sequential Optimization

When faced with a general nonlinear model $\mathcal{G}$, exact characterization of the posterior can be challenging. One may then opt for an optimization framework and solve the inverse problem (2.1) by minimizing a user-chosen objective function. Starting from a prior ensemble

$\{u_n\}_{n=1}^N$ drawn from a measure $\mathbb{P}_u$ that encodes prior beliefs about u, an ensemble Kalman update returns an updated ensemble $\{\upsilon_n\}_{n=1}^N$ whose sample mean approximates the desired minimizer. The process can be iterated by taking the updated ensemble to be the prior ensemble of a new ensemble Kalman update. Under suitable conditions on $\mathcal{G}$, and after a sufficient number of such updates, all particles in the ensemble collapse into the minimizer of the objective. Ensemble Kalman optimization algorithms are derivative-free methods, and are therefore particularly useful when derivatives of the model $\mathcal{G}$ are unavailable or expensive to compute. As for posterior-approximation algorithms, implementing each update has low computational and memory cost when the ensemble size N is small. In Section 2.3 we will establish non-asymptotic error bounds that ensure that if N is larger than a suitably defined *effective dimension*, then each particle update $u_n \mapsto \upsilon_n$, $1 \leq n \leq N$, approximates well an idealized mean-field update computed with an infinite number of particles; this suggests that the evolution of particles along an ensemble-based sequential optimizer is close to an idealized mean-field evolution. We refer to methods that solve the inverse problem (2.1) by minimization of an objective function as *sequential-optimization* algorithms.

2.1.2 Summary of Contributions and Outline

- Section 2.2 is concerned with posterior-approximation algorithms. The main results, Theorems 2.2.3 and 2.2.5, give non-asymptotic bounds on the estimation of the posterior mean and covariance in terms of a standard notion of effective dimension that accounts for spectrum decay in the prior covariance model. Our analysis explains the statistical advantage of square root updates over perturbed observation ones. We also discuss the deterioration of our bounds in small noise limits where the prior and the posterior become mutually singular.

- Section 2.3 is concerned with sequential-optimization algorithms. The main results, Theorems 2.3.5 and 2.3.7, give non-asymptotic bounds on the approximation of mean-field particle updates using ensemble Kalman updates with and without localization. Our analysis explains the advantage of localized updates if the prior covariance satisfies a soft-sparsity condition. For the study of localized updates, we show in Theorems 2.3.1 and 2.3.3 new dimension-free covariance estimation bounds in terms of a new notion of effective dimension that simultaneously accounts for spectrum decay and approximate sparsity in the prior covariance model.

- Section 2.4 concludes with a summary of our work and several research directions that stem from our non-asymptotic analysis of ensemble Kalman updates. We also discuss the potential and limitations of localization in posterior-approximation algorithms.

- The proofs of all our results are deferred to three appendices.

2.1.3 Related Work

Ensemble Kalman methods —overviewed in Evensen [2009], Katzfuss et al. [2016], Houtekamer and Zhang [2016], Roth et al. [2017], Chada et al. [2021], Sanz-Alonso et al. [2023a]— first appeared as filtering algorithms in the data assimilation literature Evensen [1995], Evensen and Leeuwen [1996], Burgers et al. [1998], Houtekamer and Derome [1995], Houtekamer and Mitchell [1998]. The goal of data assimilation is to estimate a time-evolving state as new observations become available Reich and Cotter [2015], Asch et al. [2016], Law et al. [2015], Majda and Harlim [2012], Leeuwen et al. [2015], Särkkä [2013], Sanz-Alonso et al. [2023a]. Ensemble Kalman filters (EnKFs) solve an inverse problem of the form (2.1) every time a new observation is acquired. In that filtering context, (2.1) encodes the relationship between the state u and observation y at a given time t, and the prior on u is specified by propagating a probabilistic estimate of the state at time $t-1$ through the dynamical system that governs the state evolution. To approximate this prior, EnKFs propagate an ensemble of N particles through the dynamics, and subsequently update this prior *forecast* ensemble into an updated *analysis* ensemble that assimilates the new observation. Thus, an ensemble Kalman update is performed every time a new observation is acquired. The goal is that the sample mean and sample covariance of the updated ensemble approximate well the mean and covariance of the filtering distribution, that is, the conditional distribution of the state at time t given all observations up to time t. While only giving provably accurate posterior approximation in linear settings Ernst et al. [2015], EnKFs are among the most popular methods for high-dimensional nonlinear filtering, in particular in numerical weather forecasting. In such applications the state dimension can be very large, but the effective dimension of the filter update is often much lower due to smoothness of the state and decay of correlations in space. Moreover, in practice the analysis step can be constrained to the subspace determined by the expanding directions of the dynamics Trevisan and Uboldi [2004].

The papers Gu and Oliver [2007], Li and Reynolds [2007], Reynolds et al. [2006] introduced ensemble Kalman methods for inverse problems in petroleum engineering and the geophysical sciences. Application-agnostic ensemble Kalman methods for inverse problems were developed in Iglesias et al. [2013], Iglesias [2016], inspired by classical regularization schemes Hanke [1997]. Since then, a wide range of sequential-optimization algorithms for inverse problems have been proposed that differ in the objective function they seek to minimize and in how ensemble Kalman updates are implemented. We refer to Subsection 2.2.1 for further background and to Chada et al. [2021] for a review.

Ensemble Kalman methods for inverse problems and data assimilation have been studied extensively from a large N asymptotic point of view, see e.g. Li and Xiu [2008], Le Gland et al. [2009], Mandel et al. [2011], Kwiatkowski and Mandel [2015], Ernst et al. [2015], Del Moral and Tugaut [2018], Herty and Visconti [2019], Law et al. [2016b], Garbuno-Inigo et al. [2020], Bishop and Del Moral [2023], Chen et al. [2022], Ding and Li [2021]. A complementary line of work Harlim and Majda [2010], Gottwald and Majda [2013], Kelly et al. [2015], Tong et al. [2015, 2016] has focused on challenges faced by ensemble Kalman methods, including loss of stability and catastrophic filter divergence. Two overarching themes that underlie large N asymptotic analyses are to ensure consistency and to derive equations for the mean-field evolution of the ensemble. Related to this second theme, several works (e.g. Schillings and Stuart [2017], Blömker et al. [2018], Blömker et al. [2019], Guth et al. [2020], Chada et al. [2021], Tong and Morzfeld [2023]) set the analysis in a *continuous time limit*; the idea is to view Kalman updates as occurring over an artificial discrete-time variable, and then take the time between updates to be infinitesimally small to formally derive differential equations for the evolution of the ensemble or its density. Large N asymptotics and continuous time limits have resulted in new theoretical insights and practical advancements. However, an important caveat of these results is that they cannot tell apart implementations of ensemble Kalman methods that become equivalent in large N or continuous time asymptotic regimes. Moreover, several papers (e.g. Bergemann and Reich [2010a,b], Kelly and Stuart [2014], Schillings and Stuart [2017], Majda and Tong [2018]) have noted that large N asymptotic analyses fail to explain empirical results that report good performance with a moderately sized ensemble in problems with high state dimension; for instance, $d \sim 10^9$ and $N \sim 10^2$ in operational numerical weather prediction. Finally, the note Nüsken and Reich [2019] shows subtle but important differences in the evolution of interacting particle systems with finite ensemble size when compared to their mean-field counterparts Garbuno-Inigo et al. [2020].

In this chapter we adopt a non-asymptotic viewpoint to establish sufficient conditions on the ensemble size for posterior-approximation and sequential-optimization algorithms. Empirical evidence in Ott et al. [2004] suggests that there is a sample size N^* above which ensemble Kalman methods are effective. The seminal work Furrer and Bengtsson [2007] conducts insightful explicit calculations that motivate our more general theory. Following the analysis of ensemble Kalman methods in Furrer and Bengtsson [2007] and the study of importance sampling and particle filters in Agapiou et al. [2017], Sanz-Alonso [2018], Sanz-

Alonso and Wang [2021], Bickel et al. [1994], Snyder et al. [2016], Bengtsson et al. [2008], Snyder [2011], Chorin and Morzfeld [2013], Snyder et al. [2015], we focus on analyzing a single ensemble Kalman update rather than on investigating the propagation of error across multiple updates. While in practice ensemble Kalman methods for posterior approximation in data assimilation and for sequential optimization in inverse problems often perform many updates, focusing on a single update enables us to clearly demonstrate the tight connection between the sample complexity of ensemble updates and the effective dimension of the prior; additionally, for some posterior-approximation algorithms our theory generalizes in a straightforward way to multi-step implementations, as we shall demonstrate in Section 2.2. More importantly, the focus on a single update allows us to tell apart, on accuracy grounds, perturbed observations and square root implementations of ensemble Kalman updates, as well as implementations with and without localization. Similar considerations motivate the study of sufficient sample size for importance sampling in Morzfeld et al. [2017], Snyder et al. [2016], Agapiou et al. [2017], Chatterjee and Diaconis [2018], Sanz-Alonso [2018], Sanz-Alonso and Wang [2021], where the focus on a single update facilitates establishing clear comparisons between standard and optimal proposals, and identifying meaningful notions of dimension to characterize necessary and sufficient conditions on the required sample size. Our work builds on and develops tools from high-dimensional probability and statistics Wainwright [2019], Vershynin [2018], Bickel and Levina [2008a], Levina and Vershynin [2012], Chen et al. [2012], Cai and Yuan [2012], Cai and Zhou [2012a]. In particular, we bring to bear thresholded Bickel and Levina [2008a], Cai and Yuan [2012] and masked covariance estimators Levina and Vershynin [2012], Chen et al. [2012] to the understanding of localization in ensemble Kalman methods. In so doing, we establish new dimension-free covariance and cross-covariance estimation bounds under approximate sparsity —see Theorems 2.3.1 and 2.3.3.

2.1.4 Notation

Given two positive sequences $\{a_n\}$ and $\{b_n\}$, the relation $a_n \lesssim b_n$ denotes that $a_n \leq cb_n$ for some constant $c > 0$. If the constant c depends on some quantity τ, then we write $a \lesssim_\tau b$. If both $a_n \lesssim b_n$ and $b_n \lesssim a_n$ hold simultaneously, then we write $a_n \asymp b_n$. Throughout, we denote positive universal constants by c, c_1, c_2, c_3, c_4, and the value of a universal constant may differ from line to line. For a vector $v \in \mathbb{R}^N$, $\|v\|_p^p = \sum_{n=1}^N |v_n|^p$. For a matrix $A \in \mathbb{R}^{n\times m}$, the operator norm is given by $\|A\| = \sup_{\|v\|_2=1} \|Av\|_2$. $\mathcal{S}_+^d$ denotes the set of $d \times d$ symmetric positive-semidefinite matrices, and $\mathcal{S}_{++}^d$ denotes the set of $d \times d$ symmetric positive-definite matrices. $A^\dagger$ denotes the pseudo-inverse of A. 1_N denotes the N-dimensional vector vector of ones, 0_d denotes the d-dimensional vector of zeroes, and $O_{d\times k}$ is the $d \times k$ matrix of zeroes. $\mathbf{1}_B$ denotes the indicator of the set B. $\equiv$ denotes a definition. $\circ$ denotes the matrix Hadamard or Schur (elementwise) product. Given a non-decreasing, non-zero convex function $\psi : [0,\infty] \to [0,\infty]$ with $\psi(0) = 0$, the Orlicz norm of a real random variable X is $\|X\|_\psi = \inf\{t > 0 : \mathbb{E}[\psi(t^{-1}|X|)] \leq 1\}$. In particular, for the choice $\psi_p(x) \equiv e^{x^p} - 1$ for $p \geq 1$, real random variables that satisfy $\|X\|_{\psi_2} < \infty$ are referred to as sub-Gaussian. A random vector X is sub-Gaussian if $\|v^\top X\|_{\psi_2} < \infty$ for any v such that $\|v\|_2 = 1$. For a differentiable function $g : \mathbb{R}^d \to \mathbb{R}^k$, $Dg \in \mathbb{R}^{d\times k}$ denotes the Jacobian of g.

All the methods we study have the same starting point of a prior ensemble

$$u_1, \ldots, u_N \overset{\text{i.i.d.}}{\sim} \mathcal{N}(m, C),$$

and observed data y generated according to (2.1), which are to be used in generating an updated ensemble $\{\upsilon_n\}_{n=1}^N$. We denote the prior sample means by

$$\widehat{m} \equiv \frac{1}{N}\sum_{n=1}^N u_n, \qquad \bar{\mathcal{G}} \equiv \frac{1}{N}\sum_{n=1}^N \mathcal{G}(u_n),$$

and the prior sample covariances by

$$\widehat{C} \equiv \frac{1}{N-1}\sum_{n=1}^{N}(u_n - \widehat{m})(u_n - \widehat{m})^\top, \qquad \widehat{C}^{\,pp} \equiv \frac{1}{N-1}\sum_{n=1}^{N}(\mathcal{G}(u_n) - \overline{\mathcal{G}})(\mathcal{G}(u_n) - \overline{\mathcal{G}})^\top,$$
$$\widehat{C}^{\,up} \equiv \frac{1}{N-1}\sum_{n=1}^{N}(u_n - \widehat{m})(\mathcal{G}(u_n) - \overline{\mathcal{G}})^\top. \tag{2.3}$$

The population versions will be denoted by

$$C^{\,pp} \equiv \mathbb{E}\left[\big(\mathcal{G}(u_n) - \mathbb{E}[\mathcal{G}(u_n)]\big)\big(\mathcal{G}(u_n) - \mathbb{E}[\mathcal{G}(u_n)]\big)^\top\right],$$
$$C^{\,up} \equiv \mathbb{E}\left[\big(u_n - m\big)\big(\mathcal{G}(u_n) - \mathbb{E}[\mathcal{G}(u_n)]\big)^\top\right].$$

2.2 Ensemble Kalman Updates: Posterior Approximation Algorithms

In posterior-approximation algorithms we consider the inverse problem (2.1) with a linear forward model, i.e.

$$y = Au + \eta, \qquad \eta \sim \mathcal{N}(0, \Gamma). \tag{2.4}$$

In order to establish comparisons between different posterior-approximation algorithms, as well as to streamline our analysis, we follow the exposition in Kwiatkowski and Mandel [2015] and introduce three operators that are central to the theory: the *Kalman gain* operator $\mathscr{K}$, the *mean-update* operator $\mathscr{M}$, and the *covariance-update* operator $\mathscr{C}$, defined respectively by

$$\mathscr{K} : \mathcal{S}_+^d \to \mathbb{R}^{d\times k}, \qquad \mathscr{K}(C; A, \Gamma) = \mathscr{K}(C) = CA^\top (ACA^\top + \Gamma)^{-1}, \tag{2.5}$$

$$\mathscr{M} : \mathbb{R}^d \times \mathcal{S}_+^d \to \mathbb{R}^d, \qquad \mathscr{M}(m, C; A, y, \Gamma) = \mathscr{M}(m, C) = m + \mathscr{K}(C; A, \Gamma)(y - Am), \tag{2.6}$$

$$\mathscr{C} : \mathcal{S}_+^d \to \mathcal{S}_+^d, \qquad \mathscr{C}(C; A, \Gamma) = \mathscr{C}(C) = \big(I - \mathscr{K}(C; A, \Gamma)A\big)C. \tag{2.7}$$

The pointwise continuity and boundedness of all three operators was established in Kwiatkowski and Mandel [2015], and we summarize these results in Lemmas 2.5.4, 2.5.5, and 2.5.6. We note that the Kalman update (2.2) can be rewritten succinctly as

$$\begin{aligned} \mu &= \mathscr{M}(m, C), \\ \Sigma &= \mathscr{C}(C). \end{aligned} \tag{2.8}$$

2.2.1 Ensemble Algorithms for Posterior Approximation

We study two main classes of posterior-approximation algorithms based on Perturbed Observation (PO) and Square Root (SR) ensemble Kalman updates. In both implementations, the updated ensemble has sample mean $\widehat{\mu}$ and sample covariance $\widehat{\Sigma}$ that are, by design, consistent estimators of the posterior mean μ and covariance Σ in (2.8). Although PO and SR updates are asymptotically equivalent, differences between the two algorithms do exist in finite ensembles, and this difference is captured in our non-asymptotic analysis in Subsection 2.2.3.

Perturbed Observation Update

The PO update, introduced in Evensen [1995], transforms each particle of the prior ensemble according to

$$\begin{aligned} \upsilon_n &= u_n + \mathscr{K}(\widehat{C})\big(y - Au_n - \eta_n\big) \\ &= \mathscr{M}(u_n, \widehat{C}) - \mathscr{K}(\widehat{C})\eta_n, \qquad \eta_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Gamma), \quad 1 \le n \le N. \end{aligned}$$

The form of the update is similar to the Kalman mean update (2.8) albeit with the n-th ensemble member being assigned a perturbed observation $y - \eta_n$. Consequently, denoting the sample mean of the perturbations by $\bar{\eta} \equiv N^{-1}\sum_{n=1}^{N} \eta_n$, the updated ensemble has sample mean

$$\widehat{\mu} \equiv \frac{1}{N}\sum_{n=1}^{N} \upsilon_n = \mathscr{M}(\widehat{m}, \widehat{C}) - \mathscr{K}(\widehat{C})\bar{\eta},$$

and sample covariance

$$\begin{aligned} \widehat{\Sigma} &\equiv \frac{1}{N-1}\sum_{n=1}^{N} \big(\upsilon_n - \widehat{\mu}\big)\big(\upsilon_n - \widehat{\mu}\big)^\top \\ &= \big(I - \mathscr{K}(\widehat{C})A\big)\widehat{C}\big(I - \mathscr{K}(\widehat{C})A\big)^\top + \mathscr{K}(\widehat{C})\widehat{\Gamma}\mathscr{K}^\top(\widehat{C}) \\ &\qquad - \big(I - \mathscr{K}(\widehat{C})A\big)\widehat{C}^{u\eta}\mathscr{K}^\top(\widehat{C}) - \mathscr{K}(\widehat{C})(\widehat{C}^{u\eta})^\top\big(I - A^\top\mathscr{K}^\top(\widehat{C})\big), \end{aligned} \tag{2.9}$$

where

$$\widehat{\Gamma} \equiv \frac{1}{N-1}\sum_{n=1}^{N}(\eta_n - \bar{\eta}_N)(\eta_n - \bar{\eta})^\top, \qquad \text{and} \qquad \widehat{C}^{u\eta} \equiv \frac{1}{N-1}\sum_{n=1}^{N}(u_n - \widehat{m})(\eta_n - \bar{\eta})^\top.$$

To facilitate comparison with the Kalman update in (2.8), we rewrite the PO update as follows:

$$\begin{aligned}\widehat{\mu} &= \mathscr{M}(\widehat{m}, \widehat{C}) - \mathscr{K}(\widehat{C})\bar{\eta}, \\ \widehat{\Sigma} &= \mathscr{C}(\widehat{C}) + \widehat{O},\end{aligned} \tag{2.10}$$

where the *offset* term $\widehat{O}$, obtained as the difference between (2.9) and $\mathscr{C}(\widehat{C})$, is given by

$$\widehat{O} = \mathscr{K}(\widehat{C})(\widehat{\Gamma} - \Gamma)\mathscr{K}^\top(\widehat{C}) - \big(I - \mathscr{K}(\widehat{C})A\big)\widehat{C}^{u\eta}\mathscr{K}^\top(\widehat{C}) - \mathscr{K}(\widehat{C})(\widehat{C}^{u\eta})^\top\big(I - A^\top\mathscr{K}^\top(\widehat{C})\big). \tag{2.11}$$

The offset term $\widehat{O}$ was introduced in [Furrer and Bengtsson, 2007, Proposition 4]. The addition of perturbations serves the purpose of correcting the sample covariance, in the sense that without perturbations the sample covariance is an inconsistent estimator of Σ. To see the consistency of the PO covariance estimator $\widehat{\Sigma}$ in (2.10), note that by Lemma 2.5.6 the map $\mathscr{C}$ is continuous, and so the continuous mapping theorem together with the fact that $\widehat{C}$ is consistent for C imply that $\mathscr{C}(\widehat{C}) \xrightarrow{p} \mathscr{C}(C) = \Sigma$. Further, the offset $\widehat{O}$ converges in probability to zero, which can be shown using that $\widehat{\Gamma} \xrightarrow{p} \Gamma$, $\widehat{C}^{u\eta} \xrightarrow{p} O_{d\times k}$, and the continuity of $\mathscr{K}$ established in Lemma 2.5.4.

Square Root Update

The PO update relies crucially on the added perturbations to maintain consistency and, as noted for example in Evensen [2004], Tippett et al. [2003], Bishop et al. [2001], is asymptotically equivalent to the exact posterior update (2.2). However, for a finite ensemble of size N, the addition of random perturbations introduces an extra source of error into the ensemble Kalman update. The SR update, introduced in Evensen [2004] and surveyed in Tippett et al. [2003], Lange and Stannat [2021], is a deterministic alternative to the PO update. It updates the prior ensemble in a manner that ensures that $\widehat{\Sigma} \equiv \mathscr{C}(\widehat{C})$. This is achieved by first identifying a map $g : \mathbb{R}^{d\times N} \to \mathbb{R}^{d\times N}$ such that $\widehat{\Pi} = g(\widehat{P})$, where

$$\widehat{C} = \widehat{P}\widehat{P}^\top, \qquad \text{and} \qquad \mathscr{C}(\widehat{C}) = \widehat{\Pi}\widehat{\Pi}^\top,$$

with both factorizations guaranteed to exist since $\widehat{C}, \mathscr{C}(\widehat{C}) \in \mathcal{S}^d_+$. Consistency of $\widehat{\Sigma}$ can then be ensured by choosing g to satisfy $g(\widehat{P})g(\widehat{P})^\top \equiv \mathscr{C}(\widehat{C})$, with this being referred to as the *consistency condition* in Lange and Stannat [2021]. There are infinitely many such g, each of which lead to a variant of the SR update. Here we describe two of the most popular variants in the literature as outlined in Tippett et al. [2003]: the Ensemble Transform Kalman update Bishop et al. [2001] and the Ensemble Adjustment Kalman update Anderson [2001] with respective transformations $g_T(\widehat{P}) = \widehat{P}T$ and $g_A(\widehat{P}) = B\widehat{P}$, for matrices T and B. Both g_T and g_A are therefore linear maps, with g_T post-multiplying $\widehat{P}$, which implies a transformation on the N-dimensional space spanned by the ensemble, and g_A pre-multiplying $\widehat{P}$, so that the transformation is applied to the d-dimensional state-space instead. In both approaches we identify the relevant matrix by first writing

$$\widehat{\Pi}\widehat{\Pi}^\top = \mathscr{C}(\widehat{C}) = \widehat{P}(I - VD^{-1}V^\top)\widehat{P}^\top,$$

where $V = (A\widehat{P})^\top$ and $D = V^\top V + \Gamma$.

1. *Ensemble Transform Kalman Update:* taking $\widehat{\Pi} = \widehat{P}FU$ for any F satisfying $FF^\top = I - VD^{-1}V^\top$ and arbitrary orthogonal U satisfies the consistency condition. One approach for finding such a matrix F is by rewriting

$$I - VD^{-1}V^\top = (I + \widehat{P}^\top A^\top \Gamma^{-1} A\widehat{P})^{-1} = E(I+\Lambda)^{-1/2}(I+\Lambda)^{-1/2}E^\top = FF^\top,$$

where the first equality follows by the Sherman-Morrison formula, and $E\Lambda E^\top$ is the eigenvalue decomposition of $\widehat{P}^\top A^\top \Gamma^{-1} A\widehat{P}$. In summary, we have $g_T(\widehat{P}) = \widehat{P}E(I +$

$\Lambda)^{-1/2}U$.

2. *Ensemble Adjustment Kalman Update:* Introducing $M = V\Gamma^{-1/2}$, we can write

$$\widehat{P}(I - VD^{-1}V^\top)\widehat{P}^\top = \widehat{P}(I + MM^\top)^{-1}\widehat{P}^\top.$$

Noting that $\widehat{P}$ has full column rank, we may then define $B = \widehat{P}(I+MM^\top)^{-1/2}(\widehat{P}^\top)^\dagger$, and so

$$g_A(\widehat{P}) = B\widehat{P} = \widehat{P}(I + MM^\top)^{-1/2}(\widehat{P}^\top)^\dagger\widehat{P} = \widehat{P}(I + MM^\top)^{-1/2}.$$

Once a choice of g has been made, and an estimate $\widehat{\Sigma}$ has been computed, the updated ensemble has first two moments given by

$$\begin{aligned}\widehat{\mu} &= \mathscr{M}(\widehat{m}, \widehat{C}),\\ \widehat{\Sigma} &= \mathscr{C}(\widehat{C}).\end{aligned} \tag{2.12}$$

Frequently, only $\widehat{\mu}, \widehat{\Sigma}$ are of concern to the practitioner, but it is still possible to *back-out* the individual members of the updated ensemble as they may be of interest. It is clear that one choice for $\widehat{P}$ is

$$\widehat{P} = \frac{1}{\sqrt{N-1}}\left[u_1 - \widehat{m}, \cdots, u_N - \widehat{m}\right],$$

in which case it holds that $\widehat{P}1_N = 0_d$, and so

$$\upsilon_n = \sqrt{N-1}[\widehat{\Pi}]_n + \mathscr{M}(\widehat{m}, \widehat{C}), \qquad 1 \le n \le N, \tag{2.13}$$

where $[\widehat{\Pi}]_n$ denotes the n-th column of $\widehat{\Pi}$.

In Subsection 2.2.3 we establish error bounds for the approximation of the posterior mean and covariance (μ, Σ) in (2.2) by $(\widehat{\mu}, \widehat{\Sigma})$ as estimated using the PO and SR updates in (2.10) and (2.12). It is clear from (2.12) that as long as the choice of g is valid, in the sense that the resulting $\widehat{\Sigma}$ is consistent, then the specific choice of g is irrelevant to the accuracy of a single SR update. We therefore make no assumptions in our subsequent analysis of the SR algorithm beyond that of g satisfying the consistency condition. Note that, when compared to the SR update in (2.12), the PO update in (2.10) contains additional stochastic terms that will, as our bounds indicate, hinder the estimation of (μ, Σ). As noted in the literature, for example in Tippett et al. [2003], the PO update increases the probability of underestimating the analysis error covariance. While our presentation and analysis of PO and SR updates is carried out in the linear-Gaussian setting, both updates are frequently utilized in nonlinear and non-Gaussian settings, with empirical evidence suggesting that the PO updates can outperform SR updates Lawson and Hansen [2004], Leeuwenburgh et al. [2005]. In fact, the consistency argument outlined above is only valid in the linear case $\mathcal{G}(u) = Au$, and the statistical advantage of SR implementations in linear settings may not translate into the nonlinear case.

2.2.2 Dimension-Free Covariance Estimation

We define the *effective dimension* Wainwright [2019] of a matrix $Q \in \mathcal{S}_+^d$ by

$$r_2(Q) \equiv \frac{\mathrm{Tr}(Q)}{\|Q\|}. \tag{2.14}$$

The effective dimension quantifies the number of directions where Q has significant spectral content Tropp [2015]. The monographs Tropp [2015], Vershynin [2018] refer to $r_2(Q)$ as the intrinsic dimension, while Koltchinskii and Lounici [2017] uses the term effective rank. This terminology is motivated by the observation that $1 \leq r_2(Q) \leq \mathrm{rank}(Q) \leq d$ and that $r_2(Q)$ is insensitive to changes in the scale of Q, see Tropp [2015]. In situations where the eigenvalues of Q decay rapidly, $r_2(Q)$ is a better measure of dimension than the state dimension d. The following result [Koltchinskii and Lounici, 2017, Theorem 9] gives a non-asymptotic sufficient sample size requirement for accurate covariance estimation in terms of the effective dimension of the covariance matrix. We recall that the sample covariance estimator $\widehat{C}$ is defined in (2.3).

Proposition 2.2.1 (Covariance Estimation with Sample Covariance —Unstructured Case)**.** *Let $u_1, \dots, u_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[u_1] = m$ and var$[u_1] = C$. Then, for all $t \geq 1$, it holds with probability at least $1 - ce^{-t}$ that*

$$\|\widehat{C} - C\| \lesssim \|C\| \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

Remark 2.2.2 (Effective Dimension and Smoothness)**.** *Proposition 2.2.1 motivates defining $r_2(C) \equiv Tr(C)/\|C\|$ to be the effective dimension of a d-dimensional sub-Gaussian random vector u with var$[u] = C$. As in the definition for matrices, $r_2(C)$ quantifies the number of directions where the distribution of u has significant spread. Proposition 2.2.1 and our results in Subsection 2.2.3 may be extended to sub-Gaussian random variables defined in an infinite-dimensional separable Hilbert space, say $\mathcal{H} = L^2(0,1)$. It is then illustrative to note that any Gaussian measure $\mathcal{N}(m, C)$ in $\mathcal{H}$ satisfies that $Tr(C) < \infty$; in other words, all Gaussian measures have finite effective dimension. In this context, $r_2(C)$ is related to the rate of decay of the eigenvalues of C, and hence to the almost sure Sobolev regularity of functions u drawn from the Gaussian measure $\mathcal{N}(m, C)$ on $\mathcal{H} = L^2(0,1)$, see e.g. Bogachev [1998], Stuart [2010]. In computational inverse problems and data assimilation, u is often a d-dimensional vector that represents a fine discretization of a Gaussian random field; then, $r_2(C)$ quantifies the smoothness of the undiscretized field.*

2.2.3 Main Results: Posterior Approximation with Finite Ensemble

In this subsection we state finite ensemble approximation results for the posterior mean and covariance with PO and SR ensemble updates. To highlight some key insights, including the dependence of the bounds on the effective dimension of C and the differences between PO and SR updates, we opt to present expectation bounds in Theorems 2.2.3 and 2.2.5 that are less notationally cumbersome than the stronger exponential tail bounds in Theorems 2.5.8 and 2.5.9 in Appendix 2.5.3. Throughout this section, the data y is treated as a fixed quantity.

Theorem 2.2.3 (Posterior Mean Approximation with Finite Ensemble —Expectation Bound)**.** *Consider the PO and SR ensemble Kalman updates given by* (2.10) *and* (2.12)*, respectively, leading to an estimate* $\widehat{\mu}$ *of the posterior mean* μ *defined in* (2.2)*. Set* $\varphi = 1$ *for the PO update and* $\varphi = 0$ *for the SR update. Then, for any* $p \geq 1$*,*

$$\left[\mathbb{E}\|\widehat{\mu} - \mu\|_2^p\right]^{1/p} \lesssim_p c_1 \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^{3/2}\right) + \varphi c_2 \left(\sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(C)}{N}\sqrt{\frac{r_2(\Gamma)}{N}}\right), \tag{2.15}$$

where $c_1 = c_1(\|C\|, \|A\|, \|\Gamma^{-1}\|, \|y - Am\|_2)$ *and* $c_2 = c_2(\|C\|, \|A\|, \|\Gamma^{-1}\|)$*.*

Importantly, the bound (2.15) does not depend on the dimension d of the state-space, and the only dependence on C is through its operator norm and the effective dimension $r_2(C)$. The term multiplied by φ in the PO update accounts for the additional error incurred by the presence of the offset term (2.11) in the PO update (2.10). The following remark discusses another important consequence of Theorem 2.2.3: the stable performance of ensemble Kalman updates in small noise regimes when compared with other sampling algorithms.

Remark 2.2.4 (Dependence of Constants on Model Parameters)**.** *The proof of Theorem 2.2.3 in Appendix 2.5 provides an explicit definition of* c_1 *and* c_2 *up to constants, i.e. it describes* *how these quantities rely on their arguments, and Theorem 2.5.8 establishes a high probability bound on* $\|\widehat{\mu} - \mu\|_2$. *In particular, it is important to note that the constants* c_1 *and* c_2 *in Theorem 2.2.3 deteriorate in the small noise limit where the observation noise goes to zero, and* c_2 *deteriorates with* $r_2(\Gamma)$. *In the small noise limit, the posterior and prior distribution become mutually singular, and it is hence expected for ensemble updates to be unstable. To illustrate this intuition in a concrete setting, assume that* $\Gamma = \gamma I$ *for a positive constant* γ, *and, for simplicity, that* $N \geq r_2(C)$ *as well as* $\|C\| = \|A\| = \|y - Am\|_2 = 1$. *Then, the expression for* c_1 *established in Theorem 2.2.3 implies that for the SR update, for any error* $\varepsilon > 0$ *and* $p \geq 1$,

$$N \gtrsim \frac{r_2(C)}{\varepsilon^2 \gamma^4} \implies \mathbb{E}\big[\|\widehat{\mu} - \mu\|_2^p\big]^{1/p} \lesssim_p \varepsilon.$$

Similarly, the expressions for c_1 *and* c_2 *imply that for the PO update,*

$$N \gtrsim \frac{r_2(C)}{\varepsilon^2 \gamma^4} \vee \frac{k}{\varepsilon^2 \gamma} \implies \mathbb{E}\big[\|\widehat{\mu} - \mu\|_2^p\big]^{1/p} \lesssim_p \varepsilon,$$

where we recall that k *denotes the dimension of the data* y. *The papers Agapiou et al. [2017], Sanz-Alonso and Wang [2021] show the need to increase the sample size along small noise limits in importance sampling when target and proposal are given, respectively, by posterior and prior. While our bounds here only give sufficient rather than necessary conditions on the ensemble size, it is noteworthy that, for fixed* k, *the scaling of* N *as* $\gamma \to 0$ *shown here is independent of* k. *In contrast, necessary sample size conditions for importance sampling show a polynomial dependence on* k, *see Sanz-Alonso and Wang [2021].*

Theorem 2.2.5 (Posterior Covariance Approximation with Finite Ensemble —Expectation Bound)**.** *Consider the PO and SR ensemble Kalman updates given by* (2.10) *and* (2.12), *respectively, leading to an estimate* $\widehat{\Sigma}$ *of the posterior covariance* Σ *defined in* (2.2). *Set*

$\varphi = 1$ for the PO update and $\varphi = 0$ for the SR update. Then, for any $p \geq 1$,

$$\big[\mathbb{E}\|\widehat{\Sigma} - \Sigma\|^p\big]^{1/p} \lesssim_p c_1 \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^2 \right) + \varphi \mathscr{E},$$

where

$$\mathscr{E} = c_2 \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^3 \vee \left(\sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(\Gamma)}{N} \right) \left(1 \vee \left(\frac{r_2(C)}{N}\right)^2 \right) \right),$$

where $c_1 = c_1(\|C\|, \|A\|, \|\Gamma^{-1}\|)$ and $c_2 = c_2(\|C\|, \|A\|, \|\Gamma^{-1}\|, \|\Gamma\|)$.

As in Theorem 2.2.3, the bound in Theorem 2.2.5 does not depend on the dimension d of the state-space, and the dependence on C is through the operator norm and the effective dimension $r_2(C)$.

Remark 2.2.6 (Dependence of Constants on Model Parameters)**.** *The proof of Theorem 2.2.5 in Appendix 2.5 provides an explicit definition of c_1 and c_2 up to constants and Theorem 2.5.9 establishes a high probability bound on $\|\widehat{\Sigma} - \Sigma\|$. As discussed in Remark 2.2.4, these bounds may be used to establish sufficient ensemble size requirements in small noise limits and other singular limits of practical importance.*

Remark 2.2.7 (Comparison to the Literature)**.** *The results in this section complement many of the existing analyses of ensemble Kalman updates in the literature. In one direction, our Theorems 2.2.3 and 2.2.5 can be viewed in the context of [Furrer and Bengtsson, 2007, Section 3.4], which claims that for finite ensembles the square root filter is always more efficient than the perturbed observation filter, since the latter introduces additional variability through noisy perturbations of the data. Our results quantify this additional variability both in probability and in expectation. In Majda and Tong [2018], the authors put forward a non-asymptotic analysis of a multi-step* EnKF *augmented by a spectral projection step in which the Kalman gain matrix is projected onto the linear span of its leading eigenvalues exceeding a* *threshold level. They refer to the dimension* d_{subs} *of this subspace as the effective dimension and provide guarantees on the performance of the algorithm so long as the ensemble size scales with* d_{subs}*. In contrast, our one-step analysis does not require any augmentation of standard implementations (see e.g. Furrer and Bengtsson [2007]) of the ensemble update. They also employ (forecast) covariance inflation, which is a de-biasing technique standard in the literature, see for example [Furrer and Bengtsson, 2007, Section 5], which our results do not require. In another direction, our results can be directly compared to [Kwiatkowski and Mandel, 2015, Theorem 6.1], which states that for iteration* t *of the square root* EnKF *and for any* $p \geq 1$

$$\Big[\mathbb{E}\|\hat{\mu}^{(t)} - \mu^{(t)}\|_2^p\Big]^{1/p} \leq \frac{c(p,t)}{\sqrt{N}} \qquad \textit{and} \qquad \Big[\mathbb{E}\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|^p\Big]^{1/p} \leq \frac{c(p,t)}{\sqrt{N}}, \tag{2.16}$$

where $\hat{\mu}^{(t)}$ *and* $\widehat{\Sigma}^{(t)}$ *are the sample mean and covariance of the updated (analysis) ensemble at iteration* t*, and* $\mu^{(t)}$ *and* $\Sigma^{(t)}$ *are the corresponding Kalman Filter posterior mean and covariance, respectively. The term* $c(p,t)$ *that arises in both of their bounds denotes a constant that depends only on* p*, the iteration index* t*, and the norms of the non-random inputs of the algorithm, but do not depend on dimension or ensemble size. Importantly, they do not distinguish between settings with different effective dimensions as our bounds do. In Appendix 2.5.4, we provide an explicit outline of the multi-step algorithm considered in their paper along with definitions of all quantities described here. As previously noted, our bounds cover the perturbed observation setting whereas* (2.16) *is specific to the square root setting. In Appendix 2.5.4 we also establish (see Corollary 2.5.12) a simple extension of Theorems 2.2.3 and 2.2.5 to the multi-step square root setting, which shows that for any* $p \geq 1$*, iteration* t*,* *and assuming for simplicity that $N \gtrsim r_2(\Sigma^{(0)})$, then*

$$\Big[\mathbb{E}\|\hat{\mu}^{(t)} - \mu^{(t)}\|_2^p\Big]^{1/p} \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|, \|y^{(l)} - A^{(l)}m^{(l)}\|_2\}_{l=1}^t, \|\Gamma^{-1}\|),$$
$$\Big[\mathbb{E}\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|^p\Big]^{1/p} \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|\}_{l=1}^t, \|\Gamma^{-1}\|). \tag{2.17}$$

Our bounds therefore refine those in Kwiatkowski and Mandel [2015] as they explicitly capture the dependence on the state dimension through the effective dimension of the initial distribution, $r_2(\Sigma^{(0)})$. It follows then that in the case of the square root EnKF*, it suffices to use an ensemble on the order of the effective dimension of $\Sigma^{(0)}$ multiplied by constants depending on the operator norms of the forward model matrices $\{\|A^{(l)}\|\}_{l=1}^t$, analysis covariance matrices $\{\|\Sigma^{(l)}\|\}_{l=1}^t$, inverse of the noise covariance, $\|\Gamma^{-1}\|$ and ℓ_2-norm of the model errors $\{\|y^{(l)} - A^{(l)}m^{(l)}\|_2\}_{l=1}^t$. We note that extensions to the multi-step setting for other variants of the* EnKF *that do not use SR updates may not follow as easily. In this direction, the recent work Al-Ghattas et al. [2024a] studies a multi-step* EnKF *with PO updates which incorporates an additional resampling step.*

2.3 Ensemble Kalman Updates: Sequential Optimization Algorithms

In the optimization approach, the solution to the inverse problem (2.1) is found by minimizing an objective function. As discussed in Chada et al. [2021], an entire suite of ensemble algorithms have been derived that differ in the choice of objective function and optimization scheme. In this subsection we introduce the Ensemble Kalman Inversion (EKI) algorithm Iglesias et al. [2013] and a new localized implementation of EKI, which we call localized EKI (LEKI) following Tong and Morzfeld [2023]. Both EKI and LEKI use an ensemble approximation of a Levenberg-Marquardt (LM) optimization scheme to minimize a data-misfit objective

$$\mathsf{J}(u) = \frac{1}{2}\|\Gamma^{-1/2}\big(y - \mathcal{G}(u)\big)\|_2^2, \tag{2.18}$$

which promotes fitting the data y. Before deriving EKI the in Subsection 2.3.1 and LEKI in Subsection 2.3.1, we give some background that will help us interpret both methods as ensemble-based implementations of classical gradient-based LM schemes. The finite ensemble approximation of an idealized mean-field EKI update using EKI and LEKI updates will be studied in Subsection 2.3.3.

Recall that classical iterative optimization algorithms choose an initialization $u^{(0)}$ and set

$$u^{(t+1)} = u^{(t)} + w^{(t)}, \qquad t = 0, 1, \dots, \tag{2.19}$$

until a pre-specified convergence criterion is met. Here, $w^{(t)}$ is some favorable direction determined by the optimization algorithm at iteration t, given the current estimate $u^{(t)}$. In the case that the inverse problem is ill-posed, directly minimizing (2.18) leads to a solution that over-fits the data. Then, implicit regularization can be achieved through the optimization scheme used to obtain the update $w^{(t)}$. Under the assumption that $r(u) \equiv y - \mathcal{G}(u)$ is differentiable, the Levenberg-Marquardt (LM) algorithm chooses $w^{(t)}$ by solving the constrained minimization problem

$$\min_{w} \mathsf{J}_t^{\text{lin}}(w) \quad \text{subject to} \quad \|C^{-1/2}w\|_2^2 \leq \delta_l,$$

where

$$\mathsf{J}_t^{\mathrm{lin}}(w) \equiv \frac{1}{2}\|Dr(u^{(t)})w + r(u^{(t)})\|_2^2,$$

and Dr denotes the Jacobian of r. The LM algorithm belongs to the class of trust region optimization methods, and it chooses each increment to minimize a linearized objective, $\mathsf{J}_t^{\mathrm{lin}}$, but with the added constraint that the minimizer belongs to the ball $\{\|C^{-1/2}w\|^2 \leq \delta_l\}$, in which we *trust* that the objective may be replaced by its linearization. Equivalently, $w^{(t)}$ can be viewed as the unconstrained minimizer of a regularized objective,

$$\min_w \mathsf{J}_t^{\mathrm{U}}(w), \qquad \mathsf{J}_t^{\mathrm{U}}(w) \equiv \mathsf{J}_t^{\mathrm{lin}}(w) + \frac{1}{2\alpha_t}\|C^{-1/2}w\|_2^2, \tag{2.20}$$

where $\alpha_t > 0$ acts as a Lagrange multiplier.

We are interested in ensemble sequential-optimization algorithms, which instead of updating a single estimate $u^{(t)}$ —as in (2.19)— propagate an *ensemble* of estimates. Ensemble-based optimization schemes often rely on *statistical linearization* to avoid the computation of derivatives. Underpinning this idea Ungarala [2012], Chada et al. [2021], Kim et al. [2023] is the argument that if $\mathcal{G}(u) = Au$ were linear, then $\widehat{C}^{\,up} = \widehat{C}A^\top$, leading to the approximation in the general nonlinear case

$$D\mathcal{G}(u_n) \approx (\widehat{C}^{\,up})^\top \widehat{C}^\dagger \equiv G. \tag{2.21}$$

This approximation motivates the derivative-free label often attached to ensemble-based algorithms Kovachki and Stuart [2019], and we note that they may be employed whenever computing $D\mathcal{G}(u)$ is expensive or when $\mathcal{G}$ is not differentiable. For the remainder, our analysis focuses on a single step of EKI and LEKI, and so we drop the iteration index t from our notation; we will use instead our previous terminology of prior ensemble and updated ensemble. Finally, similar to our presentation of posterior-approximation algorithms, our exposition is simplified by introducing the *nonlinear gain-update* operator $\mathscr{P}$,

$$\mathscr{P} : \mathbb{R}^{d\times k} \times \mathcal{S}_+^k \to \mathbb{R}^{d\times k}, \qquad \mathscr{P}(C^{up}, C^{pp}; \Gamma) = \mathscr{P}(C^{up}, C^{pp}) = C^{up}(C^{pp} + \Gamma)^{-1}, \tag{2.22}$$

which is shown to be both pointwise continuous and bounded in Lemma 2.5.7.

2.3.1 Ensemble Algorithms for Sequential Optimization

Ensemble Kalman Inversion Update

In the EKI, each particle in the prior ensemble is updated according to the LM algorithm, so that

$$v_n = u_n + w_n, \qquad 1 \le n \le N,$$

where w_n is the minimizer of a linearized and regularized data-misfit objective

$$\mathsf{J}_n^{\text{lin}}(w) = \frac{1}{2}\|\Gamma^{-1/2}\big(y - \eta_n - \mathcal{G}(u_n) - Gw\big)\|_2^2 + \frac{1}{2\alpha}\|\widehat{C}^{-1/2}w\|_2^2, \qquad \eta_n \sim \mathcal{N}(0, \Gamma). \tag{2.23}$$

Following Iglesias et al. [2013], we henceforth set $\alpha = 1$, but note that our main results can be readily extended to any $\alpha > 0$. Note that each ensemble member solves the optimization (2.23) with a perturbed observation $y - \eta_n$, similar in spirit to the PO update of Subsection 2.2.1. The minimizer of (2.23) (with $\alpha = 1$) is given by

$$w_n = \widehat{C}G^\top(G\widehat{C}G^\top + \Gamma)^{-1}\big(y - \eta_n - \mathcal{G}(u_n)\big).$$

Substituting $\widehat{C}G^\top = \widehat{C}^{\,up}$, and approximating

$$G\widehat{C}G^\top = G\widehat{C}^{\,up} = (\widehat{C}^{\,up})^\top \widehat{C}^\dagger \widehat{C}^{\,up} \approx \widehat{C}^{\,pp}$$

leads to the EKI update

$$\upsilon_n = u_n + \mathscr{P}(\widehat{C}^{\,up}, \widehat{C}^{\,pp})\big(y - \mathcal{G}(u_n) - \eta_n\big), \qquad 1 \le n \le N. \tag{2.24}$$

In the linear forward-model setting, $\mathscr{P}(\widehat{C}^{\,up}, \widehat{C}^{\,pp}) = \mathscr{K}(\widehat{C})$, and (2.24) takes on a form identical to the PO update in (2.10). We further define the *mean-field* EKI update

$$\upsilon_n^* = u_n + \mathscr{P}(C^{\,up}, C^{\,pp})\big(y - \mathcal{G}(u_n) - \eta_n\big), \qquad 1 \le n \le N, \tag{2.25}$$

which is the update that would be performed if one had access to the population quantities C^{up} and C^{pp} or, equivalently, to an infinite ensemble. We will analyze the approximation of the update (2.24) to the mean-field update (2.25) in Subsection 2.3.3. The study of mean-field ensemble Kalman methods of the form (2.25) was proposed in Herty and Visconti [2019] and is overviewed in Calvello et al. [2022]. While mean-field algorithms are not useful for practical implementation, they facilitate a transparent mathematical analysis that can provide understanding on the performance of practical ensemble approximations. Desirable properties of mean-field algorithms include convergence to the desired target in a continuous-time limit Carrillo and Vaes [2021], a gradient flow structure Garbuno-Inigo et al. [2020], or the ability to approximate derivative-based optimization algorithms Chada et al. [2021]. The transfer of theoretical insights from mean-field algorithms to particle-based algorithms tacitly presupposes, however, that the ensemble is large enough for ensemble-based updates to approximate well idealized mean-field updates. In this direction, Ding and Li [2021] establishes a $\mathcal{O}(N^{-1/2})$ rate for an approximation of a mean-field evolution equation in terms of the ensemble size N. Our first main result of this section, Theorem 2.3.5, will show that the mean-field update (2.25) can be well approximated with the EKI update (2.24) with a number of particles of the order of the *effective dimension* of the problem, which is defined as for posterior-approximation algorithms.

Localized Ensemble Kalman Inversion Update

In practice, ensemble-based algorithms are often implemented with $N \ll d$, that is, with an ensemble that is much smaller than the state dimension. In this setting, the update is augmented with an additional *localization* procedure applied to $\widehat{C}$ in the case of linear forward model, and to both $\widehat{C}^{pp}$ and $\widehat{C}^{up}$ in the case of a nonlinear forward model. In either case, localization is seen as an approach to deal both with the extreme rank deficiency and the sampling error that arise from using an ensemble that is significantly smaller than the dimension of the state and/or the dimension of the observation, see for example Houtekamer and Mitchell [2001], Houtekamer and Zhang [2016], Farchi and Bocquet [2019]. Localization is also useful when the state u, or the transformed state $\mathcal{G}(u)$, has elements $\mathcal{E}(i)$ and $\mathcal{E}(j)$ that represent the values of a variable of interest at physical locations that are a known distance $\mathsf{d}(i,j)$ apart: correlations may decay quickly with the physical distance of the variables and localization may help to remove spurious correlations in the sample covariance estimator. In ensemble Kalman methods, localization has most commonly been carried out via the Schur (elementwise) product of the estimator and a positive-semidefinite matrix M of equal dimension. In the vast majority of cases, the elements of M are taken to be $\mathsf{M}_{ij} = \kappa(\mathsf{d}(i,j)/b)$, where κ is a locally supported correlation function —usually the Gaspari Cohn 5^{th}-order compact piecewise polynomial Gaspari and Cohn [1999]— and $b > 0$ is a length-scale parameter chosen by the practitioner. Since κ tapers off to zero as its argument becomes larger, i.e. when the underlying variables are further apart, the Schur-product operation zeroes out the corresponding elements of the estimator, and the rate at which this tapering occurs is controlled by the size of the length-scale. The localized EKI (LEKI), recently studied in Tong and Morzfeld [2023], replaces both $\widehat{C}^{pp}$ and $\widehat{C}^{up}$ with their localized counterparts, $\mathsf{M}_1 \circ \widehat{C}^{pp}$ and $\mathsf{M}_2 \circ \widehat{C}^{up}$, where M_1 and M_2 are localization matrices of appropriate dimension. Two important issues have, in our opinion, hindered the rigorous study of localized ensemble algorithms, and we highlight these next before moving on to introduce our localization framework.

1. Optimality: The justification outlined above for localization in the ensemble Kalman literature has been largely heuristic, and relying on these arguments alone one cannot hope to define a localization procedure that is demonstrably optimal. Notably, the widespread usage of the Gaspari-Cohn correlation function is not rooted in any sense of optimality. Generally, focusing solely on a band of entries near the diagonal is a sub-optimal approach to covariance estimation, as noted in the high-dimensional covariance estimation literature, see for example Chen et al. [2012], Levina and Vershynin [2012], Bickel and Levina [2008b]. Moreover, even in cases where focusing on elements near the diagonal is justified, for example by assuming that the underlying target is a banded matrix, the bandwidth $b > 0$ must be chosen carefully as a function of the ensemble size, problem dimension, and dependence structure Bickel and Levina [2008a]. This type of analysis has, to the best of our knowledge, not been carried out for the Gaspari-Cohn localization scheme. An important message in the covariance estimation literature is that localization —regardless of how it is employed— can only be optimal if the target of estimation itself is sparse, and such sparsity assumptions must be made explicit in order to facilitate a rigorous mathematical analysis of the procedure. The difficulty of optimal localization in ensemble updates has also been highlighted in Furrer and Bengtsson [2007], where the authors derive an optimal localization matrix M under the unrealistic assumption that C is a diagonal matrix.

2. Schur-Product Approximations: In the literature on ensemble Kalman methods, a

consensus has not been reached on how best to apply localization in practice. The issue here can be sufficiently described by deferring to the linear forward-model setting, i.e. $\mathcal{G}(u) = Au$, in which the Kalman gain is a central quantity. As mentioned for example in Houtekamer and Mitchell [2001], in a localized update, the Kalman gain operator should in theory be applied to $\mathsf{M} \circ \widehat{C}$, i.e. one should study the quantity

$$\mathscr{K}(\mathsf{M} \circ \widehat{C}) = (\mathsf{M} \circ \widehat{C})A^\top \big(A(\mathsf{M} \circ \widehat{C})A^\top + \Gamma\big)^{-1},$$

although their experimental results are based on the more computationally convenient approximation

$$\mathscr{K}(\mathsf{M} \circ \widehat{C}) \approx \big(\mathsf{M} \circ (\widehat{C}A^\top)\big)\big(\mathsf{M} \circ (A\widehat{C}A^\top) + \Gamma\big)^{-1}, \tag{2.26}$$

which, as they mention, is a reasonable approximation in the case that A is diagonal. Subsequently, much of the literature on localization in ensemble Kalman updates has adopted this or similar approximations, as discussed in greater depth in [Petrie, 2008, Section 3.3]. In general, however, approximations made on the Schur product are difficult to justify without strong assumptions on the forward model $\mathcal{G}$.

With these issues in mind, we opt to study an alternative, data-driven approach to localization often employed in the high-dimensional covariance estimation literature Bickel and Levina [2008a], Cai and Zhou [2012a,b], where it is referred to as *thresholding.* We ground our analysis in the assumption that the target of estimation belongs to the following soft sparsity matrix class:

$$\mathscr{U}_{d_1,d_2}(q, R_q) \equiv \left\{ B \in \mathbb{R}^{d_1 \times d_2} : \max_{i \le d_1} \sum_{j=1}^{d_2} |B_{ij}|^q \le R_q \right\}, \tag{2.27}$$

where $q \in [0, 1)$ and $R_q > 0$, and write $\mathscr{U}_d(q, R_q)$ in the case $d_1 = d_2 = d$. In the special case

$q = 0$, matrices in $\mathscr{U}_{d_1,d_2}(0, R_0)$ possess rows that have no more than R_0 non-zero entries —a special case of which are banded matrices— which is the classical hard-sparsity constraint. In contrast, for $q \in (0, 1)$, the class $\mathscr{U}_{d_1,d_2}(q, R_q)$ contains matrices with rows belonging to the ℓ_q ball of radius R_q^q. This includes matrices with rows that contain possibly many non-zero entries so long as their magnitudes decay sufficiently rapidly, and so is often referred to as a *soft*-sparsity constraint. Importantly, the class $\mathscr{U}_d(q, R_q)$ is sufficiently rich to capture the motivating intuition that correlations decay with physical distance in a rigorous manner that avoids the optimality issues mentioned above. Structured covariance matrices, such as those belonging to $\mathscr{U}_{d_1,d_2}(q, R_q)$ are optimally estimated using localized versions of their sample covariances. To this end, we study the localized matrix estimator $B_{\rho_N} \equiv \mathcal{L}_{\rho_N}(B)$, where $\mathcal{L}_{\rho_N}(u) = u\mathbf{1}_{\{|u|\geq\rho_N\}}$ is a localization operator with localization radius ρ_N, and which is applied elementwise to its argument B. In Section 2.3 we detail how the localization radius ρ_N can be chosen optimally in terms of the parameters of the inverse problem (2.1) and the ensemble size N.

Throughout our analysis, we refrain from using approximations such as the one outlined in (2.26); that is, our analysis of localization replaces all non-localized quantities in the original update (2.24) with their localized counterparts. We introduce the LEKI update:

$$\upsilon_n^\rho = u_n + \mathscr{P}(\widehat{C}^{up}_{\rho_{N,1}}, \widehat{C}^{pp}_{\rho_{N,2}})\big(y - \mathcal{G}(u_n) - \eta_n\big), \qquad 1 \leq n \leq N, \tag{2.28}$$

where $\rho_{N,1}$ and $\rho_{N,2}$ are two, potentially different localization radii. As in the non-localized case, in Subsection 2.3.3 we provide finite sample bounds on the deviation of the LEKI update from the mean-field update of (2.25), and describe in detail how the additional structure imposed on C^{up} and C^{pp} leads to improved bounds relative to the non-localized setting. Our second main result of this section, Theorem 2.3.7 will be based on new covariance estimation bounds that may be of independent interest, and on a suitable notion of effective dimension that we introduce in Subsection 2.3.2. Our theory explains the improved sample complexity that can be achieved by simultaneously exploiting spectral decay and sparsity of the covariance model.

An important issue that warrants discussion is that of positive-semidefiniteness of the estimator $\widehat{B}_{\rho_N}$ when the target B is a square covariance matrix. In the case of the Schur-product estimator, any localization matrix M derived from a valid correlation function κ is guaranteed to be positive-semidefinite by definition Gaspari and Cohn [1999], and so by the Schur-product Theorem [Horn and Johnson, 2012, Theorem 7.5.3] the estimator $\mathsf{M} \circ \widehat{B}$ is positive-semidefinite as well. In contrast, the localization operator $\mathcal{L}_{\rho_N}$ thresholds the sample covariance $\widehat{B}$ elementwise and does not in general preserve positive-semidefiniteness. As discussed in El Karoui [2008], Cai and Zhou [2012b], $\widehat{B}_{\rho_N}$ is positive-semidefinite with high probability, but in practice one may opt to use an augmented estimator that guarantees positive-semidefiniteness. We describe this estimator here for completeness: let $\widehat{B}_{\rho_N} = \sum_{j=1}^d \hat{\lambda}_j v_j v_j^\top$ be the eigen-decomposition of $\widehat{B}_{\rho_N}$, so that λ_j, v_j are the j-th eigenvalue and eigenvector of $\widehat{B}_{\rho_N}$. Consider then the positive-part estimator $\widehat{B}^+_{\rho_N} \equiv \sum_{j=1}^d (0 \vee \hat{\lambda}_j) v_j v_j^\top$. Clearly then, $\widehat{B}^+_{\rho_N}$ is positive-semidefinite, and furthermore it achieves the same rate as $\widehat{B}_{\rho_N}$ since

$$\|\widehat{B}^+_{\rho_N} - B\| \leq \|\widehat{B}^+_{\rho_N} - \widehat{B}_{\rho_N}\| + \|\widehat{B}_{\rho_N} - B\| \leq \max_{j:\hat{\lambda}_j<0} |\hat{\lambda}_j - \lambda_j| + \|\widehat{B}_{\rho_N} - B\| \leq 2\|\widehat{B}_{\rho_N} - B\|,$$

where λ_j is the j-th eigenvalue of B. In light of this fact, we abuse notation slightly and assume that B_{ρ_N} is positive-semidefinite throughout this work.

2.3.2 Dimension-Free Covariance Estimation Under Soft Sparsity

For the covariance estimation problem under (approximate) sparsity, there are estimators that significantly improve upon the sample covariance. In particular, [Wainwright, 2019, Chapter 6.5] notes that for sub-Gaussian data the operator-norm covariance estimation error depends logarithmically on the state dimension d for localized estimators, while the error depends linearly on d for the sample covariance. If no sparse structure is assumed, the effective dimension r_2 defined in (2.14) characterizes the error of the sample covariance estimator, as described in Proposition 2.2.1. We introduce an analogous notion of effective dimension that is more suitable than r_2 in the sparse covariance estimation problem, termed the *max-log effective dimension* and which, for $Q \in \mathcal{S}_+^d$, is given by

$$r_\infty(Q) \equiv \frac{\max_{j \le d} Q_{(j)} \log(j+1)}{Q_{(1)}},$$

where $Q_{(1)} \ge Q_{(2)} \ge \ldots \ge Q_{(d)}$ is the decreasing rearrangement of the diagonal entries of Q. To the best of our knowledge, this notion of dimension has not been previously considered in the literature, and, as will be shown, refines the rate of covariance estimation under sparsity by incorporating intrinsic properties of the underlying matrix, albeit differently to (2.14). In particular, $r_\infty(Q)$ is small whenever Q exhibits a decay of the ordered elements $Q_{(1)}, Q_{(2)}, \ldots$ that is faster than $\log(j+1)$. We use the subscript ∞ to highlight that the quantity r_∞ is related to the dimension-free sub-Gaussian maxima result of Lemma 2.6.6. Similarly, we use the subscript 2 to draw the connection between r_2 and the sub-Gaussian 2-norm concentration of Theorem 2.5.1. Importantly, bounds based on r_∞ will be dimension-free, in the sense that they exhibit no dependence on the state dimension d. The next result is our analog of Proposition 2.2.1 for estimation under sparsity using the localized sample covariance estimator. Recall that $C_{(1)}$ denotes the largest element on the diagonal of C, $\widehat{C}_{\rho_N} \equiv \mathcal{L}_{\rho_N}(\widehat{C})$ denotes the localized sample covariance matrix, and $\mathscr{U}_d(q, R_q)$ is the sparse matrix class defined in (2.27). All proofs in this subsection have been deferred to Appendix 2.6.1.

Theorem 2.3.1 (Covariance Estimation with Localization —Soft Sparsity Assumption)**.** *Let $u_1, \ldots, u_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[u_1] = m$ and $var[u_1] = C$. Further, assume that $C \in \mathscr{U}_d(q, R_q)$ for some $q \in [0, 1)$ and $R_q > 0$. For any*

$t \geq 1$, set

$$\rho_N \asymp C_{(1)} \left(\sqrt{\frac{r_\infty(C)}{N}} \vee \frac{tr_\infty(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right)$$

and let $\widehat{C}_{\rho_N} \equiv \mathcal{L}_{\rho_N}(\widehat{C}_{\rho_N})$ be the localized sample covariance estimator. There exists a constant $c > 0$ such that, with probability at least $1 - ce^{-t}$,

$$\|\widehat{C}_{\rho_N} - C\| \lesssim R_q \rho_N^{1-q}.$$

Remark 2.3.2 (Max-Log Effective Dimension)**.** *The proof of Theorem 2.3.1 can be found in Section 2.6.1 and, up to the choice of ρ_N, follows an identical approach to the standard proof for localized covariance estimators in the literature, for example [Wainwright, 2019, Theorem 6.27]. The result depends crucially on the order of the maximum elementwise distance between the sample and true covariance matrices, $\|\widehat{C} - C\|_{\max}$, which is where our analysis differs from the exiting literature. Our proof utilizes techniques in Koltchinskii and Lounici [2017] combined with the dimension-free sub-Gaussian maxima bound of Lemma 2.6.6 to obtain a bound in terms of r_∞. In the worst case, for example when $C = cI_d$ for some constant $c > 0$ so that the ordered diagonal elements of C exhibit no decay, we recover exactly the standard logarithmic dependence on the state dimension. In particular, when $N \geq r_\infty(C)(= \log d)$, Theorem 2.3.1 matches the result for recovering C in operator norm in the sub-Gaussian setting over the class $\mathscr{U}_d(q, R_q)$, as shown in [Bickel and Levina, 2008a, Theorem 1]. If the ordered variances exhibit sufficiently fast decay, our upper bound is significantly better. (Recall that in many applications $d \sim 10^9$ and $N \sim 10^2$, and so the logarithmic dependence on d may play a significant role in determining a sufficient ensemble size.) Importantly, many of the results in the structured covariance estimation literature rely similarly on the maximum elementwise norm, and so our results can be utilized to achieve refined bounds on the estimation error of the localized estimator under structural assumptions on C*

that differ from the soft-sparsity assumption considered in this work.

A result analogous to Theorem 2.3.1 holds for cross-covariance estimation under sparsity. For a formal statement we refer to Theorem 2.6.11 whose proof is based on a deep generic chaining bound for product empirical processes [Mendelson, 2016, Theorem 1.12]. Here we present a cross-covariance estimation result that is specific to the LEKI setting in that it relies on a smoothness assumption on the forward model.

Theorem 2.3.3 (Cross-Covariance Estimation with Localization —Soft Sparsity Assumption)**.** *Let $u_1, \dots, u_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[u_1] = m$ and $var[u_1] = C$. Let $\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^k$ be a Lipschitz continuous forward model and assume that $C^{up} \in \mathcal{U}_{d,k}(q_1, R_{q_1})$ and $C^{pu} \in \mathcal{U}_{k,d}(q_2, R_{q_2})$ where $q_1, q_2 \in [0, 1)$ and R_{q_1}, R_{q_2} are positive constants. For any $t \geq 1$, set*

$$\rho_N \asymp (C_{(1)} \vee C^{pp}_{(1)}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(C)} \vee \sqrt{r_\infty(C^{pp})} \right) \vee \sqrt{\frac{r_\infty(C)}{N}} \sqrt{\frac{r_\infty(C^{pp})}{N}} \right),$$

and let $\widehat{C}^{up}_{\rho_N} \equiv \mathcal{L}_{\rho_N}(\widehat{C}^{up})$ be the localized sample cross-covariance estimator. There exist positive universal constants c_1, c_2 such that, with probability at least $1 - c_1 e^{-c_2 t}$,

$$\|\widehat{C}^{up}_{\rho_N} - C^{up}\| \lesssim R_{q_1} \rho_N^{1-q_1} \vee R_{q_2} \rho_N^{1-q_2}.$$

Remark 2.3.4 (Sparsity of the Cross-Covariance)**.** *To the best of our knowledge, estimation of the cross-covariance matrix under structural assumptions has not been a point of focus in the literature. Indeed, one may implicitly estimate the cross-covariance by applying Theorem 2.3.1 to the full covariance matrix*

$$\begin{bmatrix} C & C^{up} \\ C^{pu} & C^{pp} \end{bmatrix}$$

of the sub-Gaussian vector $[u^\top, \mathcal{G}(u)^\top]^\top$*, and extracting a bound on* $\|C^{up}_{\rho_N} - C^{up}\|$*. This approach however requires one to place sparsity assumptions on the full covariance matrix, making the result potentially less useful in practice. That is, one may wish to make structural assumptions on* C^{up} *and* C^{pp} *without imposing any restrictions on* C*, which our result allows for.*

2.3.3 Main Results: Approximation of Mean-Field Particle Updates with Finite Ensemble Size

In this subsection we state finite ensemble approximation results for EKI and LEKI updates. The main results, Theorems 2.3.5 and 2.3.7, showcase the dependence on the effective dimension of C and C^{pp} for EKI and on the max-log dimension of these matrices for LEKI. For both algorithms, we study the update of a generic particle u_n and the analysis is carried out conditional on both u_n and the noise perturbation η_n.

Theorem 2.3.5 (Approximation of Mean-Field EKI with EKI —Operator-Norm Bound)**.** *Let* y *be generated according to* (2.1) *with Lipschitz forward model* $\mathcal{G}: \mathbb{R}^d \to \mathbb{R}^k$*. Let* υ_n *and* υ_n^* *be the EKI and mean-field EKI updates defined in* (2.24) *and* (2.25) *respectively. Then, for any* $t \geq 1$*, there exists universal positive constants* c_1, c_2 *such that, with probability at least* $1 - c_1 e^{-c_2 t}$*,*

$$\|\upsilon_n - \upsilon_n^*\|_2 \leq c_1 \left(\frac{c_2}{N} \vee \sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{r_2(C^{pp})}{N}} \vee \frac{r_2(C^{pp})}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right),$$

where $c_1 = c_1(\|y - \mathcal{G}(u_n) - \eta_n\|_2, \|\Gamma^{-1}\|, \|C\|, \|C^{up}\|, \|C^{pp}\|)$ *and for* $u \sim \mathcal{N}(m, C)$*,* $c_2 = c_2(\|u_n\|_2, \|m\|_2, \|\mathcal{G}(u_n)\|_2, \|\mathbb{E}[\mathcal{G}(u)]\|_2)$*.*

Remark 2.3.6 (Dependence of Constants on Model Parameters)**.** *The proof of Theorem 2.3.5 in Appendix 2.6.2 gives an explicit expression for the dependence of c on its arguments.*

These bounds may be used to establish the sufficient ensemble size to ensure that the EKI update approximates well the mean-field EKI update in the unstructured covariance setting.

Theorem 2.3.7 (Approximation of Mean-Field EKI with LEKI —Operator-Norm Bound)**.** *Let y be generated according to (2.1) with Lipschitz forward model $\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^k$. Assume that $C^{up} \in \mathscr{U}_{d,k}(q_1, R_{q_1})$, $C^{pu} \in \mathscr{U}_{k,d}(q_2, R_{q_2})$ and $C^{pp} \in \mathscr{U}_k(q_3, R_{q_3})$ for $q_1, q_2, q_3 \in [0, 1)$, and positive constants $R_{q_1}, R_{q_2}, R_{q_3}$. Let υ_n^ρ and υ_n^* be the LEKI and mean-field EKI updates outlined in (2.28) and (2.25) respectively. For any $t \geq 1$, set*

$$\rho_{N,1} = \rho_{N,2}$$
$$\asymp \frac{c_1}{N} + (C_{(1)} \vee C_{(1)}^{pp}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(C)} \vee \sqrt{r_\infty(C^{pp})} \right) \vee \sqrt{\frac{r_\infty(C)}{N}} \sqrt{\frac{r_\infty(C^{pp})}{N}} \right),$$

and

$$\rho_{N,3} \asymp \frac{c_2}{N} + C_{(1)}^{pp} \left(\sqrt{\frac{r_\infty(C^{pp})}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{t r_\infty(C^{pp})}{N} \right),$$

where $c_1 = c_1(\|u_n\|_\infty, \|m\|_\infty, \|\mathcal{G}(u_n)\|_\infty, \|\mathbb{E}[\mathcal{G}(u)]\|_\infty)$ and $c_2 = c_2(\|\mathcal{G}(u_n)\|_\infty, \|\mathbb{E}[\mathcal{G}(u)]\|_\infty)$, with $u \sim \mathcal{N}(m, C)$. There exist positive universal constants c_3, c_4 such that, with probability at least $1 - c_3 e^{-c_4 t}$,

$$\|\upsilon_n^\rho - \upsilon_n^*\|_2 \leq c_5 (R_{q_1} \rho_{N,1}^{1-q_1} \vee R_{q_2} \rho_{N,2}^{1-q_2} \vee R_{q_3} \rho_{N,3}^{1-q_3}),$$

where $c_5 = c_5(\|y - \mathcal{G}(u_n) - \eta_n\|_2, \|\Gamma^{-1}\|, \|C^{up}\|)$.

Remark 2.3.8 (Dependence of Constants on Model Parameters)**.** *The proof of Theorem 2.3.7 in Appendix 2.6.2 gives an explicit expression for the dependence of c on its arguments. As discussed in Remark 2.3.6, these bounds may be used to establish the sufficient ensemble size to ensure that the LEKI update approximates well the mean-field EKI update in the structured covariance setting.*

Remark 2.3.9 (On the Soft-Sparsity Assumptions). *Importantly, Theorem 2.3.7 makes no assumptions on the covariance matrix C, and so can be used even in cases where C is dense, but the covariances C^{up}, C^{pu}, and C^{pp} can be reasonably assumed to be sparse. In the case that sparsity assumptions on C are appropriate, then an interesting question is: what (explicit) assumptions on $\mathcal{G}$ ensure sparsity of C^{up}, C^{pu}, and C^{pp}? We provide here two simple arguments that may provide some insight. Throughout, c_1, c_2, c_2, c_4, c_5 are arbitrary positive constants independent of both state and observation dimensions d and k, and $q \in [0, 1)$.*

1. *Suppose $C \in \mathscr{U}_d(q, c_1)$ and $\mathbb{E}[D\mathcal{G}]^\top \in \mathscr{U}_{d,k}(q, c_2)$. Then there exists c_3 such that $C^{up} \in \mathscr{U}_{d,k}(q, c_3)$. We provide a formal statement of this result in Lemma 2.6.14. Similarly, if $\mathbb{E}[D\mathcal{G}] \in \mathscr{U}_{k,d}(q, c_4)$, then there exists c_5 such that $C^{pu} \in \mathscr{U}_{k,d}(q, c_5)$. The assumptions on the expected Jacobian $\mathbb{E}[D\mathcal{G}]$ can be understood as the requirements that, in expectation:*

   (a) Any coordinate function $\mathcal{G}_j$ of $\mathcal{G}$ depends on its input u only through a subset of u whose size does not grow with k nor d.

   (b) Any state coordinate u_j of u is acted on only by a subset of the coordinate-functions of $\mathcal{G}$ whose size does not grow with k nor d.

   For example, a Jacobian that is banded in expectation would satisfy these two properties.

2. *Suppose $C \in \mathscr{U}_d(q, c_1)$. Then there exists c_2 such that $C^{pp} \in \mathscr{U}_k(q, c_2)$ whenever $\mathcal{G}(u) = Au$ is a linear map with $A \in \mathscr{U}_{k,d}(q, c_3)$ and $A^\top \in \mathscr{U}_{d,k}(q, c_4)$, i.e. whenever A has both rows and columns that are sparse. This condition holds, for example, for banded A. We provide a formal statement of this result in Lemma 2.6.16.*

The two arguments above indicate that if $\mathcal{G}$ acts on local subsets of u, which holds for instance for convolution or moving average operators, then one can expect the sparsity of C to carry on to C^{up}, C^{pu}, and C^{pp}.

Remark 2.3.10 (Comparison to the Literature)**.** *Although the focus of this subsection is the LEKI, it is useful to compare our Theorems 2.3.5 and 2.3.7 to existing results for the performance of ensemble based algorithms with localization. In this regard, our results are closest to those of Tong [2018], which shows that an ensemble that scales with the logarithm of the state dimension times a localization radius suffices for good performance of the localized* EnKF *(LEnKF). They study performance over multiple time steps and linear dynamics under a stability assumption which enforces control over the model matrices as well as a sparse ($q = 0$) structure of the underlying true covariance matrices. They consider domain localization whereas we study covariance localization. In contrast to our results, Tong [2018] employs covariance localization and utilizes a Schur-product localization scheme in which elements whose indices are beyond a certain bandwidth are set to zero, whereas we study localization via thresholding (recall our discussion comparing these two approaches in Subsection 2.3.1). Consequently, our required localization radius is in terms of the max-log effective dimension whereas theirs is in terms of the bandwidth of the underlying covariance matrix. Our results are dimension-free in that they do not rely on the state dimension d, and so as noted in Remark 2.3.2, our bounds can have significantly better than logarithmic dependence on dimension. Our setting also differs from Tong [2018] in that our dynamics are allowed to be nonlinear, and our prior ensemble can be sub-Gaussian as opposed to Gaussian. Related to this point is that the analysis in Tong [2018] does not account for noise introduced from adding perturbations to the ensemble update, which is justified by a law of large numbers argument; however in the non-asymptotic and nonlinear settings, it is likely that one must account for this noise especially when considering the covariance between the current ensemble and the perturbation noise at a given iteration of the algorithm. We view it as an important avenue to extend the results of this subsection to a multi-step analysis, and a particularly important question is whether dimension-free control of the LEnKF can be rigorously shown utilizing a combination of our results and those of Tong [2018]. The LEKI* *has also been recently studied in Tong and Morzfeld [2023] under a nonlinear, multi-step setting. The authors study convergence of the iterates to a global minimizer and the rate of collapse of the ensemble. They argue that localization is a remedy for the "subspace property" of the EKI, which refers to the fact that ensembles at any given iteration live in the linear subspace spanned by the initial ensemble, which cannot capture the true state if $N < d$. Their analysis differs from ours in that they study the continuous-time setting whereas we analyze discrete time updates as implemented in practice. Further, while they discuss that the size of the ensemble may be much smaller than the state dimension, as well as illustrate this with simulations, they do not provide an explicit characterization of the sufficient ensemble size. Our results also show that the LEKI is close to the mean field version of the problem, which is not considered in their set-up. An interesting open question is whether the results of this section can be used in conjunction with results in Tong and Morzfeld [2023] to provide a sufficient ensemble size for LEKI over multiple iterations.*

2.4 Conclusions, Discussion, and Future Directions

This chapter has introduced a non-asymptotic approach to the study of ensemble Kalman methods. Our theory explains why these algorithms may be accurate provided that the ensemble size is larger than a suitable notion of effective dimension, which may be dramatically smaller than the state dimension due to spectrum decay and/or approximate sparsity. Our non-asymptotic results in Section 2.2 tell apart PO and SR updates for posterior approximation, and our results in Section 2.3 demonstrate the potential advantage of using localization in sequential-optimization algorithms.

As discussed in Subsection 2.3.1, localization is also often used in posterior-approximation algorithms. For instance, one may define a localized PO update by

$$\begin{aligned} \widehat{\mu} &= \mathscr{M}(\widehat{m}, \widehat{C}_{\rho_N}) - \mathscr{K}(\widehat{C}_{\rho_N})\overline{\eta}, \\ \widehat{\Sigma} &= \mathscr{C}(\widehat{C}_{\rho_N}) + \widehat{O}_{\rho_N}, \end{aligned} \tag{2.29}$$

where $\widehat{O}_{\rho_N}$ is defined replacing $\widehat{C}$ with $\widehat{C}_{\rho_N}$ in (2.11). Similarly, one may define a localized SR update by

$$\begin{aligned} \widehat{\mu} &= \mathscr{M}(\widehat{m}, \widehat{C}_{\rho_N}), \\ \widehat{\Sigma} &= \mathscr{C}(\widehat{C}_{\rho_N}). \end{aligned} \tag{2.30}$$

It is then natural to ask if localized PO and SR updates can yield better approximation of the posterior mean and covariance than those without localization in Theorems 2.2.3 and 2.2.5. The answer for the posterior mean seems to be negative.

To see why, consider for intuition that we are given a random sample $X_1, \ldots, X_N$ from a normal distribution with mean μ^X and covariance Σ^X with the objective to estimate μ^X. Standard results, see e.g. [L. E. Lehmann, and G. Casella, 2006, Example 1.14], show that the sample mean $\overline{X}$ is minimax optimal for ℓ_2-loss regardless of whether or not Σ^X is known. In other words, the minimax rate of estimating μ^X can be achieved without making use of information regarding Σ^X. It follows then that placing assumptions on Σ^X can lead to impressive improvements in the covariance estimation problem (as shown in Section 2.3) but cannot be expected to affect the mean estimation problem. Similarly, in our inverse problem setting, sparsity assumptions on the prior covariance C cannot be expected to translate into a better bound on $\|\widehat{\mu} - \mu\|_2$: this quantity is a function of both the covariance deviation $\|\widehat{C}_{\rho_N} - C\|$ and the prior mean deviation $\|\widehat{m} - m\|_2$ and since the latter is unaffected it dominates the overall bound, yielding an error bound of the same order as that in Theorem 2.2.3. As discussed in Remark 2.2.7, a potential avenue for future investigation is to utilize techniques introduced in this manuscript to study alternative localization schemes in the posterior approximation setting, such as *domain localization* considered in Tong [2018]. In short, covariance localization as defined in (2.30) does not lead to improved bounds for the posterior-approximation problem.

Similar issues to those arising in the estimation of the posterior mean affect the analysis of the localized offset $\widehat{O}_{\rho_N}$, and we therefore do not expect improvement on the bound in Theorem 2.2.5 for covariance estimation with the localized PO update. We note, however, that for localized SR it is possible to derive an analog to the high probability version of Theorem 2.2.5 (see Theorem 2.5.9) with an improved error bound, which we present in Theorem 2.7.2.

Our discussion here should not be taken to imply that localization in posterior-approximation algorithms is not useful; it is plausible that localization in one step of the algorithm can lead to improved bounds in later steps, and we leave this multi-step analysis of localized posterior approximation ensemble updates as an important line for future work. A related phenomenon is known to occur in sequential Monte Carlo, where a proposal density that may be optimal for one step of the filter may not be optimal over multiple steps Agapiou et al. [2017]. Another interesting direction for future study is the non-asymptotic analysis of ensemble Kalman methods for likelihood approximations in state-space models Chen et al. [2022]. Finally, we envision that the non-asymptotic approach set forth here may be adopted to design and analyze new multi-step methods for posterior-approximation and sequential-optimization in inverse problems and data assimilation.

Proofs

We provide proofs of all theorems in the main body. We will use the following result extensively and summarise it here for brevity. Given events $E_1, \ldots, E_J$ that each occur with probability at least $1 - ce^{-t}$, where $t \geq 1$ and $c > 0$ is a universal constant that may be different for each event, then

$$\mathbb{P}\left(\bigcap_{j=1}^{J} E_j\right) = 1 - \mathbb{P}\left(\bigcup_{j=1}^{J} \bar{E}_j\right) \geq 1 - \sum_{j=1}^{J} \mathbb{P}(\bar{E}_j) \geq 1 - ce^{-t}.$$

2.5 Proofs: Section 2.2

This appendix contains the proofs of all the theorems in Section 2.2. Background results on covariance estimation are reviewed in Subsection 2.5.1 and the continuity and boundedness of the Kalman gain, mean-update, covariance-update, and nonlinear gain-update operators are summarized in Subsection 2.5.2. These preliminary results are used in Subsection 2.5.3 to establish our main theorems.

2.5.1 Preliminaries: Concentration and Covariance Estimation

Theorem 2.5.1 (Sub-Gaussian Norm Concentration, [Vershynin, 2018, Exercise 6.3.5])**.** *Let X be a d-dimensional sub-Gaussian random vector with $\mathbb{E}[X] = \mu^X$, $var[X] = \Sigma^X$. Then, for any $t \geq 1$, with probability at least $1 - ce^{-t}$ it holds that*

$$\|X - \mu^X\|_2 \lesssim \sqrt{Tr(\Sigma^X)} + \sqrt{t\|\Sigma^X\|} \lesssim \sqrt{\|\Sigma^X\|(r_2(\Sigma^X) \vee t)}\,.$$

Proof of Proposition 2.2.1. For $n = 1, \ldots, N$, let $u_n = Z_n + m$, where Z_n is a centered sub-Gaussian random vector with $\text{var}[Z_n] = C$. Then we may write

$$\widehat{C} = \frac{1}{N-1}\sum_{n=1}^{N}(Z_n - \overline{Z})(Z_n - \overline{Z})^\top \asymp \frac{1}{N}\sum_{n=1}^{N} Z_n Z_n^\top - \overline{Z}\,\overline{Z}^\top \equiv \widehat{C}^0 - \overline{Z}\,\overline{Z}^\top.$$

Therefore,

$$\|\widehat{C} - C\| \leq \|\widehat{C}^0 - C\| + \|\overline{Z}\,\overline{Z}^\top\| = \|\widehat{C}^0 - C\| + \|\overline{Z}\|_2^2.$$

Let E_1 denote the event on which

$$\|\widehat{C}^0 - C\| \lesssim \|C\| \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right),$$

and E_2 the event on which

$$\|\overline{Z}\|_2^2 \lesssim \|C\| \left(\frac{r_2(C)}{N} \vee \frac{t}{N} \right).$$

Then by Theorem 9 of Koltchinskii and Lounici [2017], $\mathbb{P}(E_1) \geq 1 - e^{-t}$, and by Theorem 2.5.1, $\mathbb{P}(E_2) \geq 1 - e^{-t}$. Therefore, the result holds on $E_1 \cap E_2$, which has probability at least $1 - ce^{-t}$. $\square$

Lemma 2.5.2 (Sample Covariance Operator Norm Bound)**.** *Let $u_1, \ldots, u_N$ and $\widehat{C}$ be as in Proposition 2.2.1. Then, for any $t \geq 1$, it holds with probability at least $1 - ce^{-t}$ that*

$$\|\widehat{C}\| \lesssim \|C\| \left(1 \vee \frac{r_2(C)}{N} \vee \frac{t}{N} \right).$$

Proof. By the triangle inequality $\|\widehat{C}\| \leq \|\widehat{C} - C\| + \|C\|$. The result follows by Proposition 2.2.1 noting that, for any $x \geq 0$, $1 \vee \sqrt{x} \vee x = 1 \vee x$. $\square$

Lemma 2.5.3 (Cross-Covariance Estimation —Unstructured Case)**.** *Let $u_1, \ldots, u_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[u_1] = m$ and $var[u_1] = C$. Let $\eta_1, \ldots, \eta_N$ be k-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[\eta_1] = 0$ and $var[\eta_1] = \Gamma$, and assume that the two sequences are independent. Consider the estimator*

$$\widehat{C}^{u\eta} = \frac{1}{N-1} \sum_{n=1}^{N} (u_n - \widehat{m})(\eta_n - \overline{\eta})^\top$$

of the cross-covariance $C^{u\eta} \equiv \mathbb{E}\big[(u_1 - m)\eta_1^\top\big]$. Then there exists a constant c such that, for *all $t \geq 1$, it holds with probability at least $1 - ce^{-t}$ that*

$$\|\widehat{C}^{u\eta} - C^{u\eta}\| \lesssim (\|C\| \vee \|\Gamma\|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(\Gamma)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

Proof. First, we note that

$$\widehat{C}^{u\eta} \asymp \frac{N-1}{N} \left(\frac{1}{N} \sum_{n=1}^{N} (u_n - \widehat{m})(\eta_n - \bar{\eta})^\top \right) \equiv \frac{N-1}{N} \widetilde{C}^{u\eta},$$

and so it suffices to prove the claim for the biased sample covariance estimator, which we denote by $\widetilde{C}^{u\eta}$. Letting $Z_n = u_n - m$, it follows that

$$\|\widetilde{C}^{u\eta}\| = \left\| \frac{1}{N} \sum_{n=1}^{N} Z_n \eta_n^\top - \overline{Z} \bar{\eta}^\top \right\| \leq \left\| \frac{1}{N} \sum_{n=1}^{N} Z_n \eta_n^\top \right\| + \|\overline{Z} \bar{\eta}^\top\|. \tag{2.31}$$

For the second term in the right-hand side of (2.31), let E_1 denote the event on which

$$\|\overline{Z}\|_2 \lesssim \sqrt{\|C\| \left(\frac{r_2(C)}{N} \vee \frac{t}{N} \right)},$$

and E_2 the event on which

$$\|\bar{\eta}\|_2 \lesssim \sqrt{\|\Gamma\| \left(\frac{r_2(\Gamma)}{N} \vee \frac{t}{N} \right)},$$

each of which have probability at least $1 - e^{-t}$ by Theorem 2.5.1. Therefore, the event $E_1 \cap E_2$ occurs with probability at least $1 - ce^{-t}$, and on which it follows that

$$\|\overline{Z} \bar{\eta}^\top\| = \|\overline{Z}\|_2 \|\bar{\eta}\|_2 \lesssim (\|C\| \vee \|\Gamma\|) \left(\frac{r_2(C)}{N} \vee \frac{r_2(\Gamma)}{N} \vee \frac{t}{N} \right),$$

where the inequality follows since $\sqrt{ab} \lesssim a \vee b$ for $a, b \geq 0$. To control the first term in the right-hand side of (2.31), we define the vector

$$W_n = \begin{bmatrix} Z_n \\ \eta_n \end{bmatrix} \in \mathbb{R}^{d+k}, \qquad 1 \le n \le N,$$

and note that $W_1, \dots, W_N$ is an i.i.d. sub-Gaussian sequence with $\mathbb{E}[W_1] = [m^\top, 0_k^\top]^\top$ and variance $C^W = \operatorname{diag}(C, \Gamma)$. Let E_3 denote the event on which

$$\begin{aligned}
\left\| \frac{1}{N} \sum_{n=1}^{N} W_n W_n^\top - C^W \right\| &\lesssim \|C^W\| \left(\sqrt{\frac{r_2(C^W)}{N}} \vee \frac{r_2(C^W)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \\
&\lesssim (\|C\| \vee \|\Gamma\|) \left(\left(\sqrt{\frac{\operatorname{Tr}(C)}{N\|C\|}} + \sqrt{\frac{\operatorname{Tr}(\Gamma)}{N\|\Gamma\|}} \right) \vee \frac{\operatorname{Tr}(C) + \operatorname{Tr}(\Gamma)}{N(\|C\| \vee \|\Gamma\|)} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \\
&\lesssim (\|C\| \vee \|\Gamma\|) \left(\left(\sqrt{\frac{r_2(C)}{N}} + \sqrt{\frac{r_2(\Gamma)}{N}} \right) \vee \left(\frac{r_2(C)}{N} + \frac{r_2(\Gamma)}{N} \right) \vee \left(\sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \right) \\
&\lesssim (\|C\| \vee \|\Gamma\|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(\Gamma)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).
\end{aligned}$$

By Proposition 2.2.1, it holds for any $t \ge 1$ that $\mathbb{P}(E_3) \ge 1 - e^{-t}$. Note that we can express

$$\mathcal{P} \equiv \frac{1}{N} \sum_{n=1}^{N} W_n W_n^\top - \begin{bmatrix} C & O \\ O & \Gamma \end{bmatrix} = \begin{bmatrix} N^{-1} \sum_{n=1}^{N} Z_n Z_n^\top - C & N^{-1} \sum_{n=1}^{N} Z_n \eta_n^\top \\ N^{-1} \sum_{n=1}^{N} \eta_n Z_n^\top & N^{-1} \sum_{n=1}^{N} \eta_n \eta_n^\top - \Gamma \end{bmatrix},$$

and that

$$\left\| \frac{1}{N} \sum_{n=1}^{N} Z_n \eta_n^\top \right\| = \|E_{11} \mathcal{P} E_{12}\| \le \|E_{11}\| \|\mathcal{P}\| \|E_{12}\| = \|\mathcal{P}\|,$$

where E_{11}, E_{12} are block *selection* matrices that pick the relevant sub-block matrix of $\mathcal{P}$.

Therefore, it holds on E_3 that

$$\left\| \frac{1}{N} \sum_{n=1}^{N} Z_n \eta_n^\top \right\| \lesssim (\|C\| \vee \|\Gamma\|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(\Gamma)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

The final result follows by noting that the intersection $E_1 \cap E_2 \cap E_3$ has probability at least $1 - ce^{-t}$. □

2.5.2 Continuity and Boundedness of Update Operators

The next three lemmas, shown in Kwiatkowski and Mandel [2015], ensure the continuity and boundedness of the Kalman gain, mean-update, and covariance-update operators introduced in Section 2.2. We include them here for completeness. Lemma 2.5.7 below establishes similar properties for the nonlinear gain-update operator introduced in Section 2.3.

Lemma 2.5.4 (Continuity and Boundedness of Kalman Gain Operator [Kwiatkowski and Mandel, 2015, Lemma 4.1 & Corollary 4.2])**.** *Let $\mathscr{K}$ be the Kalman gain operator defined in* (2.5)*. Let $P, Q \in \mathcal{S}_+^d$, $\Gamma \in \mathcal{S}_{++}^k$, and $A \in \mathbb{R}^{k \times d}$. The following hold:*

$$\begin{aligned}
\|\mathscr{K}(Q) - \mathscr{K}(P)\| &\leq \|Q - P\| \|A\| \|\Gamma^{-1}\| \Big(1 + \min\left(\|P\|, \|Q\|\right) \|A\|^2 \|\Gamma^{-1}\|\Big), \\
\|\mathscr{K}(Q)\| &\leq \|Q\| \|A\| \|\Gamma^{-1}\|, \\
\|I - \mathscr{K}(Q)A\| &\leq 1 + \|Q\| \|A\|^2 \|\Gamma^{-1}\|.
\end{aligned}$$

Lemma 2.5.5 (Continuity and Boundedness of Mean-Update Operator [Kwiatkowski and Mandel, 2015, Corollary 4.3 & Lemma 4.7])**.** *Let $\mathscr{M}$ be the mean-update operator defined in*

(2.6). *Let* $P, Q \in \mathcal{S}_+^d$*,* $\Gamma \in \mathcal{S}_{++}^k$*,* $A \in \mathbb{R}^{k\times d}$*,* $y \in \mathbb{R}^k$*, and* $m, m' \in \mathbb{R}^d$*. The following hold:*

$$\begin{aligned}
\|\mathscr{M}(m,Q)\| &\leq \|m\| + \|Q\|\|A\|\|\Gamma^{-1}\|\,\|y - Am\|_2\,,\\
\left\|\mathscr{M}(m,Q) - \mathscr{M}(m',P)\right\| &\leq \left\|m - m'\right\|\left(1 + \|A\|^2\|\Gamma^{-1}\|\|Q\|\right)\\
&\quad + \|Q - P\|\|A\|\|\Gamma^{-1}\|\left(1 + \|A\|^2\|\Gamma^{-1}\|\|P\|\right)\left\|y - Am'\right\|_2.
\end{aligned}$$

Lemma 2.5.6 (Continuity and Boundedness of Covariance-Update Operator [Kwiatkowski and Mandel, 2015, Lemma 4.4 & Lemma 4.6])**.** *Let* $\mathscr{C}$ *be the covariance-update operator defined in* (2.7)*. Let* $P, Q \in \mathcal{S}_+^d$*,* $\Gamma \in \mathcal{S}_{++}^k$*,* $A \in \mathbb{R}^{k\times d}$*,* $y \in \mathbb{R}^k$*, and* $m, m' \in \mathbb{R}^d$*. The following hold:*

$$\begin{aligned}
\|\mathscr{C}(Q) - \mathscr{C}(P)\| &\leq \|Q - P\|\Big(1 + \|A\|^2\|\Gamma^{-1}\|(\|Q\| + \|P\|) + \|A\|^4\|\Gamma^{-1}\|^2\|Q\|\|P\|\Big),\\
0 \preccurlyeq \mathscr{C}(Q) &\preccurlyeq Q,\\
\|\mathscr{C}(Q)\| &\leq \|Q\|.
\end{aligned}$$

Lemma 2.5.7 (Continuity and Boundedness of Nonlinear Gain-Update Operator)**.** *Let* $\mathscr{P}$ *be the nonlinear gain-update operator defined in* (2.22)*. Let* $P, \widetilde{P} \in \mathbb{R}^{d\times k}$*,* $Q, \widetilde{Q} \in \mathcal{S}_+^k$*, and* $\Gamma \in \mathcal{S}_{++}^k$*. The following hold:*

$$\begin{aligned}
\|\mathscr{P}(P,Q) - \mathscr{P}(\widetilde{P},\widetilde{Q})\| &\leq \|\Gamma^{-1}\|\|P - \widetilde{P}\| + \|\Gamma^{-1}\|^2\|P\|\|Q - \widetilde{Q}\|,\\
\|\mathscr{P}(P,Q)\| &\leq \|\Gamma^{-1}\|\|P\| + \|\Gamma^{-1}\|^2\|Q\|.
\end{aligned}$$

Proof. The proof follows in similar style to Lemma 4.1 in Kwiatkowski and Mandel [2015]. We note that

$$\begin{aligned}
\|P(Q+\Gamma)^{-1} - \widetilde{P}(\widetilde{Q}+\Gamma)^{-1}\| &\leq \|P(Q+\Gamma)^{-1} - P(\widetilde{Q}+\Gamma)^{-1}\| + \|P(\widetilde{Q}+\Gamma)^{-1} - \widetilde{P}(\widetilde{Q}+\Gamma)^{-1}\|\\
&\leq \|P\|\|(Q+\Gamma)^{-1} - (\widetilde{Q}+\Gamma)^{-1}\| + \|\widetilde{P} - P\|\|(Q+\Gamma)^{-1}\|.
\end{aligned}$$

Since $\Gamma \succ 0$ and $Q \succeq 0$, it holds that $Q + \Gamma \succeq \Gamma$ and so $(Q+\Gamma)^{-1} \preccurlyeq \Gamma^{-1}$, which in turn implies $\|(Q+\Gamma)^{-1}\| \le \|\Gamma^{-1}\|$. Further,

$$
\begin{aligned}
\|(Q+\Gamma)^{-1} - (\widetilde{Q}+\Gamma)^{-1}\| &= \|\Gamma^{-1/2}[(\Gamma^{-1/2}Q\Gamma^{-1/2}+I)^{-1} - (\Gamma^{-1/2}\widetilde{Q}\Gamma^{-1/2}+I)^{-1}]\Gamma^{-1/2}\| \\
&\le \|\Gamma^{-1}\|\|(\Gamma^{-1/2}Q\Gamma^{-1/2}+I)^{-1} - (\Gamma^{-1/2}\widetilde{Q}\Gamma^{-1/2}+I)^{-1}\| \\
&\le \|\Gamma^{-1}\|\|\Gamma^{-1/2}Q\Gamma^{-1/2} - \Gamma^{-1/2}\widetilde{Q}\Gamma^{-1/2}\| \\
&\le \|\Gamma^{-1}\|^2\|Q-\widetilde{Q}\|,
\end{aligned}
$$

where the second to last equality follows by the fact that $\|(I+A)^{-1} - (I+B)^{-1}\| \le \|B-A\|$ for $A, B \in \mathcal{S}_+^k$. To prove the pointwise boundedness of $\mathscr{P}$, take $\widetilde{P}$ to be the $d \times k$ matrix of zeroes, and $\widetilde{Q}$ to be the $k \times k$ matrix of zeroes, and plug these values into the continuity bound. □

2.5.3 Proof of Main Results in Section 2.2

Theorem 2.5.8 (Posterior Mean Approximation with Finite Ensemble —High Probability Bound)**.** *Consider the PO and SR ensemble Kalman updates given by* (2.10) *and* (2.12)*, respectively, leading to an estimate $\widehat{\mu}$ of the posterior mean μ defined in* (2.2)*. Set $\varphi = 1$ for the PO update and $\varphi = 0$ for the SR update. Then there exists a constant c such that, for all $t \ge 1$, it holds with probability at least $1 - ce^{-t}$ that*

$$
\begin{aligned}
\|\widehat{\mu} - \mu\|_2 &\lesssim (\|C\|^{1/2} \vee \|C\|^2)(\|A\| \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|y - Am\|_2) \\
&\times \left(\sqrt{\frac{r_2(C)}{N}} \vee \sqrt{\frac{t}{N}} \vee \left(\frac{r_2(C)}{N}\right)^{3/2} \vee \left(\frac{t}{N}\right)^{3/2} \vee \frac{r_2(C)}{N}\sqrt{\frac{t}{N}} \vee \sqrt{\frac{r_2(C)}{N}}\frac{t}{N} \right) + \varphi\mathscr{E},
\end{aligned}
$$

where

$$\mathscr{E} = \|A\|\|\Gamma^{-1}\|\|\Gamma\|^{1/2}\|C\|$$
$$\times \left(\sqrt{\frac{r_2(\Gamma)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{r_2(C)}{N}\sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(C)}{N}\sqrt{\frac{t}{N}} \vee \frac{t}{N}\sqrt{\frac{r_2(\Gamma)}{N}} \vee \left(\frac{t}{N}\right)^{3/2} \right).$$

Proof. It follows from Lemma 2.5.5 that

$$\begin{aligned}
\|\widehat{\mu} - \mu\|_2 &= \|\mathscr{M}(\widehat{m}, \widehat{C}) - \varphi\mathscr{K}(\widehat{C})\bar{\eta} - \mathscr{M}(m, C)\|_2 \\
&\le \|\mathscr{M}(\widehat{m}, \widehat{C}) - \mathscr{M}(m, C)\|_2 + \varphi\|\mathscr{K}(\widehat{C})\bar{\eta}\|_2 \\
&\le \|\widehat{m} - m\|_2 \left(1 + \|A\|^2\|\Gamma^{-1}\|\|\widehat{C}\|\right) && (2.32) \\
&+ \|\widehat{C} - C\|\|A\|\|\Gamma^{-1}\| \left(1 + \|A\|^2\|\Gamma^{-1}\|\|C\|\right) \|y - Am\|_2 && (2.33) \\
&+ \varphi\|\mathscr{K}(\widehat{C})\| \, \|\bar{\eta}\|_2 \, . && (2.34)
\end{aligned}$$

We now control each of the terms in equations (2.32), (2.33), and (2.34) separately. For (2.32), we note that $\widehat{m} - m \sim \mathcal{N}(0, C/N)$. Let E_1 be the set on which

$$\|\widehat{m} - m\|_2 \lesssim \sqrt{\|C\| \left(\frac{r_2(C)}{N} \vee \frac{t}{N}\right)},$$

let E_2 be the set on which

$$\|\widehat{C} - C\| \lesssim \|C\| \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right)$$

and

$$\|\widehat{C}\| \lesssim \|C\| \left(1 \vee \frac{r_2(C)}{N} \vee \frac{t}{N}\right),$$

and let E_3 be the set on which

$$\|\bar{\eta}\|_2 \lesssim \sqrt{\|\Gamma\| \left(\frac{r_2(\Gamma)}{N} \vee \frac{t}{N} \right)}.$$

By Theorem 2.5.1, Proposition 2.2.1, and Lemma 2.5.2, the set $E = E_1 \cap E_2 \cap E_3$ has probability at least $1 - ce^{-t}$, and it holds on this set that (2.32) is bounded above by

$$(\|C\|^{1/2} \vee \|C\|^{3/2})(1 \vee \|A\|^2 \|\Gamma^{-1}\|) \Bigg(\sqrt{\frac{r_2(C)}{N}} \vee \sqrt{\frac{t}{N}} \vee \left(\frac{r_2(C)}{N} \right)^{3/2} \vee \left(\frac{t}{N} \right)^{3/2} \vee \frac{r_2(C)}{N} \sqrt{\frac{t}{N}} \vee \sqrt{\frac{r_2(C)}{N} \frac{t}{N}} \Bigg). \tag{2.35}$$

Further, on the set E we can bound (2.33) above by

$$(\|C\| \vee \|C\|^2)(\|A\| \vee \|A\|^3) \left(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2 \right) \|y - Am\| \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right). \tag{2.36}$$

Finally, for (2.34), it follows from Lemma 2.5.4,

$$\|\mathscr{K}(\widehat{C})\| \, \|\bar{\eta}\| \leq \|A\| \|\Gamma^{-1}\| \|\widehat{C}\| \, \|\bar{\eta}\|$$

and so on the set E, we can show that (2.34) is bounded above by $\mathcal{E}$. Putting the bounds (2.35), (2.36) together we see that on E, it holds that

$$\begin{aligned} \|\widehat{\mu} - \mu\|_2 \lesssim & (\|C\|^{1/2} \vee \|C\|^2)(\|A\| \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|y - Am\|_2) \\ & \times \left(\sqrt{\frac{r_2(C)}{N}} \vee \sqrt{\frac{t}{N}} \vee \left(\frac{r_2(C)}{N} \right)^{3/2} \vee \left(\frac{t}{N} \right)^{3/2} \vee \frac{r_2(C)}{N} \sqrt{\frac{t}{N}} \vee \sqrt{\frac{r_2(C)}{N} \frac{t}{N}} \right) \\ & + \varphi \mathcal{E}. \end{aligned}$$

□

Proof of Theorem 2.2.3. Recall that from Theorem 2.5.8, for all $t \geq 1$ with probability at least $1 - ce^{-t}$,

$$
\begin{aligned}
\|\widehat{\mu} - \mu\|_2 \lesssim\ & (\|C\|^{1/2} \vee \|C\|^2)(\|A\| \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|y - Am\|_2) \\
& \times \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^{3/2} \vee \sqrt{\frac{t}{N}} \vee \left(\frac{t}{N}\right)^{3/2} \vee \frac{r_2(C)}{N}\sqrt{\frac{t}{N}} \vee \sqrt{\frac{r_2(C)}{N}\frac{t}{N}} \right) + \varphi\mathscr{E}.
\end{aligned}
$$

For notational brevity, let

$$
\mathcal{W} \equiv (\|C\|^{1/2} \vee \|C\|^2)(\|A\| \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|y - Am\|_2),
$$

and let $B \equiv \mathcal{W}\left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^{3/2}\right)$. Then, for $\varphi = 0$ and $p \geq 1$,

$$
\begin{aligned}
\mathbb{E}\big[\|\widehat{\mu} - \mu\|_2^p\big] &= p\int_0^\infty x^{p-1}\mathbb{P}(\|\widehat{\mu} - \mu\|_2 > x)dx \\
&\leq p\int_0^B x^{p-1}dx + p\int_B^\infty x^{p-1}\mathbb{P}(\|\widehat{\mu} - \mu\|_2 > x)dx \\
&\lesssim B^p + p\int_0^\infty x^{p-1}\exp\left(-\min\left(\frac{Nx^2}{\mathcal{W}^2}, \frac{Nx^{2/3}}{\mathcal{W}^{2/3}}, \frac{N^3x^2}{\mathcal{W}^2 r_2^2(C)}, \frac{N^{3/2}x}{\mathcal{W}\sqrt{r_2(C)}}\right)\right)dx \\
&= B^p + p\max\left\{\frac{1}{2}\Gamma\left(\frac{p}{2}\right)\left(\frac{\mathcal{W}}{\sqrt{N}}\right)^p, \frac{1}{2}\Gamma\left(\frac{3p}{2}\right)\left(\frac{\mathcal{W}}{N^{3/2}}\right)^p,\right. \\
&\qquad\qquad \left.\frac{1}{2}\Gamma\left(\frac{p}{2}\right)\left(\frac{\mathcal{W}r_2(C)}{N^{3/2}}\right)^p, \Gamma(p)\left(\frac{\mathcal{W}\sqrt{r_2(C)}}{N^{3/2}}\right)^p\right\},
\end{aligned}
$$

where the final equality follows by direct integration. It follows then that

$$
\big[\mathbb{E}\|\widehat{\mu} - \mu\|_2^p\big]^{1/p} \lesssim B + c(p)\mathcal{W}\max\left(\frac{1}{\sqrt{N}}, \frac{1}{N^{3/2}}, \frac{r_2(C)}{N^{3/2}}, \frac{\sqrt{r_2(C)}}{N^{3/2}}\right) \lesssim c(p)B,
$$

where the final inequality holds since $r_2(C) \geq 1$. The result for the $\varphi = 1$ case is identical and thus omitted. The constants in the statement of the result are then:

$$c_1 = (\|C\|^{1/2} \vee \|C\|^2)(\|A\| \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|y - Am\|_2),$$
$$c_2 = \|A\|\|\Gamma^{-1}\|\|\Gamma\|^{1/2}\|C\|. \qquad \square$$

Theorem 2.5.9 (Posterior Covariance Approximation with Finite Ensemble —High Probability Bound)**.** *Consider the PO and SR ensemble Kalman updates given by* (2.10) *and* (2.12)*, respectively, leading to an estimate* $\widehat{\Sigma}$ *of the posterior covariance* Σ *defined in* (2.2)*. Set* $\varphi = 1$ *for the PO update and* $\varphi = 0$ *for the SR update. For any* $t \geq 1$*, it holds with probability at least* $1 - ce^{-t}$ *that*

$$\begin{aligned}\|\widehat{\Sigma} - \Sigma\| \lesssim\, &(\|C\| \vee \|C\|^3)(\|A\|^2 \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)\\ &\times \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^2 \vee \sqrt{\frac{t}{N}} \vee \left(\frac{t}{N}\right)^2\right) + \varphi\mathscr{E},\end{aligned}$$

where

$$\begin{aligned}\mathscr{E} = &(\|A\| \vee \|A\|^3)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(\|C\| \vee \|\Gamma\|)(\|C\| \vee \|C\|^2)\\ \times &\left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^3 \vee \sqrt{\frac{t}{N}} \vee \left(\frac{t}{N}\right)^3 \vee \left(\sqrt{\frac{r_2(\Gamma)}{N}} \vee \frac{r_2(\Gamma)}{N}\right)\left(1 \vee \left(\frac{r_2(C)}{N}\right)^2 \vee \left(\frac{t}{N}\right)^2\right)\right).\end{aligned}$$

Proof. From Proposition 4 of Furrer and Bengtsson [2007], for the PO-ensemble Kalman update we may write

$$\widehat{\Sigma} = \mathscr{C}(\widehat{C}) + \widehat{O},$$

while for the SR-ensemble Kalman update we have $\widehat{\Sigma} = \mathscr{C}(\widehat{C})$. We deal initially with the $\mathscr{C}(\widehat{C})$ term that is common to both expressions, and then proceed to show how the operator norm of the additional $\widehat{O}$ term can be controlled. From Lemma 2.5.6, the continuity of $\mathscr{C}$

immediately implies that

$$\begin{aligned}\|\mathscr{C}(\widehat{C}) - \mathscr{C}(C)\| &\leq \|\widehat{C} - C\| \left(1 + \|A\|^2 \|\Gamma^{-1}\| \left(\|\widehat{C}\| + \|C\|\right) + \|A\|^4 \|\Gamma^{-1}\|^2 \|\widehat{C}\| \|C\|\right) \\ &= \left[\|A\|^2 \|\Gamma^{-1}\| + \|A\|^4 \|\Gamma^{-1}\|^2 \|C\|\right] \|\widehat{C} - C\| \|\widehat{C}\| \\ &+ \left[1 + \|A\|^2 \|\Gamma^{-1}\| \|C\|\right] \|\widehat{C} - C\|.\end{aligned}$$

For any $N \in \mathbb{N}$ and $a > 0$, let $\mathscr{R}_N(a) \equiv \sqrt{\frac{a}{N}} \vee \frac{a}{N}$. Let E_1 be the set on which both

$$\|\widehat{C} - C\| \lesssim \|C\| \left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t)\right), \qquad \text{and} \qquad \|\widehat{C}\| \lesssim \|C\| \left(1 \vee \mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t)\right).$$

Let E_2 be the set on which

$$\|\hat{\Gamma} - \Gamma\| \lesssim \|\Gamma\| \left(\mathscr{R}_N(r_2(\Gamma)) \vee \mathscr{R}_N(t)\right),$$

and E_3 the set on which

$$\|\widehat{C}^{\,u\eta} - C^{\,u\eta}\| \lesssim \left(\|C\| \vee \|\Gamma\|\right) \left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(r_2(\Gamma)) \vee \mathscr{R}_N(t)\right).$$

Then, by Proposition 2.2.1 applied separately to E_1 and E_2, and Lemma 2.5.3 applied to E_3, the intersection $E = E_1 \cap E_2 \cap E_3$ has probability at least $1 - ce^{-t}$. It follows that on E:

$$\|\widehat{C}-C\|\|\widehat{C}\| \lesssim \|C\|^2 \left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t)\right)\left(1 \vee \mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t)\right)$$

$$\lesssim \|C\|^2 \left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t) \vee \mathscr{R}^2_{N,2}(C) \vee \mathscr{R}^2_{N,2}(t)\right), \tag{2.37}$$

$$\|\hat{\Gamma}-\Gamma\|\|\widehat{C}\|^2 \lesssim \|C\|^2\|\Gamma\| \left(\mathscr{R}_N(r_2(\Gamma)) \vee \mathscr{R}_N(t)\right)\left(1 \vee \mathscr{R}^2_{N,2}(C) \vee \mathscr{R}^2_{N,2}(t)\right), \tag{2.38}$$

$$\|\widehat{C}^{u\eta}-C^{u\eta}\|\|\widehat{C}\| \lesssim \|C\|(\|C\| \vee \|\Gamma\|) \tag{2.39}$$

$$\times \left(1 \vee \mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t)\right)\left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(r_2(\Gamma)) \vee \mathscr{R}_N(t)\right), \tag{2.40}$$

$$\|\widehat{C}^{u\eta}-C^{u\eta}\|\|\widehat{C}\|^2 \lesssim \|C\|^2(\|C\| \vee \|\Gamma\|) \tag{2.41}$$

$$\times \left(1 \vee \mathscr{R}^2_{N,2}(C) \vee \mathscr{R}^2_{N,2}(t)\right)\left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(r_2(\Gamma)) \vee \mathscr{R}_N(t)\right). \tag{2.42}$$

Using (2.37), it follows that on E,

$$\begin{aligned}\|\widehat{\Sigma}-\Sigma\| &\lesssim (\|C\| \vee \|C\|^3)(\|A\|^2 \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)\\ &\quad\times \left(\mathscr{R}_N(r_2(C)) \vee \mathscr{R}_N(t) \vee \mathscr{R}^2_{N,2}(C) \vee \mathscr{R}^2_{N,2}(t)\right)\\ &= (\|C\| \vee \|C\|^3)(\|A\|^2 \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)\\ &\quad\times \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^2 \vee \sqrt{\frac{t}{N}} \vee \left(\frac{t}{N}\right)^2\right)\end{aligned}$$

Next, for the PO-ensemble Kalman update, it follows by the triangle inequality that

$$\|\widehat{O}\| \leq \|\mathscr{K}(\widehat{C})(\widehat{\Gamma}-\Gamma)\mathscr{K}^\top(\widehat{C})\| \tag{2.43}$$

$$+\|(I-\mathscr{K}(\widehat{C})A)\widehat{C}^{u\eta}\mathscr{K}^\top(\widehat{C})\| \tag{2.44}$$

$$+\|\mathscr{K}(\widehat{C})(\widehat{C}^{u\eta})^\top(I-A^\top\mathscr{K}^\top(\widehat{C}))\|, \tag{2.45}$$

and so we may proceed by bounding each of the three terms (2.43), (2.44), and (2.45) separately. For (2.43), invoking first the bound on $\mathscr{K}$ from Lemma 2.5.4 as well as the inequality in (2.38), it holds on E that

$$\begin{aligned}
\|\mathscr{K}(\widehat{C})(\widehat{\Gamma}-\Gamma)\mathscr{K}^\top(\widehat{C})\| &\le \|\mathscr{K}(\widehat{C})\|^2\|\widehat{\Gamma}-\Gamma\| \\
&\le \|A\|^2\|\Gamma^{-1}\|^2\|\widehat{C}\|^2\|\widehat{\Gamma}-\Gamma\| \\
&\lesssim \|A\|^2\|\Gamma^{-1}\|^2\|C\|^2\|\Gamma\|\left(\mathscr{R}_N(r_2(\Gamma))\vee\mathscr{R}_N(t)\right)\left(1\vee\mathscr{R}^2_{N,2}(C)\vee\mathscr{R}^2_{N,2}(t)\right)
\end{aligned}$$

Both (2.44) and (2.45) are equal in operator norm, and so we consider only (2.44). We use Lemma 2.5.4 and Lemma 2.5.2, along with the inequalities (2.40) and (2.42) to show that on E,

$$\begin{aligned}
\|(I-\mathscr{K}(\widehat{C})A)\widehat{C}^{u\eta}\mathscr{K}^\top(\widehat{C})\| &\le \|\mathscr{K}(\widehat{C})\|\|I-\mathscr{K}(\widehat{C})A\|\|\widehat{C}^{u\eta}\| \\
&\le \|\mathscr{K}(\widehat{C})\|\left(1+\|\mathscr{K}(\widehat{C})\|\|A\|\right)\|\widehat{C}^{u\eta}\| \\
&\le \|A\|\|\Gamma^{-1}\|\|\widehat{C}\|\left(1+\|A\|^2\|\Gamma^{-1}\|\|\widehat{C}\|\right)\|\widehat{C}^{u\eta}\| \\
&\lesssim (\|A\|\vee\|A\|^3)(\|\Gamma^{-1}\|\vee\|\Gamma^{-1}\|^2)[\|\widehat{C}\|+\|\widehat{C}\|^2]\|\widehat{C}^{u\eta}\| \\
&\lesssim (\|A\|\vee\|A\|^3)(\|\Gamma^{-1}\|\vee\|\Gamma^{-1}\|^2)(\|C\|\vee\|\Gamma\|)(\|C\|\vee\|C\|^2) \\
&\times\left(1\vee\mathscr{R}_N(r_2(C))\vee\mathscr{R}^2_{N,2}(C)\vee\mathscr{R}_N(t)\vee\mathscr{R}^2_{N,2}(t)\right) \\
&\times\left(\mathscr{R}_N(r_2(C))\vee\mathscr{R}_N(r_2(\Gamma))\vee\mathscr{R}_N(t)\right).
\end{aligned}$$

Some algebra shows that

$$\begin{aligned}
&\left(1\vee\mathscr{R}_N(r_2(C))\vee\mathscr{R}^2_{N,2}(C)\vee\mathscr{R}_N(t)\vee\mathscr{R}^2_{N,2}(t)\right)\left(\mathscr{R}_N(r_2(C))\vee\mathscr{R}_N(r_2(\Gamma))\vee\mathscr{R}_N(t)\right) \\
&=\left(\sqrt{\frac{r_2(C)}{N}}\vee\left(\frac{r_2(C)}{N}\right)^3\vee\sqrt{\frac{t}{N}}\vee\left(\frac{t}{N}\right)^3\vee\mathscr{R}_N(r_2(\Gamma))\left(1\vee\left(\frac{r_2(C)}{N}\right)^2\vee\left(\frac{t}{N}\right)^2\right)\right),
\end{aligned}$$

and so

$$
\begin{aligned}
\|\widehat{O}\| &\lesssim (\|A\| \vee \|A\|^3)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(\|C\| \vee \|\Gamma\|)(\|C\| \vee \|C\|^2) \\
&\quad \times \left(\sqrt{\frac{r_2(C)}{N}} \vee \left(\frac{r_2(C)}{N}\right)^3 \vee \sqrt{\frac{t}{N}} \vee \left(\frac{t}{N}\right)^3 \vee \mathscr{R}_N(r_2(\Gamma)) \left(1 \vee \left(\frac{r_2(C)}{N}\right)^2 \vee \left(\frac{t}{N}\right)^2\right)\right) \square
\end{aligned}
$$

Proof of Theorem 2.2.5. The proof follows similarly to that of Theorem 2.2.3 and is therefore omitted. The constants in the statement of the result are:

$$
\begin{aligned}
c_1 &= (\|C\| \vee \|C\|^3)(\|A\|^2 \vee \|A\|^4)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2), \\
c_2 &= (\|A\| \vee \|A\|^3)(\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(\|C\| \vee \|\Gamma\|)(\|C\| \vee \|C\|^2). \qquad \square
\end{aligned}
$$

2.5.4 Multi-Step Analysis of the Square Root Ensemble Kalman Filter

Here we provide a description of the multi-step EnKF algorithm discussed in Remark 2.2.7. As described there, we focus on the square root EnKF studied in Kwiatkowski and Mandel [2015]. Given an initial ensemble $\{v_n^{(0)}\}_{n=1}^N$, the algorithm iterates the steps of the square root ensemble update (2.12) with new observations $y^{(t)}$ and with possibly varying model matrices $A^{(t)}$. We assume that the noise distribution does not change over time, though this assumption can easily be relaxed at the expense of more cumbersome notation. We summarize both the Kalman filter and the square root EnKF in Table 2.1. In this filtering set-up, $M^{(t)} \in \mathbb{R}^{d\times d}$ is the dynamics map and $A^{(t)} \in \mathbb{R}^{k\times d}$ is the observation map at time $t \geq 1$. As detailed in Sanz-Alonso et al. [2023a], such a filtering set-up leads to a sequence of inverse problems of the form (2.4), where the forward model is given by the observation map, and the prior *forecast distribution* blends the dynamics map with previous probabilistic estimates. Throughout this subsection, we write $\|\widehat{\mu} - \mu^{(t)}\|_p \equiv \left[\mathbb{E}\|\widehat{\mu}^{(t)} - \mu^{(t)}\|_2^p\right]^{1/p}$ and $\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|_p \equiv \left[\mathbb{E}\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|^p\right]^{1/p}$.

We will use two auxiliary lemmas to prove the main result of this subsection, Corol-

	Kalman filter	Square root EnKF
Input	$\{y^{(t)}, A^{(t)}, M^{(t)}\}_{t=1}^T$, Γ, $\mu^{(0)}, \Sigma^{(0)}$	$\{y^{(t)}, A^{(t)}, M^{(t)}\}_{t=1}^T$, Γ, $\{\upsilon_n^{(0)}\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu^{(0)}, \Sigma^{(0)})$
Forecast	$m^{(t)} = M^{(t)}\mu^{(t-1)}$ $C^{(t)} = M^{(t)}\Sigma^{(t-1)}(M^{(t)})^\top$	$u_n^{(t)} = M^{(t)}\upsilon_n^{(t-1)}, \quad n = 1, \dots, N$ $\widehat{m}^{(t)} = \frac{1}{N}\sum_{n=1}^N u_n^{(t)}$ $\widehat{C}^{(t)} = \frac{1}{N-1}\sum_{n=1}^N (u_n^{(t)} - \widehat{m}^{(t)})(u_n^{(t)} - \widehat{m}^{(t)})^\top$
Analysis	$\mu^{(t)} = \mathscr{M}(m^{(t)}, C^{(t)}; A^{(t)}, y^{(t)}, \Gamma)$ $\Sigma^{(t)} = \mathscr{C}(C^{(t)}; A^{(t)}, \Gamma)$	$\upsilon_n^{(t)} = \mathscr{M}(u_n^{(t)}, \widehat{C}^{(t)}; A^{(t)}, y^{(t)}, \Gamma), \quad n = 1, \dots, N$ $\widehat{\mu}^{(t)} = \frac{1}{N}\sum_{n=1}^N \upsilon_n^{(t)}$ $\widehat{\Sigma}^{(t)} = \frac{1}{N-1}\sum_{n=1}^N (\upsilon_n^{(t)} - \widehat{\mu}^{(t)})(\upsilon_n^{(t)} - \widehat{\mu}^{(t)})^\top$
Output	$\{\mu^{(t)}, \Sigma^{(t)}\}_{t=1}^T$	$\{\widehat{\mu}^{(t)}, \widehat{\Sigma}^{(t)}\}_{t=1}^T$

Table 2.1: Comparison of the Kalman filter and square root EnKF considered in Kwiatkowski and Mandel [2015]. The forecast and analysis steps are to be repeated for $t = 1, \dots, T$ iterations.

lary 2.5.12 below.

Lemma 2.5.10 (Continuity and Boundedness of Covariance-Update Operator in L^p [Kwiatkowski and Mandel, 2015, Corollary 4.8])**.** *Let $\mathscr{C}$ be the covariance-update operator defined in* (2.7)*. Let $Q \in \mathcal{S}_+^d$ be a random matrix and $P \in \mathcal{S}_+^d$ be a deterministic matrix, $\Gamma \in \mathcal{S}_{++}^k$, $A \in \mathbb{R}^{k\times d}$, $y \in \mathbb{R}^k$, and $m, m' \in \mathbb{R}^d$. Then, for any $1 \le p < \infty$, the following holds:*

$$\|\mathscr{C}(Q) - \mathscr{C}(P)\|_p \le \|Q - P\|_p(1 + \|A\|^2\|\Gamma^{-1}\|\|P\|) \\ + (\|A\|^2\|\Gamma^{-1}\| + \|A\|^4\|\Gamma^{-1}\|^2\|P\|)\|Q\|_{2p}\|Q - P\|_{2p}.$$

Lemma 2.5.11 (Continuity and Boundedness of Mean-Update Operator in L^p [Kwiatkowski and Mandel, 2015, Corollary 4.10])**.** *Let $\mathscr{M}$ be the mean-update operator defined in* (2.6)*. Let $P, Q \in \mathcal{S}_+^d$, $\Gamma \in \mathcal{S}_{++}^k$, $A \in \mathbb{R}^{k\times d}$, $y \in \mathbb{R}^k$, and $m, m' \in \mathbb{R}^d$. Assume that Q and m are* *random, and that P and m' are deterministic. The following holds:*

$$\left\|\mathscr{M}(m,Q)-\mathscr{M}(m',P)\right\|_p \le \left\|m-m'\right\|_p + \|A\|^2\|\Gamma^{-1}\|\|Q\|_{2p}\left\|m-m'\right\|_{2p} + \|Q-P\|_p\|A\|\|\Gamma^{-1}\|\big(1+\|A\|^2\|\Gamma^{-1}\|\|P\|\big)\left\|y-Am'\right\|_2 .$$

The next result shows how our one-step bounds in Theorems 2.2.3 and 2.2.5 can be extended to provide non-asymptotic bounds on the performance of the multi-step square root EnKF. The proof follows a similar argument to the proof of [Kwiatkowski and Mandel, 2015, Theorem 6.1].

Corollary 2.5.12. *Consider the square root* EnKF *defined in Table 2.1. Suppose that $N \gtrsim r_2(\Sigma^{(0)})$. Then, for any $t \ge 1$ and $p \ge 1$,*

$$\|\widehat{\mu}^{(t)}-\mu^{(t)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|,\|A^{(l)}\|,\|\Sigma^{(l-1)}\|,\|y^{(l)}-A^{(l)}m^{(l)}\|\}_{l=1}^t,\|\Gamma^{-1}\|),$$

$$\|\widehat{\Sigma}^{(t)}-\Sigma^{(t)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|,\|A^{(l)}\|,\|\Sigma^{(l-1)}\|\}_{l=1}^t,\|\Gamma^{-1}\|).$$

Proof. The proof follows by strong induction on the predicate in the statement of the theorem. To that end, the base case ($t=1$) holds by Theorems 2.2.3 and 2.2.5, which state that, for any $p \ge 1$,

$$\|\widehat{\mu}^{(1)}-\mu^{(1)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\|M^{(1)}\|,\|A^{(1)}\|,\|\Sigma^{(0)}\|,\|y^{(1)}-A^{(1)}m^{(1)}\|,\|\Gamma^{-1}\|),$$

$$\|\widehat{\Sigma}^{(1)}-\Sigma^{(1)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\|M^{(1)}\|,\|A^{(1)}\|,\|\Sigma^{(0)}\|,\|\Gamma^{-1}\|).$$

Suppose now that the claim holds for $l = 2,\dots,t-1$. Then, for $l=t$, we have by

Lemma 2.5.10

$$
\begin{aligned}
\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|_p &= \|\mathscr{C}(\widehat{C}^{(t)}) - \mathscr{C}(C^{(t)})\|_p \\
&\le \|\widehat{C}^{(t)} - C^{(t)}\|_p (1 + \|A^{(t)}\|^2 \|\Gamma^{-1}\| \|C^{(t)}\|) \\
&+ (\|A^{(t)}\|^2 \|\Gamma^{-1}\| + \|A^{(t)}\|^4 \|\Gamma^{-1}\|^2 \|C^{(t)}\|) \|\widehat{C}^{(t)}\|_{2p} \|\widehat{C}^{(t)} - C^{(t)}\|_{2p}. \quad (2.46)
\end{aligned}
$$

By the definition of $\widehat{C}^{(t)}, C^{(t)}$ together with the inductive hypothesis, it follows that, for $\mathfrak{p} \in \{p, 2p\}$,

$$
\begin{aligned}
\|\widehat{C}^{(t)} - C^{(t)}\|_{\mathfrak{p}} &= \|M^{(t)} (\widehat{\Sigma}^{(t-1)} - \Sigma^{(t-1)}) (M^{(t)})^\top \|_{\mathfrak{p}} \\
&\le \|M^{(t)}\|^2 \|\widehat{\Sigma}^{(t-1)} - \Sigma^{(t-1)}\|_{\mathfrak{p}} \\
&\lesssim_p \|M^{(t)}\|^2 \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|\}_{l=1}^{t-1}, \|\Gamma^{-1}\|) \\
&= \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|\}_{l=1}^{t}, \|\Gamma^{-1}\|).
\end{aligned}
$$

Further, we have

$$
\begin{aligned}
\|\widehat{C}^{(t)}\|_{2p} &\le \|\widehat{C}^{(t)} - C^{(t)}\|_{2p} + \|C^{(t)}\| \\
&\lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|\}_{l=1}^{t}, \|\Gamma^{-1}\|) + \|M^{(t)}\|^2 \|\Sigma^{(t-1)}\|.
\end{aligned}
$$

Plugging these two results into (2.46) gives

$$
\|\widehat{\Sigma}^{(t)} - \Sigma^{(t)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \times c(\{\|M^{(l)}\|, \|A^{(l)}\|, \|\Sigma^{(l-1)}\|\}_{l=1}^{t}, \|\Gamma^{-1}\|).
$$

Similarly, by Lemma 2.5.11 we have

$$\begin{aligned}\|\widehat{\mu}^{(t)}-\mu^{(t)}\|_p &= \|\mathscr{M}(\widehat{m}^{(t)},\widehat{C}^{(t)})-\mathscr{M}(m^{(t)},C^{(t)})\|_p \\ &\le \|\widehat{m}^{(t)}-m^{(t)}\|_p+\|A^{(t)}\|^2\|\Gamma^{-1}\|\|\widehat{C}^{(t)}\|_{2p}\|\widehat{m}^{(t)}-m^{(t)}\|_{2p} \\ &\quad+\|\widehat{C}^{(t)}-C^{(t)}\|_p\|A^{(t)}\|\|\Gamma^{-1}\|\big(1+\|A^{(t)}\|^2\|\Gamma^{-1}\|\|C^{(t)}\|\big)\|y^{(t)}-A^{(t)}m^{(t)}\|_2.\end{aligned} \tag{2.47}$$

By the definition of $\widehat{m}^{(t)}, m^{(t)}$ together with the inductive hypothesis, we have, for $\mathfrak{p}\in\{p,2p\}$,

$$\begin{aligned}\|\widehat{m}^{(t)}-m^{(t)}\|_{\mathfrak{p}} &= \|M^{(t)}(\widehat{\mu}^{(t-1)}-\mu^{(t-1)})\|_{\mathfrak{p}} \\ &\le \|M^{(t)}\|\|\widehat{\mu}^{(t-1)}-\mu^{(t-1)}\|_{\mathfrak{p}} \\ &\lesssim_p \|M^{(t)}\|\sqrt{\frac{r_2(\Sigma^{(0)})}{N}}\times c(\{\|M^{(l)}\|,\|A^{(l)}\|,\|\Sigma^{(l-1)}\|,\|y^{(l)}-A^{(l)}m^{(l)}\|\}_{l=1}^{t-1},\|\Gamma^{-1}\|) \\ &= \sqrt{\frac{r_2(\Sigma^{(0)})}{N}}\times c(\{\|M^{(l)}\|,\|A^{(l)}\|,\|\Sigma^{(l-1)}\|,\|y^{(l)}-A^{(l)}m^{(l)}\|\}_{l=1}^{t},\|\Gamma^{-1}\|).\end{aligned}$$

Plugging this bound and the one for $\|\widehat{C}^{(t)}\|_{2p}$ derived previously in the proof into (2.47) yields

$$\|\widehat{\mu}^{(t)}-\mu^{(t)}\|_p \lesssim_p \sqrt{\frac{r_2(\Sigma^{(0)})}{N}}\times c(\{\|M^{(l)}\|,\|A^{(l)}\|,\|\Sigma^{(l-1)}\|,\|y^{(l)}-A^{(l)}m^{(l)}\|\}_{l=1}^{t},\|\Gamma^{-1}\|). \quad \square$$

2.6 Proofs: Section 2.3

This appendix contains the proofs of all the theorems in Section 2.3. Results on covariance estimation are in Subsection 2.6.1 and our main results on ensemble Kalman updates are in Subsection 2.6.2.

2.6.1 Covariance Estimation

Here we establish Theorems 2.3.1 and 2.3.3. We first collect some required technical results in Subsection 2.6.1. Next we study covariance and cross-covariance estimation under soft sparsity in Subsections 2.6.1 and 2.6.1, respectively.

Background and Preliminaries

Definition 2.6.1 ([Talagrand, 2014, Definition 2.2.17]). *Given a set T, an admissible sequence of partitions of T is an increasing sequence (Δ_n) of partitions of T such that card$(\Delta_0) = 1$ and card$(\Delta_n) \leq 2^{2^n}$ for $n \geq 1$.*

The notion of an admissible sequence of partitions allows us to define the following notion of complexity of a set T, often referred to as *generic complexity.*

Definition 2.6.2 ([Talagrand, 2014, Definition 2.2.19]). *Let (T, d) be a possibly infinite metric space, and define*

$$\gamma_2(T, \mathsf{d}) = \inf \sup_{t \in T} \sum_{n \geq 0} 2^{n/2} Diam\big(\Delta_n(t)\big),$$

where $\Delta_n(t)$ denotes the unique element of the partition to which t belongs, and the infimum is taken over all admissible sequences of partitions.

The following theorem is known as the Majorizing Measure Theorem and provides upper and lower bounds for centered Gaussian processes in terms of the generic complexity.

Theorem 2.6.3 ([Talagrand, 2014, Theorem 2.4.1]). *Let X_t, $t \in T$ be a centered Gaussian process which induces a metric $\mathsf{d}_X : T \times T \to [0, \infty]$ defined by*

$$\mathsf{d}_X^2(s,t) = \mathbb{E}\left[(X_s - X_t)^2\right].$$

Then there exists an absolute constant $L > 0$ such that

$$\frac{1}{L}\gamma_2(T, \mathsf{d}_X) \leq \mathbb{E}\left[\sup_{t\in T} X_t\right] \leq L\gamma_2(T, \mathsf{d}_X).$$

We will be primarily interested in the case that $T = \mathcal{F}$ is some function class on the probability space $(\mathcal{X}, \mathcal{A}, \mathbb{P})$, and with d being the metric induced either by $\|\cdot\|_{L_2}$ or $\|\cdot\|_{\psi_2}$. We denote these spaces by $(\mathcal{F}, L_2)$ and $(\mathcal{F}, \psi_2)$ respectively throughout this section. The next result is an exponential generic chaining bound, which was introduced in [Dirksen, 2015, Corollary 5.7] and described in [Koltchinskii and Lounici, 2017, Theorem 8]. We present it as it was described in the latter reference.

Theorem 2.6.4 ([Koltchinskii and Lounici, 2017, Theorem 8])**.** *Let $(\mathcal{X}, \mathcal{A}, \mathbb{P})$ be a probability space and consider the random sample $X, X_1, \ldots, X_N \overset{i.i.d.}{\sim} \mathbb{P}$. Let $\mathcal{F}$ be a class of measurable functions on $(\mathcal{X}, \mathcal{A})$. There exists a universal constant $c > 0$ such that, for all $t \geq 1$, it holds with probability at least $1 - e^{-t}$ that*

$$\sup_{f\in\mathcal{F}} \left| \frac{1}{N}\sum_{n=1}^{N} f^2(X_n) - \mathbb{E}[f^2(X)] \right| \\ \leq c\left(\sup_{f\in\mathcal{F}} \|f\|_{\psi_2} \frac{\gamma_2(\mathcal{F}, \psi_2)}{\sqrt{N}} \vee \frac{\gamma_2^2(\mathcal{F}, \psi_2)}{N} \vee \sup_{f\in\mathcal{F}} \|f\|_{\psi_2}^2 \sqrt{\frac{t}{N}} \vee \sup_{f\in\mathcal{F}} \|f\|_{\psi_2}^2 \frac{t}{N} \right).$$

Lemma 2.6.5 (Expectation Bound from Probability Bound, [Talagrand, 2014, Lemma 2.2.3])**.** *Let $Y \geq 0$ be a random variable satisfying*

$$\mathbb{P}(Y \geq r) \leq a \exp\left(-\frac{r^2}{b^2}\right), \qquad r \geq 0,$$

for certain numbers $a \geq 2$ and $b > 0$. Then there is a universal constant c such that

$$\mathbb{E}[Y] \leq cb\sqrt{\log a}.$$

Finally, we recall the following dimension-free bound for the maxima of sub-Gaussian random variables.

Lemma 2.6.6 (Dimension-Free Sub-Gaussian Maxima, [Van Handel, 2017, Lemma 2.4])**.** *Let $X_1, \dots X_N$ be not necessarily independent sub-Gaussian random variables with*

$$\mathbb{P}(X_n > x) \leq ce^{-x^2/c\sigma_n^2}, \qquad \textit{for all } x \geq 0,\ 1 \leq n \leq N,$$

where $\sigma_n \geq 0$ is given, or alternatively $\|X_n\|_{\psi_2} \lesssim \sigma_n$. Then, for any $t \geq 1$, it holds with probability at least $1 - ce^{-ct}$ that

$$\max_{n \leq N} X_n \lesssim \sqrt{t} \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)},$$

where $\sigma_{(1)} \geq \sigma_{(2)} \geq \dots \sigma_{(N)}$ is the decreasing rearrangement of $\sigma_1, \dots, \sigma_N$. Further

$$\mathbb{E}\left[\max_{n \leq N} X_n\right] \lesssim \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)}.$$

Proof. The proof of the upper bound is based on the proof of Proposition 2.4.16 in Talagrand [2014]. By permutation invariance, we can assume without loss of generality that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N$. Then

$$\begin{aligned}
\mathbb{P}\left(\max_{n \leq N} \frac{X_n}{\sigma_n \sqrt{\log(n+1)}} \geq \sqrt{t}\right) &\leq \sum_{n=1}^{N} \mathbb{P}\Big(X_n \geq \sigma_n \sqrt{t \log(n+1)}\Big) \\
&\lesssim \sum_{n=1}^{N} \exp\left(-\frac{t}{c} \log(n+1)\right).
\end{aligned}$$

For $t \geq 2c$, the final expression in the above display is finite, and we may write

$$\sum_{n=1}^{N} \exp\left(-\frac{t}{c}\log(n+1)\right) = \sum_{n=2}^{N+1} \exp\left(-\frac{t}{c}\log(n)\right)$$
$$\leq \exp\left(-\frac{t}{c}\log(2)\right) + \int_2^\infty x^{-t/c} dx \leq ce^{-t/c}.$$

Therefore, for any $t \geq 2c$, it holds with probability at least $1 - ce^{-t/c}$ that

$$\max_{n \leq N} X_n \lesssim \sqrt{t} \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)}.$$

This implies that, for any $t \geq 1$, it holds with probability at least $1 - ce^{-(t \vee 2c)/c}$ that

$$\max_{n \leq N} X_n \lesssim (\sqrt{t} \vee \sqrt{2c}) \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)} \lesssim \sqrt{t} \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)}.$$

Since $1 - ce^{-(t \vee 2c)/c} \geq 1 - ce^{-t/c}$ it holds that, for any $t \geq 1$, with probability at least $1 - ce^{-t/c}$

$$\max_{n \leq N} X_n \lesssim \sqrt{t} \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)}.$$

It follows by Lemma 2.6.5 that

$$\mathbb{E}\left[\max_{n \leq N} \frac{X_n}{\sigma_n \sqrt{\log(n+1)}}\right] \leq c,$$

which in turn implies

$$\mathbb{E}\left[\max_{n \leq N} X_n\right] \lesssim \max_{n \leq N} \sigma_n \sqrt{\log(n+1)} = \max_{n \leq N} \sigma_{(n)} \sqrt{\log(n+1)}. \qquad \square$$

Covariance Estimation under Soft Sparsity

This subsection contains the proof of Theorem 2.3.1. We follow the approach in [Koltchinskii and Lounici, 2017, Theorem 4], but we restrict our attention to finite dimensional spaces. Our proof will rely on the following max-norm covariance estimation bound, which may be of independent interest.

Theorem 2.6.7 (Covariance Estimation with Sample Covariance —Max-Norm Bound)**.** *Let $X_1, \dots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and var$(X_1) = \Sigma^X$. Let $\widehat{\Sigma}^X = (N-1)^{-1}\sum_{n=1}^N (X_n - \mu^X)(X_n - \mu^X)^\top$. Then there exists a constant c such that, for all $t \geq 1$, it holds with probability at least $1 - ce^{-t}$ that*

$$\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} \leq c\Sigma^X_{(1)} \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(\Sigma^X)}{N} \right),$$

where

$$r_\infty(\Sigma^X) \equiv \frac{\max_j \Sigma^X_{(j)} \log(j+1)}{\Sigma^X_{(1)}}.$$

Proof. The proof of this result is based on the proof of the upper bound of Theorem 4 of Koltchinskii and Lounici [2017], in conjunction with Theorem 2.6.4. We deal with the case $\mu^X = 0$ first. To this end, let $Z_1, \dots, Z_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with zero mean and var$[Z_1] = \Sigma^X$. We denote the distribution of Z_1 by $\mathbb{P}$, and note that $\|\cdot\|_{\psi_1}$, $\|\cdot\|_{\psi_2}$, and $\|\cdot\|_{L_2}$ are defined implicitly with respect to $\mathbb{P}$. Let $\widehat{\Sigma}^0 = N^{-1}\sum_{n=1}^N Z_n Z_n^\top$. We rewrite the expectation of interest as a squared empirical process term over an appropriate class of functions. For $j \geq 1$ we denote the j-th canonical vector (the vector with 1 in the j-th index and zero otherwise) by e_j. Then, we note that

$$\begin{aligned}
\|\widehat{\Sigma}^0 - \Sigma^X\|_{\max} &= \sup_{i,j} \left\langle e_i, (\widehat{\Sigma}^0 - \Sigma^X) e_j \right\rangle \\
&= \sup_{i,j} \left[\left\langle \frac{e_i + e_j}{2}, (\widehat{\Sigma}^0 - \Sigma^X) \frac{e_i + e_j}{2} \right\rangle - \left\langle \frac{e_i - e_j}{2}, (\widehat{\Sigma}^0 - \Sigma^X) \frac{e_i - e_j}{2} \right\rangle \right] \\
&\leq 2 \sup_{u \in \mathcal{U}} \left| \left\langle (\widehat{\Sigma}^0 - \Sigma^X) u, u \right\rangle \right|,
\end{aligned}$$

where $\mathcal{U} = \left\{ u \in \mathbb{R}^d : u = \pm\frac{1}{2}(e_i \pm e_j),\ 1 \leq i, j \leq d \right\}$. Define the set of functions $\mathcal{F}_U = \left\{ \langle \cdot, u \rangle : u \in \mathcal{U} \right\}$, and note that, for any $f \in \mathcal{F}_{\mathcal{U}}$, $-f \in \mathcal{F}_{\mathcal{U}}$ and $\mathbb{E}[f(Z_1)] = 0$. It then follows by Theorem 2.6.4 that for the same universal constant c in the statement of the theorem,

$$\begin{aligned}
2 \sup_{u \in \mathcal{U}} \left| \left\langle (\widehat{\Sigma}^0 - \Sigma^X) u, u \right\rangle \right| &= 2 \sup_{u \in \mathcal{U}} \left| \frac{1}{N} \sum_{n=1}^{N} \langle Z_n, u \rangle^2 - \left\langle u, \Sigma^X u \right\rangle \right| \\
&= 2 \sup_{f \in \mathcal{F}_{\mathcal{U}}} \left| \frac{1}{N} \sum_{n=1}^{N} f^2(Z_n) - \mathbb{E}[f^2(Z_1)] \right| \\
&\leq 2c \left(\sup_{f \in \mathcal{F}_{\mathcal{U}}} \|f\|_{\psi_2} \frac{\gamma_2(\mathcal{F}_{\mathcal{U}}; \psi_2)}{\sqrt{N}} \vee \frac{\gamma_2^2(\mathcal{F}_{\mathcal{U}}; \psi_2)}{N} \vee \sup_{f \in \mathcal{F}_{\mathcal{U}}} \|f\|_{\psi_2}^2 \sqrt{\frac{t}{N}} \vee \sup_{f \in \mathcal{F}_{\mathcal{U}}} \|f\|_{\psi_2}^2 \frac{t}{N} \right).
\end{aligned}$$

Using the equivalence of the ψ_2 and L_2 norms for linear functionals, we have

$$\begin{aligned}
\sup_{f \in \mathcal{F}_{\mathcal{U}}} \|f\|_{\psi_2} \lesssim \sup_{f \in \mathcal{F}_{\mathcal{U}}} \|f\|_{L_2} &= \max_{u \in \mathcal{U}} \sqrt{\mathbb{E}\left[\langle Z_1, u \rangle^2\right]} = \max_{u \in \mathcal{U}} \sqrt{\langle u, \Sigma^X u \rangle} \\
&= \frac{1}{2} \max_{i,j} \sqrt{\langle e_i \pm e_j, \Sigma^X (e_i \pm e_j) \rangle} = \frac{1}{2} \max_{i,j} \sqrt{\langle e_i, \Sigma^X e_i \rangle + \langle e_j, \Sigma^X e_j \rangle \pm 2 \langle e_i, \Sigma^X e_j \rangle} \\
&= \frac{1}{2} \max_{i,j} \sqrt{\Sigma_{ii}^X + \Sigma_{jj}^X \pm 2\Sigma_{ij}^X} \leq \sqrt{\Sigma_{(1)}^X}.
\end{aligned}$$

To control the generic complexity $\gamma_2(\mathcal{F}_U, \psi_2)$, let $Y \sim \mathcal{N}(0, \Sigma^X)$ be a d-dimensional

Gaussian vector, with induced metric

$$\mathsf{d}_Y(u,v) = \sqrt{\mathbb{E}\big[(\langle Y,u\rangle - \langle Y,v\rangle)^2\big]} = \|\langle\cdot,u\rangle - \langle\cdot,v\rangle\|_{L_2}, \qquad u,v \in \mathcal{U}.$$

Using again the equivalence of the ψ_2 and L_2 norms for linear functionals, we have that

$$\gamma_2(\mathcal{F}_{\mathcal{U}};\psi_2) \lesssim \gamma_2(\mathcal{F}_{\mathcal{U}};L_2) = \gamma_2(\mathcal{U};\mathsf{d}_Y).$$

It follows then by Theorem 2.6.3 that

$$\begin{aligned}
\gamma_2(\mathcal{U};\mathsf{d}_Y) &\lesssim \mathbb{E}\left[\sup_{u\in\mathcal{U}} \langle Y,u\rangle\right] \\
&= \mathbb{E}\left[\max_{i,j} \left\langle Y, \pm\frac{1}{2}(e_i \pm e_j)\right\rangle\right] \\
&\le \mathbb{E}\left[\max_{j} \left|\left\langle Y, e_j\right\rangle\right|\right] \\
&\lesssim \max_j \sqrt{\Sigma^X_{(j)} \log(j+1)},
\end{aligned}$$

where the final inequality follows by Lemma 2.6.6. We have shown that with probability at least $1 - e^{-t}$

$$\begin{aligned}
\|\widehat{\Sigma}^0 - \Sigma^X\|_{\max} &\lesssim \left(\sqrt{\Sigma^X_{(1)} \max_j \frac{\Sigma^X_{(j)}\log(j+1)}{N}} \vee \max_j \frac{\Sigma^X_{(j)}\log(j+1)}{N} \vee \Sigma^X_{(1)}\sqrt{\frac{t}{N}} \vee \Sigma^X_{(1)}\frac{t}{N}\right) \\
&= \Sigma^X_{(1)}\left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \frac{r_\infty(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}\right). \qquad (2.48)
\end{aligned}$$

In the un-centered case, taking $X_n = Z_n + \mu^X$, we have $\widehat{\Sigma}^X = \widehat{\Sigma}^0 - \overline{Z}\,\overline{Z}^\top$ and it follows that

$$\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} \le \|\widehat{\Sigma}^0 - \Sigma^X\|_{\max} + \|\overline{Z}\,\overline{Z}^\top\|_{\max}.$$

By Lemma 2.6.6, with probability at least $1 - ce^{-t}$

$$\|\overline{Z}\,\overline{Z}^{\top}\|_{\max} \le \|\overline{Z}\|_{\max}^2 \le \frac{t}{N} \max_{j \le d} \Sigma_{(j)}^X \log(j+1) = t\Sigma_{(1)}^X \frac{r_\infty(\Sigma^X)}{N}. \tag{2.49}$$

Denote the set on which (2.48) occurs by E_1, and the set on which (2.49) occurs by E_2. Then the intersection $E = E_1 \cap E_2$ has probability at least $1 - ce^{-t}$, and it holds on E that

$$\begin{aligned}
\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} &\lesssim \Sigma_{(1)}^X \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \frac{r_\infty(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{t r_\infty(\Sigma^X)}{N} \right) \\
&= \Sigma_{(1)}^X \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{t r_\infty(\Sigma^X)}{N} \right). \qquad \square
\end{aligned}$$

Lemma 2.6.8. *Let $X_1, \ldots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and var$[X_1] = \Sigma^X$. Let $\widehat{\Sigma}^X = (N-1)^{-1} \sum_{n=1}^N (X_n - \mu^X)(X_n - \mu^X)^\top$. Then, for any $p \ge 1$,*

$$\left[\mathbb{E}\|\widehat{\Sigma}^X - \Sigma^X\|_{\max}^p\right]^{1/p} \lesssim_p \Sigma_{(1)}^X \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \frac{r_\infty(\Sigma^X)}{N} \right).$$

Proof. To ease notation, let $B \equiv \Sigma_{(1)}^X \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \frac{r_\infty(\Sigma^X)}{N} \right)$, then using that for positive

W, $\mathbb{E}[W^p] = p\int_0^\infty w^{p-1}\mathbb{P}(W > w)\, dw$ gives

$$
\begin{aligned}
\left[\mathbb{E}\|\widehat{\Sigma} - \Sigma\|_{\max}^p\right]^{1/p} &= p\int_0^\infty x^{p-1}\mathbb{P}(\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} > x)\, dx \\
&\le p\int_0^B x^{p-1} dx + \int_B^\infty x^{p-1}\mathbb{P}(\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} > x)\, dx \\
&\lesssim B^p + p\int_0^\infty x^{p-1}\exp\left(-\min\left(\frac{Nx^2}{(\Sigma^X_{(1)})^2}, \frac{Nx}{\Sigma^X_{(1)}}, \frac{Nx}{r_\infty(\Sigma^X)\Sigma^X_{(1)}}\right)\right) dx \\
&= B^p + p\max\left(\frac{\Gamma(p/2)}{2}\left(\frac{(\Sigma^X_{(1)})^2}{N}\right)^{p/2}, \Gamma(p)\left(\frac{\Sigma^X_{(1)}}{N}\right)^p, \Gamma(p)\left(\frac{r_\infty(\Sigma^X)\Sigma^X_{(1)}}{N}\right)^p\right),
\end{aligned}
$$

where the last line follows by direct integration. We therefore have

$$
\begin{aligned}
\left[\mathbb{E}\|\widehat{\Sigma}^X - \Sigma^X\|_{\max}^p\right]^{1/p} &\lesssim B + c(p)\max\left(\frac{\Sigma^X_{(1)}}{\sqrt{N}}, \frac{\Sigma^X_{(1)}}{N}, \frac{r_\infty(\Sigma^X)\Sigma^X_{(1)}}{N}\right) \\
&\le c(p)\Sigma^X_{(1)}\left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \frac{r_\infty(\Sigma^X)}{N}\right),
\end{aligned}
$$

where the final inequality holds due to the fact that $r_\infty(\Sigma^X) \gtrsim 1$. □

Theorem 2.6.9 (Covariance Estimation with Localized Sample Covariance —Operator-Norm Bound)**.** *Let $X_1, \ldots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and $var[X_1] = \Sigma^X$. Further, assume that $\Sigma^X \in \mathscr{U}_d(q, R_q)$ for some $q \in [0, 1)$ and $R_q > 0$. Let $\widehat{\Sigma}^X = (N-1)^{-1}\sum_{n=1}^N (X_n - \overline{X})(X_n - \overline{X})^\top$ and, for any $t \ge 1$, set*

$$
\rho_N \asymp \Sigma^X_{(1)}\left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(\Sigma^X)}{N}\right)
$$

and let $\widehat{\Sigma}^X_{\rho_N}$ be the localized sample covariance estimator. There exists a constant $c > 0$ such *that, with probability at least* $1 - ce^{-t}$, *it holds that*

$$\|\widehat{\Sigma}^X_{\rho_N} - \Sigma^X\| \lesssim R_q \rho_N^{1-q}.$$

Proof. The localized sample covariance matrix has elements

$$[\widehat{\Sigma}^X_{\rho_N}]_{ij} = \widehat{\Sigma}^X_{ij} \mathbf{1}_{|\widehat{\Sigma}^X_{ij}| \geq \rho_N}, \qquad 1 \leq i, j \leq d.$$

By Theorem 2.6.7, it holds with probability at least $1 - ce^{-t}$ that

$$\|\widehat{\Sigma}^X - \Sigma^X\|_{\max} \lesssim \rho_N.$$

The remainder of the analysis is carried out conditional on this event, following the approach taken in [Wainwright, 2019, Theorem 6.27]. Define the set of indices of the i-th row of Σ^X that exceed $\rho_N/2$ by

$$\mathcal{I}_i(\rho_N/2) \equiv \left(j \in (1, \dots, d) : \left|\Sigma^X_{ij}\right| \geq \rho_N/2\right), \qquad i = 1, \dots, d.$$

We then have

$$\begin{aligned}
\|\Sigma^X - \widehat{\Sigma}^X_{\rho_N}\| &\leq \|\Sigma^X - \widehat{\Sigma}^X_{\rho_N}\|_\infty \\
&= \max_{i=1,\dots,d} \sum_{j=1}^{d} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij} \mathbf{1}_{|\widehat{\Sigma}^X_{ij}| \geq \rho_N}\right| \\
&= \max_{i=1,\dots,d} \left(\sum_{j \in \mathcal{I}_i(\rho_N/2)} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij} \mathbf{1}_{|\widehat{\Sigma}^X_{ij}| \geq \rho_N}\right| + \sum_{j \notin \mathcal{I}_i(\rho_N/2)} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij} \mathbf{1}_{|\widehat{\Sigma}^X_{ij}| \geq \rho_N}\right| \right),
\end{aligned}$$

where $\widehat{\Sigma}^X_{ij}$ is element (i,j) of $\widehat{\Sigma}^X$. For $j \in \mathcal{I}_i(\rho_N/2)$, it holds that $|\Sigma^X_{ij}| \geq \rho_N/2$ so that

$$\begin{aligned}
\sum_{j\in\mathcal{I}_i(\rho_N/2)} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right| &\leq \sum_{j\in\mathcal{I}_i(\rho_N/2)} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij}\right| + \left|\widehat{\Sigma}^X_{ij} - \widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right| \\
&\leq \sum_{j\in\mathcal{I}_i(\rho_N/2)} \|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij}\|_{\max} + \left|\widehat{\Sigma}^X_{ij} - \widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right| \\
&\leq \sum_{j\in\mathcal{I}_i(\rho_N/2)} \left(\frac{\rho_N}{2} + \rho_N\right) \\
&= |\mathcal{I}_i(\rho_N/2)|\frac{3\rho_N}{2},
\end{aligned}$$

where we have used the fact that

$$\left|\widehat{\Sigma}^X_{ij} - \widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right| = 0 \times \mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N} + \widehat{\Sigma}^X_{ij} \times \mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\leq\rho_N} \leq \rho_N.$$

Further, since

$$R_q \geq \sum_{j=1}^{d} |\Sigma^X_{ij}|^q \geq |\mathcal{I}_i(\rho_N/2)| \left(\frac{\rho_N}{2}\right)^q,$$

it follows that $|\mathcal{I}_i(\rho_N/2)| \leq 2^q\rho_N^{-q}R_q$, and so

$$\sum_{j\in\mathcal{I}_i(\rho_N/2)} \left|\Sigma^X_{ij} - \widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right| \leq |\mathcal{I}_i(\rho_N/2)|\frac{3\rho_N}{2} \leq \frac{3}{2}2^{-q}\rho_N^{1-q}R_q.$$

For $j \notin \mathcal{I}_i(\rho_N)$, then $|\Sigma^X_{ij}| \leq \rho_N/2$ and so

$$|\widehat{\Sigma}^X_{ij}| \leq |\widehat{\Sigma}^X_{ij} - \Sigma^X_{ij}| + |\Sigma^X_{ij}| \leq \|\widehat{\Sigma}^X - \Sigma^X\|_{\max} + |\Sigma^X_{ij}| \leq \frac{\rho_N}{2} + \frac{\rho_N}{2} = \rho_N.$$

This implies that $\widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N} = 0$, and therefore for $q \in [0,1)$, since $|\Sigma^X_{ij}|/(\rho_N/2) \leq 1$, it holds that

$$\sum_{j\notin\mathcal{I}_i(\rho_N/2)}\left|\Sigma^X_{ij}-\widehat{\Sigma}^X_{ij}\mathbf{1}_{|\widehat{\Sigma}^X_{ij}|\geq\rho_N}\right|\leq\sum_{j\notin\mathcal{I}_i(\rho_N/2)}|\Sigma^X_{ij}|=\frac{\rho_N}{2}\sum_{j\notin\mathcal{I}_i(\rho_N/2)}\frac{|\Sigma^X_{ij}|}{\frac{\rho_N}{2}}$$
$$\leq\frac{\rho_N}{2}\sum_{j\notin\mathcal{I}_i(\rho_N/2)}\left(\frac{|\Sigma^X_{ij}|}{\rho_N/2}\right)^q\leq\rho_N^{1-q}R_q.$$

Combining these two results gives

$$\|\Sigma^X-\widehat{\Sigma}^X_{\rho_N}\|\leq 4\rho_N^{1-q}R_q.$$ □

Proof of Theorem 2.3.1. The result follows immediately by Theorem 2.6.9. □

Cross-Covariance Estimation under Soft Sparsity

This subsection contains the proof of Theorem 2.3.3. The presentation is parallel to that in Subsection 2.6.1. We will use a max-norm cross-covariance estimation bound, analogous to Theorem 2.6.7. The proof relies on a high probability bound for product function classes that was shown in [Mendelson, 2016, Theorem 1.13]. We present here a simplified version of that more general statement that suffices for our purposes.

Theorem 2.6.10. *Let $(\mathcal{X},\mathcal{A},\mathbb{P})$ be a probability space and consider the random sample $X,X_1,\dots,X_N\overset{i.i.d.}{\sim}\mathbb{P}$. Let $\mathcal{F},\mathcal{G}$ be two classes of measurable functions on $(\mathcal{X},\mathcal{A})$ such that $0\in\mathcal{F}$ and $0\in\mathcal{G}$. There exist positive universal constants c_1,c_2,c_3 such that, for all $t\geq 1$, it holds with probability at least $1-c_1e^{-c_2t}$ that*

$$\sup_{f\in\mathcal{F},g\in\mathcal{G}}\left|\frac{1}{N}\sum_{n=1}^N f(X_n)g(X_n)-\mathbb{E}\big[f(X)g(X)\big]\right|$$
$$\leq c_3\left[\left(\frac{t}{N}\vee\sqrt{\frac{t}{N}}\right)\left(\sup_{f\in\mathcal{F}}\|f\|_{\psi_2}\gamma_2(\mathcal{G},\psi_2)\vee\sup_{g\in\mathcal{G}}\|g\|_{\psi_2}\gamma_2(\mathcal{F},\psi_2)\right)\vee\frac{\gamma_2(\mathcal{F},\psi_2)\gamma_2(\mathcal{G},\psi_2)}{N}\right].$$

Proof. For notational brevity, throughout this proof we write $\gamma_2(\mathcal{F})$ instead of $\gamma_2(\mathcal{F}, \psi_2)$ and $\mathsf{d}_{\psi_2}(\mathcal{F})$ instead of $\sup_{f\in\mathcal{F}}\|f\|$ and similarly for the class $\mathcal{G}$. The result follows by an application of [Mendelson, 2016, Theorem 1.13] and the ensuing remark, which deals with the case $\mathcal{F} = \mathcal{G}$, but is easily extended to the general case considered here. Together they imply that, for any $u \geq 1$, it holds with probability at least $1 - 2\exp\left(-cu^2\left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})} \wedge \frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right)\right)$ that for any $f \in \mathcal{F}, g \in \mathcal{G}$

$$\left|\frac{1}{N}\sum_{n=1}^{N} f(X_n)g(X_n) - \mathbb{E}\big[f(X)g(X)\big]\right| \lesssim \frac{u^2}{N}\gamma_2(\mathcal{F})\gamma_2(\mathcal{G}) + \frac{u}{\sqrt{N}}\left(\gamma_2(\mathcal{F})\mathsf{d}_{\psi_2}(\mathcal{G}) + \gamma_2(\mathcal{G})\mathsf{d}_{\psi_2}(\mathcal{F})\right). \tag{2.50}$$

We seek to rewrite (2.50) so that all problem specific terms appear only in the upper bound. To this end, let

$$t \equiv u^2\left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})} \wedge \frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right) \implies u = \sqrt{t}\left(\frac{\mathsf{d}_{\psi_2}(\mathcal{F})}{\gamma_2(\mathcal{F})} \vee \frac{\mathsf{d}_{\psi_2}(\mathcal{G})}{\gamma_2(\mathcal{G})}\right)$$

and note that since $u \geq 1$, it must hold that $t \geq \left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})} \wedge \frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right)$. Therefore, for any $t \geq \left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})} \wedge \frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right)$, we have that with probability at least $1 - 2e^{-ct}$, the right-hand side of (2.50) becomes

$$\frac{t}{N}\left(\frac{\mathsf{d}_{\psi_2}^2(\mathcal{F})}{\gamma_2^2(\mathcal{F})} \vee \frac{\mathsf{d}_{\psi_2}^2(\mathcal{G})}{\gamma_2^2(\mathcal{G})}\right)\gamma_2(\mathcal{F})\gamma_2(\mathcal{G}) + \frac{\sqrt{t}}{\sqrt{N}}\left(\frac{\mathsf{d}_{\psi_2}(\mathcal{F})}{\gamma_2(\mathcal{F})} \vee \frac{\mathsf{d}_{\psi_2}(\mathcal{G})}{\gamma_2(\mathcal{G})}\right)\left(\gamma_2(\mathcal{F})\mathsf{d}_{\psi_2}(\mathcal{G}) + \gamma_2(\mathcal{G})\mathsf{d}_{\psi_2}(\mathcal{F})\right).$$

The above implies that, for any $t \geq 1$, it holds with probability at least

$$1 - 2\exp\left(-c\left(t \vee \left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})} \wedge \frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right)\right)\right) \geq 1 - 2e^{-ct},$$

that

$$\frac{1}{N}\left(t\vee\left(\frac{\gamma_2^2(\mathcal{F})}{\mathsf{d}_{\psi_2}^2(\mathcal{F})}\wedge\frac{\gamma_2^2(\mathcal{G})}{\mathsf{d}_{\psi_2}^2(\mathcal{G})}\right)\right)\left(\frac{\mathsf{d}_{\psi_2}^2(\mathcal{F})}{\gamma_2^2(\mathcal{F})}\vee\frac{\mathsf{d}_{\psi_2}^2(\mathcal{G})}{\gamma_2^2(\mathcal{G})}\right)\gamma_2(\mathcal{F})\gamma_2(\mathcal{G}) \tag{2.51}$$

$$+\frac{1}{\sqrt{N}}\left(\sqrt{t}\vee\left(\frac{\gamma_2(\mathcal{F})}{\mathsf{d}_{\psi_2}(\mathcal{F})}\wedge\frac{\gamma_2(\mathcal{G})}{\mathsf{d}_{\psi_2}(\mathcal{G})}\right)\right)\left(\frac{\mathsf{d}_{\psi_2}(\mathcal{F})}{\gamma_2(\mathcal{F})}\vee\frac{\mathsf{d}_{\psi_2}(\mathcal{G})}{\gamma_2(\mathcal{G})}\right)\left(\gamma_2(\mathcal{F})\mathsf{d}_{\psi_2}(\mathcal{G})+\gamma_2(\mathcal{G})\mathsf{d}_{\psi_2}(\mathcal{F})\right). \tag{2.52}$$

Straightforward calculations then show that the first of the two terms, (2.51), is bounded above by

$$\frac{t}{N}\frac{\mathsf{d}_{\psi_2}^2(\mathcal{F})\gamma_2(\mathcal{G})}{\gamma_2(\mathcal{F})}\vee\frac{t}{N}\frac{\mathsf{d}_{\psi_2}^2(\mathcal{G})\gamma_2(\mathcal{F})}{\gamma_2(\mathcal{G})}\vee\frac{\gamma_2(\mathcal{F})\gamma_2(\mathcal{G})}{N},$$

and (2.52) is similarly bounded above by

$$\sqrt{\frac{t}{N}}\frac{\mathsf{d}_{\psi_2}^2(\mathcal{F})\gamma_2(\mathcal{G})}{\gamma_2(\mathcal{F})}\vee\sqrt{\frac{t}{N}}\frac{\mathsf{d}_{\psi_2}^2(\mathcal{G})\gamma_2(\mathcal{F})}{\gamma_2(\mathcal{G})}\vee\frac{\mathsf{d}_{\psi_2}(\mathcal{G})\gamma_2(\mathcal{F})}{\sqrt{N}}\vee\frac{\mathsf{d}_{\psi_2}(\mathcal{F})\gamma_2(\mathcal{G})}{\sqrt{N}}.$$

Note then that since $0\in\mathcal{F}$,

$$\mathsf{d}_{\psi_2}(\mathcal{F})=\sup_{f\in\mathcal{F}}\|f\|_{\psi_2}\le\sup_{f_1,f_2\in\mathcal{F}}\|f_1-f_2\|_{\psi_2}=\mathrm{diam}_{\psi_2}(\mathcal{F})\le\gamma_2(\mathcal{F}),$$

where the final equality holds since $\gamma_2(\mathcal{F}) = \inf\sup_{f\in\mathcal{F}}\sum_{n=0}^{\infty}2^{n/2}\mathrm{diam}_{\psi_2}(\Delta_n(f))$, and for $n=0$, $\Delta_0=\mathcal{F}$. Similarly, $\mathsf{d}_{\psi_2}(\mathcal{G})\le\gamma_2(\mathcal{G})$, and so $\frac{\mathsf{d}_{\psi_2}^2(\mathcal{F})\gamma_2(\mathcal{G})}{\gamma_2(\mathcal{F})}\le d_{\psi_2}(\mathcal{F})\gamma_2(\mathcal{G})$ and $\frac{\mathsf{d}_{\psi_2}^2(\mathcal{G})\gamma_2(\mathcal{F})}{\gamma_2(\mathcal{G})}\le d_{\psi_2}(\mathcal{G})\gamma_2(\mathcal{F})$ which along with the fact that $t\ge 1$ completes the proof. □

Theorem 2.6.11 (Cross-Covariance Estimation —Max-Norm Bound)**.** *Let $X_1,\dots,X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1]=\mu^X$ and var$[X_1]=\Sigma^X$. Let $Y_1,\dots,Y_N$ be k-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[Y_1]=\mu^Y$ and var$[Y_1]=\Sigma^Y$. Define $\Sigma^{XY}=\mathbb{E}\big[(X-\mu^X)(Y-\mu^Y)^\top\big]$ and consider the cross-covariance* *estimator*

$$\widehat{\Sigma}^{XY} = \frac{1}{N-1}\sum_{n=1}^{N}(X_n - \overline{X})(Y_n - \overline{Y})^\top.$$

Then there exist positive universal constants c_1, c_2 *such that, for all* $t \geq 1$, *it holds with probability at least* $1 - c_1 e^{-c_2 t}$ *that*

$$\begin{aligned}&\|\widehat{\Sigma}^{XY} - \Sigma^{XY}\|_{\max} \\ &\lesssim (\Sigma^X_{(1)} \vee \Sigma^Y_{(1)}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(\Sigma^X)} \vee \sqrt{r_\infty(\Sigma^Y)} \right) \vee \sqrt{\frac{r_\infty(\Sigma^X)}{N}} \sqrt{\frac{r_\infty(\Sigma^Y)}{N}} \right).\end{aligned}$$

Proof. Assume first that $\mu^X = \mu^Y = 0$. Let $Z_1, \ldots, Z_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with zero mean and $\text{var}[Z_1] = \Sigma^X$, and similarly let $V_1, \ldots, V_N$ be k-dimensional i.i.d. sub-Gaussian random vectors with zero mean and $\text{var}[V_1] = \Sigma^Y$. Further, let $W_n \equiv [Z_n^\top, V_n^\top]^\top$ for $n = 1, \ldots, N$. We denote the distribution of W_1 by $\mathbb{P}$ and note that $\|\cdot\|_{\psi_2}$ and $\|\cdot\|_{L_2}$ are defined implicitly with respect to $\mathbb{P}$ throughout this proof. Define $\widehat{\Sigma}^0 = N^{-1}\sum_{n=1}^N Z_n V_n^\top$. Define the dilation operator: $\mathcal{H} : \mathbb{R}^{d\times k} \to \mathbb{R}^{(d+k)\times(d+k)}$ by

$$\mathcal{H}(A) = \begin{bmatrix} O & A \\ A^\top & O \end{bmatrix},$$

see for example [Tropp, 2015, Section 2.1.16], and note that $\|A\|_{\max} = \|\mathcal{H}(A)\|_{\max}$. Let $\mathcal{B}^m$ be the space of standard basis vectors in m dimensions, i.e. any $b \in \mathcal{B}^m$ is an m-dimensional vector with 1 in a single coordinate and 0 otherwise. Then, for $e_i, e_j \in \mathcal{B}^{d+k}$, we have

$$\|\widehat{\Sigma}^0 - \Sigma^{XY}\|_{\max} = \|\mathcal{H}(\widehat{\Sigma}^0) - \mathcal{H}(\Sigma^{XY})\|_{\max} = \max_{1\le i,j\le d+k} \left\langle (\mathcal{H}(\widehat{\Sigma}^0) - \mathcal{H}(\Sigma^{XY}))e_i, e_j \right\rangle$$
$$\le 2 \sup_{u\in\mathcal{U}} \left| \left\langle (\mathcal{H}(\widehat{\Sigma}^0) - \mathcal{H}(\Sigma^{XY}))u, u \right\rangle \right|,$$

where

$$\mathcal{U} \equiv \left\{ u \in \mathbb{R}^{d+k} : u = \pm\frac{1}{2}(e_i \pm e_j) \text{ and } e_i, e_j \in \mathcal{B}^{d+k} \right\}.$$

Writing $u = [u_1^\top, u_2^\top]^\top$ where $u_1 \in \mathbb{R}^d$ and $u_2 \in \mathbb{R}^k$, we have

$$\left\langle \mathcal{H}(\widehat{\Sigma}^0)u, u \right\rangle = \frac{2}{N}\sum_{n=1}^{N} \langle u_1, Z_n\rangle \langle u_2, V_n\rangle = \frac{2}{N}\sum_{n=1}^{N} f_u(W_n),$$

where $f_u(W_n) \equiv \langle \mathscr{A}_1 W_n, u_1\rangle \langle \mathscr{A}_2 W_n, u_2\rangle$ and where $\mathscr{A}_1 \equiv [I_d, O_{d\times k}] \in \mathbb{R}^{d\times(d+k)}$ and $\mathscr{A}_2 \equiv [O_{k\times d}, I_k] \in \mathbb{R}^{k\times(d+k)}$ are the relevant selection matrices so that $\mathscr{A}_1 W_n = Z_n$ and $\mathscr{A}_2 W_n = V_n$. We define the class of functions

$$\mathcal{F}_{\mathcal{U}} \equiv \left\{ f_u(\cdot) = \langle \mathscr{A}_1\cdot, u_1\rangle \langle \mathscr{A}_2\cdot, u_2\rangle : u = [u_1^\top, u_2^\top]^\top \in \mathcal{U} \right\}.$$

It is clear then that $\mathcal{F}_{\mathcal{U}} \subset \mathcal{F}_1 \cdot \mathcal{F}_2$, where

$$\mathcal{U}_1 \equiv \left\{ u_1 \in \mathbb{R}^d : u_1 = \pm\frac{1}{2}(e_i \pm e_j) \text{ and } e_i, e_j \in \mathcal{B}_d \right\}, \quad \mathcal{F}_1 \equiv \{ f(\cdot) = \langle \mathscr{A}_1\cdot, u_1\rangle : u_1 \in \mathcal{U}_1 \},$$
$$\mathcal{U}_2 \equiv \left\{ u_2 \in \mathbb{R}^k : u_2 = \pm\frac{1}{2}(e_i \pm e_j) \text{ and } e_i, e_j \in \mathcal{B}_k \right\}, \quad \mathcal{F}_2 \equiv \{ f(\cdot) = \langle \mathscr{A}_2\cdot, u_2\rangle : u_2 \in \mathcal{U}_2 \},$$

and $\mathcal{F}_1 \cdot \mathcal{F}_2 \equiv \{ f(\cdot) = f_1(\cdot) f_2(\cdot) : f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2 \}$. We can then apply the product empirical process concentration bound of Theorem 2.6.10, which implies that, with probability

$1 - c_1 e^{-c_2 t}$,

$$\sup_{u\in\mathcal{U}} \left|\left\langle (\mathcal{H}(\widehat{\Sigma}^0) - \mathcal{H}(\Sigma^{XY}))u, u \right\rangle\right| = \sup_{f_u\in\mathcal{F}_\mathcal{U}} \left| \frac{1}{N}\sum_{n=1}^{N} f_u(W_n) - \mathbb{E}\big[f_u(W_n)\big] \right|$$
$$\lesssim \left(\frac{t}{N} \vee \sqrt{\frac{t}{N}}\right) \big(\mathsf{d}_{\psi_2}(\mathcal{F}_1)\gamma_2(\mathcal{F}_2) \vee \mathsf{d}_{\psi_2}(\mathcal{F}_2)\gamma_2(\mathcal{F}_1)\big) \vee \frac{\gamma_2(\mathcal{F}_1)\gamma_2(\mathcal{F}_2)}{N}, \tag{2.53}$$

where we use the notational shorthand $\gamma_2(\mathcal{F}_1) = \gamma_2(\mathcal{F}_1, \psi_2)$ and $\mathsf{d}_{\psi_2}(\mathcal{F}_1) = \sup_{f\in\mathcal{F}_1}\|f\|_{\psi_2}$, and similarly for $\mathcal{F}_2$. Following a similar approach to the one taken in the proof of Theorem 2.6.7, it follows by the equivalence of ψ_2 and L_2 norms for linear functionals that

$$d_{\psi_2}(\mathcal{F}_1) = \sup_{f_1\in\mathcal{F}_1}\|f_1\|_{\psi_2} \lesssim \sup_{f_1\in\mathcal{F}_1}\|f_1\|_{L_2} = \max_{u_1\in\mathcal{U}_1}\sqrt{\langle u_1, \Sigma^X u_1\rangle} \le \sqrt{\Sigma^X_{(1)}},$$

and similarly that $d_{\psi_2}(\mathcal{F}_2) \le \sqrt{\Sigma^Y_{(1)}}$. Further,

$$\gamma_2(\mathcal{F}_1) = \gamma_2(\mathcal{F}_1, \psi_2) \lesssim \gamma_2(\mathcal{F}_1, L_2) = \gamma_2(\mathcal{U}_1, \mathsf{d}_X),$$

where

$$\mathsf{d}_X(u,v) = \sqrt{\mathbb{E}\left[(\langle g_X, u\rangle - \langle g_X, v\rangle)^2\right]}, \qquad g_X \sim \mathcal{N}(0, \Sigma^X).$$

By Theorem 2.6.3 and Lemma 2.6.6,

$$\gamma_2(\mathcal{U}_1, \mathsf{d}_X) \lesssim \mathbb{E}\left[\sup_{u_1\in\mathcal{U}_1}\langle g_X, u_1\rangle\right] = \mathbb{E}\left[\max_{i,j\le d}\left\langle g_X, \pm\frac{1}{2}(e_i \pm e_j)\right\rangle\right]$$
$$\le \mathbb{E}\left[\max_{i\le d}\langle g_X, e_i\rangle\right] \lesssim \max_{i\le d}\sqrt{\Sigma^X_{(i)}\log(i+1)}.$$

Similarly, $\gamma_2(\mathcal{F}_2) \lesssim \max_{j\le k}\sqrt{\Sigma^Y_{(j)}\log(j+1)}$. In summary, we have that

$$
\begin{aligned}
\|\widehat{\Sigma}^0 - \Sigma^{XY}\|_{\max} &\lesssim \left(\frac{t}{N} \vee \sqrt{\frac{t}{N}}\right)\left(\sqrt{\Sigma^X_{(1)}}\max_{j\le k}\sqrt{\Sigma^Y_{(j)}\log(j+1)} \vee \sqrt{\Sigma^Y_{(1)}}\max_{i\le d}\sqrt{\Sigma^X_{(i)}\log(i+1)}\right) \\
&\vee \frac{\max_{i\le d}\sqrt{\Sigma^X_{(i)}\log(i+1)}\max_{j\le k}\sqrt{\Sigma^Y_{(j)}\log(j+1)}}{N} \\
&\lesssim \left(\frac{t}{N} \vee \sqrt{\frac{t}{N}}\right)\left(\Sigma^X_{(1)}\sqrt{r_\infty(\Sigma^X)} \vee \Sigma^Y_{(1)}\sqrt{r_\infty(\Sigma^Y)}\right) \vee \sqrt{\frac{\Sigma^X_{(1)}r_\infty(\Sigma^X)}{N}}\sqrt{\frac{\Sigma^Y_{(1)}r_\infty(\Sigma^Y)}{N}} \\
&\lesssim (\Sigma^X_{(1)} \vee \Sigma^Y_{(1)})\left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}}\right)\left(\sqrt{r_\infty(\Sigma^X)} \vee \sqrt{r_\infty(\Sigma^Y)}\right) \vee \sqrt{\frac{r_\infty(\Sigma^X)}{N}}\sqrt{\frac{r_\infty(\Sigma^Y)}{N}}\right).
\end{aligned}
$$

In the un-centered case, take $X_n = Z_n + \mu^X$ and $Y_n = V_n + \mu^Y$ for $n = 1, \ldots, N$, then $\widehat{\Sigma}^{XY} = \widehat{\Sigma}^0 - \overline{X}\,\overline{Y}^\top$, and so

$$
\|\widehat{\Sigma}^{XY} - \Sigma^{XY}\|_{\max} \le \|\widehat{\Sigma}^0 - \Sigma^{XY}\|_{\max} + \|\overline{X}\,\overline{Y}^\top\|_{\max}.
$$

The first term is controlled by appealing to the result in the centered case. For the second term, we note that by Lemma 2.6.6

$$
\begin{aligned}
\|\overline{X}\,\overline{Y}^\top\|_{\max} \le \|\overline{X}\|_{\max}\|\overline{Y}\|_{\max} &\lesssim \frac{1}{N}\max_{i\le d}\sqrt{\Sigma^X_{(i)}\log(i+1)}\max_{j\le k}\sqrt{\Sigma^Y_{(j)}\log(j+1)} \\
&\le \sqrt{\frac{\Sigma^X_{(1)}r_\infty(\Sigma^X)}{N}}\sqrt{\frac{\Sigma^Y_{(1)}r_\infty(\Sigma^Y)}{N}}.
\end{aligned}
$$

□

Theorem 2.6.12 (Cross-Covariance Estimation with Localized Sample Cross-Covariance —Operator-Norm bound)**.** *Let $X_1, \ldots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and var$[X_1] = \Sigma^X$. Let $Y_1, \ldots, Y_N$ be k-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[Y_1] = \mu^Y$ and var$[Y_1] = \Sigma^Y$. Define $\Sigma^{XY} = \mathbb{E}\big[(X -$*

$\mu^X)(Y - \mu^Y)^\top]$ *and consider the estimator*

$$\widehat{\Sigma}^{XY} = \frac{1}{N-1}\sum_{n=1}^{N}(X_n - \overline{X})(Y_n - \overline{Y})^\top.$$

Assume that $\Sigma^{XY} \in \mathcal{U}_{d,k}(q_1, R_{q_1})$ *and* $\Sigma^{YX} \in \mathcal{U}_{k,d}(q_2, R_{q_2})$ *where* $q_1, q_2 \in [0,1)$ *and* R_{q_1}, R_{q_2} *are positive constants. For any* $t \geq 1$*, set*

$$\rho_N \asymp (\Sigma^X_{(1)} \vee \Sigma^Y_{(1)}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(\Sigma^X)} \vee \sqrt{r_\infty(\Sigma^Y)} \right) \vee \sqrt{\frac{r_\infty(\Sigma^X)}{N}} \sqrt{\frac{r_\infty(\Sigma^Y)}{N}} \right),$$

and let $\widehat{\Sigma}^{XY}_{\rho_N}$ *be the localized sample cross-covariance estimator. There exist positive universal constants* c_1, c_2 *such that, with probability at least* $1 - c_1 e^{-c_2 t}$*,*

$$\|\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY}\| \lesssim R_{q_1}\rho_N^{1-q_1} \vee R_{q_2}\rho_N^{1-q_2}.$$

Proof of Theorem 2.6.12. Let E denote the event on which $\|\widehat{\Sigma}^{XY} - \Sigma^{XY}\|_{\max} = \|\widehat{\Sigma}^{YX} - \Sigma^{YX}\|_{\max} \lesssim \rho_N$. By Theorem 2.6.11, E holds with probability at least $1 - c_1 e^{-c_2 t}$. Conditional on E, and following an analysis identical to the one in the proof of Theorem 2.6.9 with $\widehat{\Sigma}^{XY}(\widehat{\Sigma}^{YX})$ and $\Sigma^{XY}(\Sigma^{YX})$ in place of $\widehat{\Sigma}^X$ and Σ^X respectively, it follows that

$$\|\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY}\|_\infty \lesssim R_{q_1}\rho_N^{1-q_1},$$

and

$$\|\widehat{\Sigma}^{YX}_{\rho_N} - \Sigma^{YX}\|_\infty \lesssim R_{q_2}\rho_N^{1-q_2}.$$

The result then follows by noting that

$$\begin{aligned}\|\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY}\| &= \|\mathcal{H}(\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY})\| \le \|\mathcal{H}(\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY})\|_\infty \\ &= \|\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY}\|_\infty \vee \|\widehat{\Sigma}^{YX}_{\rho_N} - \Sigma^{YX}\|_\infty \lesssim R_{q_1}\rho_N^{1-q_1} \vee R_{q_2}\rho_N^{1-q_2},\end{aligned}$$

where $\mathcal{H}$ is the dilation operator defined in the proof of Theorem 2.6.11. □

Proof of Theorem 2.3.3. The proof follows immediately from Theorem 2.6.12: since $u_1, \dots, u_N$ are i.i.d. Gaussian they are sub-Gaussian. Moreover, since $\mathcal{G}$ is Lipschitz, by [Vershynin, 2018, Theorem 5.2.2], $\|\mathcal{G}(u_1) - \mathbb{E}[\mathcal{G}(u_1)]\|_{\psi_2} \le \|\mathcal{G}\|_{\text{Lip}}\|C\|^{1/2} < \infty$, and so $\mathcal{G}(u_1), \dots, \mathcal{G}(u_N)$ are i.i.d. sub-Gaussian random vectors. □

Lemma 2.6.13 (Stein's Lemma Stein [1972])**.** *Let $u \sim \mathcal{N}(m, C)$ be a d-dimensional Gaussian vector. Let $h : \mathbb{R}^d \to \mathbb{R}$ such that $\partial_j h \equiv \partial h(u)/\partial u_j$ exists almost everywhere and $\mathbb{E}[|\partial_j h(u)|] < \infty$, $j = 1, \dots, d$. Then*

$$Cov\big(u_j, h(u)\big) = \sum_{l=1}^{d} C_{jl}\mathbb{E}[\partial_l h(u)].$$

Lemma 2.6.14 (Soft-Sparsity of Cross-Covariance —Nonlinear Forward Map)**.** *Let u be a d-dimensional Gaussian random vector with $\mathbb{E}[u] = m$ and $var[u] = C \in \mathscr{U}_d(q, c)$. Consider the function $\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^k$ with coordinate functions $\mathcal{G}_1, \dots, \mathcal{G}_k$. Assume that for each $i = 1, \dots, d$ and $j = 1, \dots, k$, $\mathcal{G}_j : \mathbb{R}^d \to \mathbb{R}$ for $j = 1, \dots, k$, such that $\partial_i \mathcal{G}_j \equiv \partial \mathcal{G}_j(u)/\partial u_i$ exists almost everywhere, and $\mathbb{E}[|\partial_i \mathcal{G}_j|] < \infty$. Let $D\mathcal{G} \in \mathbb{R}^{k\times d}$ denote the Jacobian of $\mathcal{G}$, and assume that $\mathbb{E}\big[(D\mathcal{G})^\top\big] \in \mathscr{U}_{d,k}(q, a)$ for some $q \in [0, 1)$ and $a > 0$. Then,*

$$C^{up} \in \mathscr{U}_{d,k}\big(q, ac\|\mathbb{E}[D\mathcal{G}]\|^{1-q}_{\max}\|C\|^{1-q}_{\max}\big).$$

Proof. By Stein's Lemma (Lemma 2.6.13), the i-th row sum of C^{up} is given by

$$\begin{aligned}\sum_{j=1}^{k} C^{up}_{ij} = \sum_{j=1}^{k}\sum_{l=1}^{d} C_{il}\mathbb{E}\big[\partial_l \mathcal{G}_j(u)\big] &= \sum_{l=1}^{d} C_{il}\sum_{j=1}^{k}\mathbb{E}\big[\partial_l \mathcal{G}_j(u)\big] \\ &= \|\mathbb{E}[D\mathcal{G}]\|_{\max}\sum_{l=1}^{d} C_{il}\sum_{j=1}^{k}\frac{\mathbb{E}[\partial_l \mathcal{G}_j(u)]}{\|\mathbb{E}[D\mathcal{G}]\|_{\max}} \\ &\le \|\mathbb{E}[D\mathcal{G}]\|_{\max}^{1-q}\sum_{l=1}^{d} C_{il}\sum_{j=1}^{k}\mathbb{E}[\partial_l \mathcal{G}_j(u)]^q \\ &\le a\|\mathbb{E}[D\mathcal{G}]\|_{\max}^{1-q}\sum_{l=1}^{d} C_{il} \\ &\le ac\|\mathbb{E}[D\mathcal{G}]\|_{\max}^{1-q}\|C\|_{\max}^{1-q},\end{aligned}$$

where the first inequality holds since $q \in [0,1)$ and $\mathbb{E}[\partial_l \mathcal{G}_j(u)] \le \|\mathbb{E}[D\mathcal{G}]\|_{\max}$. □

Lemma 2.6.15 (Product of Two Soft-Sparse Matrices)**.** *Fix $q \in [0,1)$ and let $S \in \mathscr{U}_d(q,s)$ and assume $S^\top = S$. Let $B \in \mathscr{U}_{k,d}(q,b)$. Then $BS \in \mathscr{U}_{k,d}(q, bs\|B\|_{\max}^{1-q}\|S\|_{\max}^{1-q})$.*

Proof. The (i,j)-th element of BS is given by $[BS]_{ij} = \sum_{l=1}^{d} B_{il}S_{lj}$, and so the sum of the i-th row of BS satisfies

$$\begin{aligned}\sum_{j=1}^{d}[BS]_{ij} = \sum_{j=1}^{d}\sum_{l=1}^{d} B_{il}S_{lj} = \sum_{l=1}^{d} B_{il}\sum_{j=1}^{d} S_{lj} &= \|B\|_{\max}\|S\|_{\max}\sum_{l=1}^{d}\frac{B_{il}}{\|B\|_{\max}}\sum_{j=1}^{d}\frac{S_{lj}}{\|S\|_{\max}} \\ \le \|B\|_{\max}\|S\|_{\max}\sum_{l=1}^{d}\left(\frac{B_{il}}{\|B\|_{\max}}\right)^q\sum_{j=1}^{d}\left(\frac{S_{lj}}{\|S\|_{\max}}\right)^q &\le \|B\|_{\max}^{1-q}\|S\|_{\max}^{1-q} bs,\end{aligned}$$

where the first inequality holds since $q \in [0,1)$, and the second follows by the symmetry of S. □

Lemma 2.6.16 (Product of Three Soft-Sparse Matrices)**.** *Fix $q \in [0,1)$ and let $S \in \mathscr{U}_d(q,s)$ with $S^\top = S$. Let $B \in \mathscr{U}_{k,d}(q,b_1)$ and $B^\top \in \mathscr{U}_{d,k}(q,b_2)$, that is, B is both row and column sparse. Then $BSB^\top \in \mathscr{U}_{k,d}(q, b_1b_2s\|B\|_{\max}^{2(1-q)}\|S\|_{\max}^{1-q})$.*

Proof. The (i,j)-th element of $BSB^\top$ is given by

$$[BSB^\top]_{ij} = \sum_{m=1}^{d} [BS]_{im} B^\top_{mj} = \sum_{m=1}^{d} [BS]_{im} B_{jm} = \sum_{m=1}^{d} \left(\sum_{l=1}^{d} B_{il} S_{lm} \right) B_{jm}.$$

Therefore, the sum of the i-th row of $BSB^\top$ satisfies

$$\begin{aligned}\sum_{j=1}^{k} [BSB^\top]_{ij} &= \sum_{j=1}^{k} \sum_{m=1}^{d} \sum_{l=1}^{d} B_{il} S_{lm} B_{jm} = \sum_{m=1}^{d} \sum_{l=1}^{d} B_{il} S_{lm} \sum_{j=1}^{k} B_{jm} \\ &\leq \|B\|_{\max}^{1-q} b_2 \sum_{m=1}^{d} \sum_{l=1}^{d} B_{il} S_{lm} \leq b_1 b_2 s \|B\|_{\max}^{2(1-q)} \|S\|_{\max}^{1-q},\end{aligned}$$

where the final inequality follows by Lemma 2.6.15. □

Lemma 2.6.17 (Sample Covariance Deviation)**.** *Let $X_1, \ldots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and var$[X_1] = \Sigma^X$. Let $\widehat{\Sigma}^X_N = (N-1)^{-1} \sum_{n=1}^{N} (X_n - \mu^X)(X_n - \mu^X)^\top$. Then*

$$\begin{aligned}\widehat{\Sigma}^X_N - \widehat{\Sigma}^X_{N-1} \asymp{}& \frac{1}{N} X_N X_N^\top - \frac{1}{N} \widehat{\Sigma}^0_{N-1} - \frac{1}{N^2} X_N X_N^\top - \left(\left(\frac{N-1}{N} \right)^2 - 1 \right) \overline{X}_{N-1} \overline{X}^\top_{N-1} \\ &- \left(\frac{N-1}{N^2} \right) \left(X_N \overline{X}^\top_{N-1} + \overline{X}_{N-1} X_N^\top \right),\end{aligned}$$

where $\widehat{\Sigma}^0_N = N^{-1} \sum_{n=1}^{N} X_n X_n^\top$.

Proof. We work with the biased sample covariance estimator $\frac{1}{N} \sum_{n=1}^{n} (X_n - \overline{X}_N)(X_n - \overline{X}_N)^\top$, which is equivalent to the unbiased covariance estimator up to constants. Note then that

$$\widehat{\Sigma}^X_N \asymp \frac{1}{N} \sum_{n=1}^{N} (X_n - \overline{X}_N)(X_n - \overline{X}_N)^\top = \frac{1}{N} \sum_{n=1}^{N} X_n X_n^\top - \overline{X}_N \overline{X}^\top_N = \widehat{\Sigma}^0_N - \overline{X}_N \overline{X}^\top_N.$$

We now seek to control the difference $\widehat{\Sigma}_N^X - \widehat{\Sigma}_{N-1}^X$. To that end, note that

$$\begin{aligned}\overline{X}_N\overline{X}_N^\top &= \left(\frac{1}{N}X_N + \frac{N-1}{N}\overline{X}_{N-1}\right)\left(\frac{1}{N}X_N + \frac{N-1}{N}\overline{X}_{N-1}\right)^\top \\ &= \frac{1}{N^2}X_NX_N^\top + \left(\frac{N-1}{N}\right)^2\overline{X}_{N-1}\overline{X}_{N-1}^\top + \left(\frac{N-1}{N^2}\right)\left(X_N\overline{X}_{N-1}^\top + \overline{X}_{N-1}X_N^\top\right),\end{aligned}$$

and so

$$\begin{aligned}\overline{X}_N\overline{X}_N^\top - \overline{X}_{N-1}\overline{X}_{N-1}^\top = \frac{1}{N^2}X_NX_N^\top + \left(\left(\frac{N-1}{N}\right)^2 - 1\right)\overline{X}_{N-1}\overline{X}_{N-1}^\top \\ + \left(\frac{N-1}{N^2}\right)\left(X_N\overline{X}_{N-1}^\top + \overline{X}_{N-1}X_N^\top\right). \qquad (2.54)\end{aligned}$$

Therefore,

$$\begin{aligned}\widehat{\Sigma}_N^X - \widehat{\Sigma}_{N-1}^X &\asymp \left(\frac{1}{N}\sum_{n=1}^{N}X_nX_n^\top - \overline{X}_N\overline{X}_N^\top\right) - \left(\frac{1}{N-1}\sum_{n=1}^{N-1}X_nX_n^\top - \overline{X}_{N-1}\overline{X}_{N-1}^\top\right) \\ &= \frac{1}{N}X_NX_N^\top + \left(\left(\frac{1}{N} - \frac{1}{N-1}\right)\sum_{n=1}^{N-1}X_nX_n^\top\right) + \left(\overline{X}_{N-1}\overline{X}_{N-1}^\top - \overline{X}_N\overline{X}_N^\top\right) \\ &= \frac{1}{N}X_NX_N^\top - \frac{1}{N}\widehat{\Sigma}_{N-1}^0 - \frac{1}{N^2}X_NX_N^\top - \left(\left(\frac{N-1}{N}\right)^2 - 1\right)\overline{X}_{N-1}\overline{X}_{N-1}^\top \\ &\quad - \left(\frac{N-1}{N^2}\right)\left(X_N\overline{X}_{N-1}^\top + \overline{X}_{N-1}X_N^\top\right), \qquad (2.55)\end{aligned}$$

where the last equality follows by (2.54). □

Lemma 2.6.18 (Sample Cross-Covariance Deviation)**.** *Let $X_1, \ldots, X_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[X_1] = \mu^X$ and var$[X_1] = \Sigma^X$. Let $Y_1, \ldots, Y_N$ be k-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[Y_1] = \mu^Y$ and var$[Y_1] = \Sigma^Y$. Let*

$\widehat{\Sigma}_N^{XY} = (N-1)^{-1}\sum_{n=1}^N (X_n - \mu^X)(Y_n - \mu^Y)^\top$. *Then*

$$\widehat{\Sigma}_N^{XY} - \widehat{\Sigma}_{N-1}^{XY} \asymp \frac{1}{N} X_N Y_N^\top - \frac{1}{N}\widehat{\Sigma}_{N-1}^{0,XY} - \frac{1}{N^2} X_N Y_N^\top - \left(\left(\frac{N-1}{N}\right)^2 - 1\right) \overline{X}_{N-1}\overline{Y}_{N-1}^\top$$
$$- \left(\frac{N-1}{N^2}\right)\left(X_N \overline{Y}_{N-1}^\top + \overline{X}_{N-1} Y_N^\top\right),$$

where $\widehat{\Sigma}_{N-1}^{0,XY} = N^{-1}\sum_{n=1}^N X_n Y_n$.

Proof. The result follows using the same approach utilized in the proof of Lemma 2.6.17 and is omitted for brevity. □

Lemma 2.6.19 (Covariance Estimation with Known Particle — Operator-Norm Bound)**.** *Consider the set-up in Lemma 2.6.17 and assume additionally that* X_n *is known for some* $n \in \{1, \dots, N\}$. *Then with probability at least* $1 - ce^{-t}$

$$\|\widehat{\Sigma}_N^X - \Sigma^X\| \lesssim \frac{c(\|X_n\|_2, \|\mu^X\|_2)}{N} + \|\Sigma^X\|\left(\sqrt{\frac{r_2(\Sigma^X)}{N}} \vee \frac{r_2(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}\right).$$

Proof. By symmetry, we can assume without loss of generality that $n = N$. Let E_1 denote the event on which $\|\overline{X}_{N-1} - \mu^X\|_2 \lesssim \sqrt{\|\Sigma^X\|\frac{r_2(\Sigma^X)\vee t}{N-1}} \asymp \sqrt{\|\Sigma^X\|\frac{r_2(\Sigma^X)\vee t}{N}}$. Then by

Theorem 2.5.1, $\mathbb{P}(E_1) \geq 1 - e^{-t}$ and on E_1 it holds that

$$\begin{aligned}
\|\widehat{\Sigma}_N^X - \widehat{\Sigma}_{N-1}^X\| &\lesssim \frac{1}{N}\|X_N X_N^\top\| + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0\| + \left(\left(\frac{N-1}{N}\right)^2 - 1\right)\|\overline{X}_{N-1}\overline{X}_{N-1}^\top\| \\
&+ \frac{N-1}{N^2}\left(\|X_N\overline{X}_{N-1}^\top\| + \|\overline{X}_{N-1}X_N^\top\|\right) \\
&\lesssim \frac{1}{N}\|X_N\|_2^2 + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0\| + \frac{1}{N}\|\overline{X}_{N-1}\|_2^2 + \frac{1}{N}\|X_N\|_2\|\overline{X}_{N-1}\|_2 \\
&\leq \frac{1}{N}\|X_N\|_2^2 + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0 - \Sigma^X\| + \frac{1}{N}\|\Sigma^X\| + \frac{1}{N}\|\overline{X}_{N-1} - \mu^X\|_2^2 + \frac{1}{N}\|\mu^X\|_2^2 \\
&+ \frac{1}{N}\|X_N\|_2\|\overline{X}_{N-1} - \mu^X\|_2 + \frac{1}{N}\|X_N\|_2\|\mu^X\|_2 \\
&\lesssim \frac{1}{N}\|X_N\|_2^2 + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0 - \Sigma^X\| + \frac{1}{N}\|\Sigma^X\| + \|\Sigma^X\|\frac{r_2(\Sigma^X)\vee t}{N^2} + \frac{1}{N}\|\mu^X\|_2^2 \\
&+ \|X_N\|_2\sqrt{\|\Sigma^X\|}\frac{\sqrt{r_2(\Sigma^X)\vee t}}{N^{3/2}} + \frac{1}{N}\|X_N\|_2\|\mu^X\|_2, \qquad (2.56)
\end{aligned}$$

where the first line follows by Lemma 2.6.17. Let E_2 denote the event on which

$$\begin{aligned}
\|\widehat{\Sigma}_{N-1}^X - \Sigma^X\| &\lesssim \|\Sigma^X\|\left(\sqrt{\frac{r_2(\Sigma^X)}{N-1}} \vee \frac{r_2(\Sigma^X)}{N-1} \vee \sqrt{\frac{t}{N-1}} \vee \frac{t}{N-1}\right) \\
&\asymp \|\Sigma^X\|\left(\sqrt{\frac{r_2(\Sigma^X)}{N}} \vee \frac{r_2(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}\right).
\end{aligned}$$

Then by Proposition 2.2.1, $\mathbb{P}(E_2) \geq 1 - ce^{-t}$. It holds on $E_1 \cap E_2$ that

$$
\begin{aligned}
\|\widehat{\Sigma}_N^X - \Sigma^X\| &\leq \|\widehat{\Sigma}_N^X - \widehat{\Sigma}_{N-1}^X\| + \|\widehat{\Sigma}_{N-1}^X - \Sigma^X\| \\
&\lesssim \frac{1}{N}\|X_N\|_2^2 + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0 - \Sigma^X\| + \frac{1}{N}\|\Sigma^X\| + \|\Sigma^X\|\frac{r_2(\Sigma^X) \vee t}{N^2} + \frac{1}{N}\|\mu^X\|_2^2 \\
&+ \|X_N\|_2\sqrt{\|\Sigma^X\|}\frac{\sqrt{r_2(\Sigma^X) \vee t}}{N^{3/2}} + \frac{1}{N}\|X_N\|_2\|\mu^X\|_2 \\
&+ \|\Sigma^X\| \left(\sqrt{\frac{r_2(\Sigma^X)}{N}} \vee \frac{r_2(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \\
&\lesssim \frac{c(\|X_N\|_2, \|\mu^X\|_2)}{N} + \|\Sigma^X\| \left(\sqrt{\frac{r_2(\Sigma^X)}{N}} \vee \frac{r_2(\Sigma^X)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right),
\end{aligned}
$$

where the first inequality holds by (2.56). The result follows by noting that $\mathbb{P}(E_1 \cap E_2) \geq 1 - ce^{-t}$. □

Lemma 2.6.20 (Covariance Estimation with Known Particle — Maximum-Norm Bound)**.** *Consider the set-up in Lemma 2.6.17 and assume additionally that X_n is known for some $n \in \{1, \ldots, N\}$. Then with probability at least $1 - ce^{-t}$*

$$
\|\widehat{\Sigma}_N^X - \Sigma^X\|_{\max} \lesssim \frac{c(\|X_n\|_\infty, \|\mu^X\|_\infty)}{N} + \Sigma_{(1)}^X \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{t r_\infty(\Sigma^X)}{N} \right).
$$

Proof. As in the proof of Lemma 2.6.19, we may assume that $n = N$. Let E_1 denote the event on which $\|\overline{X}_{N-1} - \mu^X\|_\infty \lesssim \sqrt{t\Sigma_{(1)}^X \frac{r_\infty(\Sigma^X)}{N-1}} \asymp \sqrt{t\Sigma_{(1)}^X \frac{r_\infty(\Sigma^X)}{N}}$. Then by Lemma 2.6.6, $\mathbb{P}(E_1) \geq 1 - ce^{-ct}$ and on E_1, using similar calculations to those used to derive (2.56), it holds that

$$
\begin{aligned}
\|\widehat{\Sigma}_N^X - \widehat{\Sigma}_{N-1}^X\|_{\max} \lesssim \frac{1}{N}\|X_N\|_\infty^2 + \frac{1}{N}\|\widehat{\Sigma}_{N-1}^0 - \Sigma^X\|_{\max} + \frac{1}{N}\|\Sigma^X\|_{\max} + t\Sigma_{(1)}^X \frac{r_\infty(\Sigma^X)}{N^2} \\
+ \frac{1}{N}\|\mu^X\|_\infty^2 + \|X_N\|_\infty \sqrt{t\Sigma_{(1)}^X}\frac{\sqrt{r_\infty(\Sigma^X)}}{N^{3/2}} + \frac{1}{N}\|X_N\|_\infty\|\mu^X\|_\infty. \quad (2.57)
\end{aligned}
$$

Let E_2 be the event on which

$$\|\widehat{\Sigma}^X_{N-1} - \Sigma^X\|_{\max} \lesssim \Sigma^X_{(1)} \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(\Sigma^X)}{N} \right).$$

By Theorem 2.6.7, $\mathbb{P}(E_2) \geq 1 - ce^{-t}$. Finally, note that the desired result holds on $E_1 \cap E_2$ and that $\mathbb{P}(E_1 \cap E_2) \geq 1 - ce^{-t}$, which completes the proof. □

Lemma 2.6.21 (Cross-Covariance Estimation with Known Particle — Maximum-Norm Bound)**.** *Consider the set-up in Lemma 2.6.18 and assume additionally that* (X_n, Y_n) *is known for some* $n \in \{1, \ldots, N\}$*. Then with probability at least* $1 - ce^{-t}$

$$\|\widehat{\Sigma}^{XY}_N - \Sigma^{XY}\|_{\max} \lesssim \frac{c(\|X_n\|_\infty, \|\mu^X\|_\infty, \|Y_n\|_\infty, \|\mu^Y\|_\infty)}{N}$$
$$+ (\Sigma^X_{(1)} \vee \Sigma^Y_{(1)}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(\Sigma^X)} \vee \sqrt{r_\infty(\Sigma^Y)} \right) \vee \sqrt{\frac{r_\infty(\Sigma^X)}{N}} \sqrt{\frac{r_\infty(\Sigma^Y)}{N}} \right).$$

Proof. The result follows using the same approach utilized in the proof of Lemma 2.6.17 and utilizing the statements of Lemma 2.6.18 and Theorem 2.6.11. We omit the details for brevity. □

Lemma 2.6.22 (Covariance Estimation with Localized Sample Covariance and with Known Particle —Operator-Norm Bound)**.** *Consider the set-up in Theorem 2.6.9 and assume additionally that* X_n *is known for some* $n \in \{1, \ldots, N\}$*. For any* $t \geq 1$*, set*

$$\rho_N \asymp \frac{c(\|X_n\|_\infty, \|\mu^X\|_\infty)}{N} + \Sigma^X_{(1)} \left(\sqrt{\frac{r_\infty(\Sigma^X)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(\Sigma^X)}{N} \right)$$

and let $\widehat{\Sigma}^X_{\rho_N}$ *be the localized sample covariance estimator. There exists a constant* $c > 0$ *such* *that, with probability at least* $1 - ce^{-t}$*, it holds that*

$$\|\widehat{\Sigma}^X_{\rho_N} - \Sigma^X\| \lesssim R_q \rho_N^{1-q}.$$

Proof. The proof follows in identical fashion to that of Theorem 2.6.9, except that we now use the max-norm bound established in Lemma 2.6.20 in place of Theorem 2.6.7. □

Lemma 2.6.23 (Cross-Covariance Estimation with Localized Sample Covariance and with Known Particle —Operator-Norm Bound)**.** *Consider the set-up in Theorem 2.6.12 and assume additionally that* (X_n, Y_n) *is known for some* $n \in \{1, \dots, N\}$*. For any* $t \geq 1$*, set*

$$\begin{aligned}\rho_N \asymp &\frac{c(\|X_n\|_\infty, \|\mu^X\|_\infty, \|Y_n\|_\infty, \|\mu^Y\|_\infty)}{N} \\ &+ (\Sigma^X_{(1)} \vee \Sigma^Y_{(1)}) \left(\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sqrt{r_\infty(\Sigma^X)} \vee \sqrt{r_\infty(\Sigma^Y)} \right) \vee \sqrt{\frac{r_\infty(\Sigma^X)}{N}} \sqrt{\frac{r_\infty(\Sigma^Y)}{N}} \right).\end{aligned}$$

There exists positive universal constants c_1, c_2 *such that, with probability at least* $1 - c_1 e^{-c_2 t}$*,*

$$\|\widehat{\Sigma}^{XY}_{\rho_N} - \Sigma^{XY}\| \lesssim R_{q_1} \rho_N^{1-q_1} \vee R_{q_2} \rho_N^{1-q_2}.$$

Proof. The proof follows in identical fashion to that of Theorem 2.6.9, except that we now use the max-norm bound established in Lemma 2.6.21 in place of Theorem 2.6.7. □

2.6.2 Proof of Main Results in Section 2.3

Proof of Theorem 2.3.5. First, we may write

$$\begin{aligned}\|\upsilon_n - \upsilon_n^*\|_2 &= \left\| \big(y - \mathcal{G}(u_n) - \eta_n\big) \big(\mathscr{P}(\widehat{C}^{up}, \widehat{C}^{pp}) - \mathscr{P}(C^{up}, C^{pp}) \big) \right\|_2 \\ &\leq \|y - \mathcal{G}(u_n) - \eta_n\|_2 \|\mathscr{P}(\widehat{C}^{up}, \widehat{C}^{pp}) - \mathscr{P}(C^{up}, C^{pp})\|_2. \qquad (2.58)\end{aligned}$$

For the second term in (2.58), it follows by Lemma 2.5.7 that

$$\|\mathscr{P}(\widehat{C}^{up}, \widehat{C}^{pp}) - \mathscr{P}(C^{up}, C^{pp}))\|_2 \le \|\Gamma^{-1}\|\|\widehat{C}^{up} - C^{up}\| + \|\Gamma^{-1}\|^2\|C^{up}\|\|\widehat{C}^{pp} - C^{pp}\|.$$

In order to control the two deviation terms, we write $W_i \equiv [u_i^\top, \mathcal{G}^\top(u_i)]^\top$ for $1 \le i \le N$. Further, let

$$\widehat{C}^W = \frac{1}{N-1}\sum_{i=1}^N (W_i - \overline{W}_N)(W_i - \overline{W}_N)^\top, \qquad C^W = \begin{bmatrix} C & C^{up} \\ C^{pu} & C^{pp} \end{bmatrix}.$$

with $\overline{W}_N = [\widehat{m}^\top, \overline{\mathcal{G}}^\top]^\top$ and $\overline{\mathcal{G}}$ the sample mean of $\{\mathcal{G}(u_n)\}_{n=1}^N$. Since $u \sim \mathcal{N}(m, C)$ and $\mathcal{G}$ is Lipschitz, by Gaussian concentration [Vershynin, 2018, Theorem 5.2.2] it holds that $\|\mathcal{G}(u) - \mathbb{E}[\mathcal{G}(u)]\|_{\psi_2} \le \|\mathcal{G}\|_{\text{Lip}}\|C\|^{1/2}$ and we can apply Lemma 2.6.19. Letting E_1 be the event on which

$$\begin{aligned}\|\widehat{C}^{up} - C^{up}\| \vee \|\widehat{C}^{pp} - C^{pp}\| &\le \|\widehat{C}^W - C^W\| \\ &\lesssim \frac{c(\|W_n\|, \|\mathbb{E}[W_n]\|)}{N} + \|C^W\|\left(\sqrt{\frac{r_2(C^W)}{N}} \vee \frac{r_2(C^W)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}\right),\end{aligned}$$

then Lemma 2.6.19 ensures that $\mathbb{P}(E_1) \ge 1 - c_1 e^{-c_2 t}$. It follows that on the event E_1, we also have

$$\begin{aligned}\|\mathscr{P}(\widehat{C}^{up}, \widehat{C}^{pp}) - \mathscr{P}(C^{up}, C^{pp}))\|_2 &\lesssim \|\Gamma^{-1}\|(1 \vee \|C^{up}\|)\|C^W\| \\ &\quad \times \left(\sqrt{\frac{r_2(C^W)}{N}} \vee \frac{r_2(C^W)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}\right).\end{aligned}$$

The expression can be simplified by noting that since $C^W \succeq 0$, $\|C^W\| \leq \|C\| + \|C^{pp}\|$ and further since $\text{Tr}(C^W) = \text{Tr}(C) + \text{Tr}(C^{pp})$

$$\begin{aligned}
&\|C^W\| \left(\sqrt{\frac{r_2(C^W)}{N}} \vee \frac{r_2(C^W)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \\
&\lesssim (\|C\| \vee \|C^{pp}\|) \left(\sqrt{\frac{\text{Tr}(C) + \text{Tr}(C^{pp})}{N(\|C\| \vee \|C^{pp}\|)}} \vee \frac{\text{Tr}(C) + \text{Tr}(C^{pp})}{N(\|C\| \vee \|C^{pp}\|)} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right) \\
&\lesssim (\|C\| \vee \|C^{pp}\|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{r_2(C^{pp})}{N}} \vee \frac{r_2(C^{pp})}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right),
\end{aligned}$$

where the last inequality follows by similar reasoning to that used in the proof of Lemma 2.5.3.

□

Proof of Theorem 2.3.7. As in the proof of Theorem 2.3.5, we have that

$$\|\upsilon_n^\rho - \upsilon_n^*\|_2 \leq \|y - \mathcal{G}(u_n) - \eta_n\|_2 \|\mathscr{P}(\widehat{C}_{\rho_N}^{up}, \widehat{C}_{\rho_N}^{pp}) - \mathscr{P}(C^{up}, C^{pp})\|_2.$$

Further, by Lemma 2.5.7,

$$\begin{aligned}
\|\mathscr{P}(\widehat{C}_{\rho_N}^{up}, \widehat{C}_{\rho_N}^{pp}) - \mathscr{P}(C^{up}, C^{pp})\|_2 &\leq (\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|C^{up}\|) \\
&\quad \times (\|\widehat{C}_{\rho_N}^{up} - C^{up}\| + \|\widehat{C}_{\rho_N}^{pp} - C^{pp}\|).
\end{aligned}$$

Let E_1 denote the event on which

$$\|\widehat{C}_{\rho_N}^{up} - C^{up}\| \lesssim R_{q_1} \rho_{N,1}^{1-q_1} \vee R_{q_2} \rho_{N,2}^{1-q_2}.$$

By Lemma 2.6.23 , E_1 has probability at least $1 - c_1 e^{-c_2 t}$. Let E_2 be the event on which

$$\|\widehat{C}_{\rho_N}^{pp} - C^{pp}\| \lesssim R_{q_3} \rho_{N,3}^{1-q_3}.$$

By Lemma 2.6.22, E_2 has probability at least $1 - c_1 e^{-c_2 t}$. Therefore, $E = E_1 \cap E_2$ has probability at least $1 - c_1 e^{-c_2 t}$, and on E it holds that

$$\begin{aligned}
&\|\mathscr{P}(\widehat{C}^{up}_{\rho_N}, \widehat{C}^{pp}_{\rho_N}) - \mathscr{P}(C^{up}, C^{pp})\|_2 \\
&\quad \lesssim (\|\Gamma^{-1}\| \vee \|\Gamma^{-1}\|^2)(1 \vee \|C^{up}\|)(R_{q_1}\rho_{N,1}^{1-q_1} + R_{q_2}\rho_{N,2}^{1-q_2} + R_{q_3}\rho_{N,3}^{1-q_3}). \qquad \square
\end{aligned}$$

2.7 Proofs: Section 4

This appendix contains the proofs of the auxiliary results discussed in Section 2.4.

Lemma 2.7.1 (Kalman Gain Deviation with Localization)**.** *Let $u_1, \dots, u_N$ be d-dimensional i.i.d. sub-Gaussian random vectors with $\mathbb{E}[u_1] = m$ and $\mathbb{E}\big[(u_1 - m)(u_1 - m)^\top\big] = C$. Assume further that $C \in \mathscr{U}_d(q, R_q)$ for some $q \in [0, 1)$ and $R_q > 0$. For any $t \geq 1$, set*

$$\rho_N \asymp C_{(1)} \left(\sqrt{\frac{r_\infty(C)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(C)}{N} \right)$$

and let $\widehat{C}_{\rho_N}$ be the localized sample covariance estimator. There exists a positive universal constant c such that, with probability at least $1 - ce^{-t}$,

$$\|\mathscr{K}(\widehat{C}_{\rho_N}) - \mathscr{K}(\widehat{C})\| \lesssim \|A\|\|\Gamma^{-1}\| R_q (1 + \|A\|^2 \|\Gamma^{-1}\| \|C\|) \rho_N^{1-q}.$$

Proof. By Lemma 2.5.4 and Theorem 2.3.1, it follows immediately that

$$\begin{aligned}
\|\mathscr{K}(\widehat{C}_{\rho_N}) - \mathscr{K}(\widehat{C})\| &\leq \|A\|\|\Gamma^{-1}\|\|\widehat{C}_{\rho_N} - C\|(1 + \|A\|^2\|\Gamma^{-1}\|\|C\|) \\
&\lesssim \|A\|\|\Gamma^{-1}\| R_q \rho_N^{1-q}(1 + \|A\|^2\|\Gamma^{-1}\|\|C\|). \qquad \square
\end{aligned}$$

Theorem 2.7.2 (Square Root Ensemble Kalman Covariance Deviation with Localization)**.** *Consider the localized SR ensemble Kalman update given by* (2.30)*, leading to an estimate*

$\widehat{\Sigma}$ of the posterior covariance Σ defined in (2.2)*. Assume that $C \in \mathscr{U}_d(q, R_q)$ for $q \in [0, 1)$ and $R_q > 0$. For any $t \geq 1$, set*

$$\rho_N \asymp C_{(1)} \left(\sqrt{\frac{r_\infty(C)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{tr_\infty(C)}{N} \right).$$

There exists a positive universal constant c such that, with probability at least $1 - ce^{-t}$,

$$\|\widehat{\Sigma} - \Sigma\| \lesssim R_q \rho_N^{1-q} \left(1 + \|A\|^2 \|\Gamma^{-1}\| \left(2\|C\| + R_q \rho_N^{1-q}\right) + \|A\|^4 \|\Gamma^{-1}\|^2 \|C\| (\|C\| + R_q \rho_N^{1-q})\right).$$

Proof. For the localized SR update we have $\widehat{\Sigma} = \mathscr{C}(\widehat{C}_{\rho_N})$. From Lemma 2.5.6, the continuity of $\mathscr{C}$ implies that

$$\|\mathscr{C}(\widehat{C}_{\rho_N}) - \mathscr{C}(C)\| \leq \|\widehat{C}_{\rho_N} - C\| \Big(1 + \|A\|^2 \|\Gamma^{-1}\| \big(\|\widehat{C}_{\rho_N}\| + \|C\|\big) + \|A\|^4 \|\Gamma^{-1}\|^2 \|\widehat{C}_{\rho_N}\| \|C\|\Big).$$

Let E denote the event on which

$$\|\widehat{C}_{\rho_N} - C\| \lesssim R_q \rho_N^{1-q}.$$

By Theorem 2.6.12, E has probability at least $1 - ce^{-t}$. It also holds on E that

$$\|\widehat{C}_{\rho_N}\| \leq \|C\| + \|\widehat{C}_{\rho_N} - C\| \lesssim \|C\| + R_q \rho_N^{1-q}.$$

Therefore, it holds on E that

$$\|\widehat{\Sigma} - \Sigma\| \lesssim R_q \rho_N^{1-q} \left(1 + \|A\|^2 \|\Gamma^{-1}\| \left(2\|C\| + R_q \rho_N^{1-q}\right) + \|A\|^4 \|\Gamma^{-1}\|^2 \|C\| (\|C\| + R_q \rho_N^{1-q})\right).$$

□

CHAPTER 3

ENSEMBLE KALMAN FILTERS WITH RESAMPLING

This chapter is adapted from the publication listed below and is used with permission of the publisher.

O. Al-Ghattas, J. Bao, and D. Sanz-Alonso, *Ensemble Kalman filters with resampling*, *SIAM/ASA Journal on Uncertainty Quantification*, vol. 12, no. 2, pp. 411–441, 2024.

3.1 Introduction

The filtering problem of estimating a time-evolving state from partial and noisy observations arises in numerous applications, including numerical weather prediction, automatic control, robotics, signal processing, machine learning, and finance Särkkä and Svensson [2023], Crisan and Rozovskii [2011], Reich and Cotter [2015], Asch et al. [2016], Law et al. [2015], Majda and Harlim [2012], Sanz-Alonso et al. [2023b]. When the state is high dimensional and the dynamics governing its evolution are complex, the method of choice is often the ensemble Kalman filter (EnKF) Evensen [1995], Evensen and Leeuwen [1996], Evensen [2009], Houtekamer and Zhang [2016], Evensen et al. [2022]. In this filtering algorithm, a Kalman gain matrix defined via the first two moments of an ensemble of particles determines the relative importance given to the dynamics and the observations in estimating the state. The size of the ensemble controls both the accuracy and the computational cost of the algorithm. Operational implementations of EnKF give accurate state estimation with a moderate ensemble size, significantly smaller than the state dimension Houtekamer and Zhang [2016]. However, non-asymptotic theory that explains the successful performance of EnKF with moderate ensemble size is still not fully developed. An important impediment to such a theory is the presence of correlations between particles, since the Kalman gain used to update each particle depends on the entire ensemble. This chapter investigates a modification of EnKF that incorporates a resampling step to break these correlations. The new algorithm is amenable to a theoretical analysis that extends and improves upon those available for filters without resampling, while also maintaining a similar empirical performance.

3.1.1 Resampling in Filtering Algorithms

Resampling techniques are routinely employed to enhance particle filtering algorithms which assimilate observations by weighting particles according to their likelihood Del Moral [2004], Doucet et al. [2009]. For particle filters, resampling converts weighted particles into unweighted ones to alleviate weight degeneracy and achieve variance reduction at later times [Chopin and Papaspiliopoulos, 2020, Chapter 9]. In contrast, EnKF assimilates observations by using *unweighted* particles and relying on a Gaussian *ansatz* and Kalman-type formulae. EnKF avoids weight degeneracy by design, but remains vulnerable to filter divergence and ensemble collapse Harlim and Majda [2010], Kelly et al. [2015]; several works have proposed using resampling to remedy these issues.

An early discussion of resampling for EnKF can be found in Anderson and Anderson [1999], which replaces the standard Gaussian *ansatz* with a more flexible sum of Gaussian kernels. The paper Zhang and Oliver [2010] introduced bootstrap methods for identifying and alleviating the impact of spurious correlations, thereby enhancing the robustness of the Kalman gain. The work Lawson and Hansen [2004] proposed a resampling scheme to improve the performance of deterministic filters in nonlinear settings. This method involves periodically resampling the ensemble based on a “bootstrapping” approach as suggested by Anderson and Anderson [1999], which is fundamentally based on a kernel density technique taken from the particle filtering literature. Closest to our work is the paper Myrseth et al. [2013], which demonstrates that resampling the Kalman gain in the conditioning step of EnKF can help prevent the ensemble from collapsing over time, consequently enhancing ensemble stability and reliability. The numerical experiments in Myrseth et al. [2013] suggest that relative to the non-resampled setting, EnKF algorithms that employ resampling give more reliable prediction intervals with a slight trade-off in the accuracy of their point predictions.

Ensemble Kalman methods are also used for offline parameter estimation and, relatedly, as numerical solvers for inverse problems, see e.g. Gu and Oliver [2007], Aanonsen et al. [2009], Li and Reynolds [2007], Iglesias et al. [2013], Chada et al. [2021]. While not the focus of this chapter, we point out that resampling techniques have also been investigated in this context. For instance, Wu et al. [2022] removes particles that significantly deviate from the posterior distribution via a resampling procedure, thus improving the performance of standard implementations. A similar idea is also considered in Wu et al. [2019], which proposes adding an extra resampling step in each iterative cycle. This method improves the convergence of the iterative EnKF by perturbing the shrinking ensemble covariances to prevent early stopping while preserving the consistent Kalman update direction of standard implementations.

3.1.2 Our Contributions

Whereas previous work investigates resampling from a methodological viewpoint Anderson and Anderson [1999], Zhang and Oliver [2010], Lawson and Hansen [2004], Myrseth et al. [2013], the primary objective of this chapter is to demonstrate that resampling strategies provide a promising approach to the design of ensemble Kalman algorithms with non-asymptotic theoretical guarantees. We consider a simple parametric resampling scheme: at the beginning of each filtering step, members of the ensemble are independently sampled from a Gaussian distribution whose mean and covariance match those of the ensemble at the previous time-step. Thereafter, the filtering step can be carried out using any of the numerous existing EnKF variants Evensen [2009], Tippett et al. [2003]. For the resulting algorithm, which we term REnKF, we establish theoretical guarantees that extend and improve upon those available for filters without resampling.

Our theoretical guarantees hold in the linear-Gaussian setting in which we provide a detailed error analysis of the ensemble mean and covariance as estimators of the mean and covariance of the filtering distributions, given by the Kalman filter Kalman [1960]. Our theory covers both stochastic and deterministic dynamical systems; in addition, it covers both stochastic implementations based on *perturbed observations* Evensen [2003] and deterministic implementations based on *square-root* filters Tippett et al. [2003], Anderson [2001], Bishop et al. [2001]. Importantly, our error-bounds are *non-asymptotic* and *dimension-free*: they hold for any given ensemble size and are written in terms of the *effective-dimension* of the covariance of the initial distribution, and of the dynamics and observation models. The non-asymptotic and dimension-free analysis of ensemble Kalman updates has recently been considered in Al-Ghattas and Sanz-Alonso [2024b], which demonstrated rigorously the success of ensemble Kalman updates whenever the ensemble size scaled with the effective dimension of the state as opposed to its ambient dimension. Given that ensemble Kalman algorithms are often employed in problems where the state dimension is very large, our results also contribute to the theoretical understanding of why ensemble methods are able to perform well even when the ensemble size is taken to be much smaller than the state dimension. This chapter extends the results in Al-Ghattas and Sanz-Alonso [2024b] by providing new bounds over multiple assimilation cycles. Our work may also be compared to Majda and Tong [2018], which puts forward a non-asymptotic and dimension-free analysis of a multi-step EnKF that utilizes a different modification than the one used to define REnKF. Specifically, Majda and Tong [2018] employs an additional projection step that determines the effective dimension of the method.

Other multi-step analyses were limited to square-root filters with deterministic dynamics Kwiatkowski and Mandel [2015], Al-Ghattas and Sanz-Alonso [2024b] and to *asymptotic* analysis of stochastic implementations Kwiatkowski and Mandel [2015], which, while ensuring consistency of the filters, do not explain their practical success when deployed with a small ensemble size. The key reason why existing non-asymptotic analyses Al-Ghattas and Sanz-Alonso [2024b] do not extend to stochastic implementations and dynamics is that these additional sources of randomness further complicate the correlations between particles, which we break via resampling.

We numerically illustrate the theory in a linear setting and also demonstrate the successful performance of REnKF on the Lorenz 96 equations Lorenz [1996], a simplified model for atmospheric dynamics widely used to test filtering algorithms Majda and Wang [2006], Majda and Harlim [2012], Law et al. [2016a], Sanz-Alonso and Stuart [2015]. In our experiments, REnKF performs similarly to standard, non-resampled EnKF in fully and partially-observed settings. Moreover, the results are robust to the noise level in the dynamics and in the observations. Python code to reproduce all numerical experiments is publicly available at https://github.com/Jiajun-Bao/EnKF-with-Resampling.

3.1.3 Outline

The rest of the chapter is organized as follows. Section 3.2 formalizes the problem setting and provides necessary background on EnKF. Section 3.3 introduces and analyzes the new REnKF algorithm. The main result, Theorem 3.3.2, gives non-asymptotic and dimension-free error bounds. We report numerical results that confirm and complement the theory in Section 3.4. Proofs are collected in Section 3.5. We close in Section 3.6 with a discussion of our results and directions for future research.

3.1.4 Notation

For a vector $u = \big(u(1), \dots, u(d)\big)^\top$ and $q \geq 1$, $|u|_q = (\sum_{i=1}^d |u(i)|^q)^{1/q}$ and $|u| = |u|_2$. For a random variable X and $q \geq 1$, we write $\|X\|_q = (\mathbb{E}|X|^q)^{1/q}$ and $\|X\| = \|X\|_2$. $X \sim \mathcal{N}(m, C)$ denotes that X is a Gaussian random vector with mean m and covariance C, and we denote its density at a point x by $\mathcal{N}(x; m, C)$. $\mathcal{S}_+^d$ denotes the set of $d \times d$ symmetric positive-semidefinite matrices, and $\mathcal{S}^d_{++}$ denotes the set of $d \times d$ symmetric positive-definite matrices. For two $d \times d$ matrices A, B, $A \succ B$ implies $A - B \in \mathcal{S}^d_{++}$ and $A \succeq B$ implies $A - B \in \mathcal{S}^d_{+}$, and similarly for $\prec, \preceq$. For a $n \times m$ matrix $A = (A_{ij})_{i=1,j=1}^{n,m}$, the operator norm is given by $|A| = \sup_{\|v\|_2=1} |Av|_2$. $\mathbf{1}\{S\}$ denotes the indicator of the set S. The identity matrix will be denoted by I, and on occasion its dimension will be made explicit with a subscript. The $n \times m$ zero matrix will be denoted by $O_{n\times m}$.

3.2 Problem Setting and Ensemble Kalman Filters

We consider a d-dimensional unobserved *state process* $\{u^{(j)}\}_{j\geq 0}$ and a k-dimensional *observation process* $\{y^{(j)}\}_{j\geq 1}$ whose relationship over discrete time j is governed by the following hidden Markov model:

$$\text{(Initialization)} \qquad u^{(0)} \sim \mathcal{N}(\mu^{(0)}, \Sigma^{(0)}), \tag{3.1}$$

$$\text{(Dynamics)} \qquad u^{(j)} = \Psi(u^{(j-1)}) + \xi^{(j)}, \qquad \xi^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Xi), \qquad j = 1, 2, \ldots \tag{3.2}$$

$$\text{(Observation)} \qquad y^{(j)} = Hu^{(j)} + \eta^{(j)}, \qquad \eta^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Gamma), \qquad j = 1, 2, \ldots \tag{3.3}$$

We assume that the initial distribution $\mathcal{N}(\mu^{(0)}, \Sigma^{(0)})$, where $\mu^{(0)} \in \mathbb{R}^d$, $\Sigma^{(0)} \in \mathcal{S}^d_{++}$, the model dynamics map $\Psi : \mathbb{R}^d \to \mathbb{R}^d$, the observation matrix $H \in \mathbb{R}^{k\times d}$, and the dynamics and observation noise covariance matrices $\Xi \in \mathcal{S}^d_{+}, \Gamma \in \mathcal{S}^k_{++}$ are known; otherwise, they may be estimated from the observations, see e.g. Evensen et al. [2022], Chen et al. [2022, 2023]. We further assume that the random variables $u^{(0)}$, $\{\xi^{(j)}\}_{j\geq 1}$, and $\{\eta^{(j)}\}_{j\geq 1}$ are mutually independent. All methods and theory presented in this chapter extend immediately to dynamics and/or observation models that are not time homogeneous at the expense of a more cumbersome notation. Additionally, nonlinear observations can be dealt with by augmenting the state, see e.g. Anderson [2001].

For a given time index $j \in \mathbb{N}$, the *filtering* goal is to compute the *filtering distribution*

$p\big(u^{(j)}|Y^{(j)}\big)$, where $Y^{(j)} := \{y^{(1)}, \dots, y^{(j)}\}$. The filtering distribution provides a probabilistic summary of the state $u^{(j)}$ conditional on observations up to time j. Given access to the filtering distribution at the preceding time-step $j-1$, $p\big(u^{(j)}|Y^{(j)}\big)$ may be obtained by the following two-step procedure:

$$\text{(Forecast)} \qquad p\big(u^{(j)}|Y^{(j-1)}\big) = \int \mathcal{N}(u^{(j)}; \Psi(u^{(j-1)}), \Xi) p\big(u^{(j-1)}|Y^{(j-1)}\big)\, du^{(j-1)}, \quad (3.4)$$

$$\text{(Analysis)} \qquad p\big(u^{(j)}|Y^{(j)}\big) \propto \mathcal{N}(y^{(j)}; Hu^{(j)}, \Gamma) p\big(u^{(j)}|Y^{(j-1)}\big). \quad (3.5)$$

The *forecast distribution* $p\big(u^{(j)}|Y^{(j-1)}\big)$ represents our knowledge of the state at time j given observations up to time $j-1$, and its computation in (3.4) utilizes the dynamics model (3.2). In the analysis step (3.5), the new observation y_j is assimilated through an application of Bayes formula with prior given by the forecast distribution and likelihood determined by the observation model (3.3). Closed-form expressions for the filtering and forecast distributions are only available for a small class of hidden Markov models Papaspiliopoulos and Ruggiero [2014]. For problems outside this class, many algorithms have been developed to approximate the filtering distributions, or, if this is too costly, to find point estimates of the state Särkkä and Svensson [2023], Sanz-Alonso et al. [2023b].

This chapter is concerned with EnKF algorithms that belong to the larger family of Kalman methods. These methods invoke a Gaussian *ansatz* for the forecast distribution, so that Bayes formula in the analysis step can be readily applied using the conjugacy of the Gaussian forecast distribution and the Gaussian likelihood model (3.3). The distinctive feature of EnKF is that the Gaussian approximation is defined using the first two moments of an ensemble of particles. Then, in the analysis step each individual particle is updated with a Kalman gain matrix which incorporates the forecast covariance. Several stochastic and deterministic implementations for the analysis step have been proposed in the literature, see e.g. Houtekamer and Zhang [2016], Tippett et al. [2003], Evensen [2009]. In Algorithm 1, an example of a stochastic implementation of EnKF —commonly referred to as the Perturbed Observation EnKF— is provided for reference, and will be our focus for this work. At time $j = 0$, an *initial ensemble* of N particles are independently drawn from the initial distribution in (3.1). These ensemble members are then sequentially passed through forecast and analysis steps: In the forecast step, the ensemble is propagated through the system dynamics yielding the j-th *forecast ensemble*. In the analysis step, the new observation $y^{(j)}$ is assimilated by updating each ensemble member according to a Kalman-type formula, yielding the j-th *analysis ensemble*. Although the initial ensemble members are mutually independent, the dependence structure of the ensemble is highly non-trivial beginning at the analysis step at time $j = 1$. Indeed, note that the Kalman Gain $K^{(1)}$ is a nonlinear transformation of the entire forecast ensemble, and this matrix is used to update each of the ensemble members when constructing the analysis ensemble. The recursive nature of the algorithm further complicates the dependence structure of the ensemble, rendering a non-asymptotic analysis highly challenging.

The stochastic variant of EnKF in Algorithm 1 is arguably the most popular in applications Evensen [1995], Van Leeuwen [2020]. Unfortunately, as noted in Furrer and Bengtsson [2007], Al-Ghattas and Sanz-Alonso [2024b] and further discussed in Section 3.3, it is harder to analyze from a non-asymptotic viewpoint than deterministic variants of the EnKF.

The output $\widehat{\mu}^{(j)}$ of EnKF gives a point estimate of the state $u^{(j)}$ at time j. For such a state-estimation task, EnKF is very effective Law and Stuart [2012]. Additionally, the output $\widehat{\Sigma}^{(j)}$ may be used to construct confidence intervals. However, as often noted in the literature Ernst et al. [2015], Law and Stuart [2012] and further discussed in Section 3.4, caution should be exercised when using ensemble Kalman algorithms for such uncertainty quantification tasks. EnKF performance for state estimation and uncertainty quantification tasks can be assessed by the error in approximating the mean and covariance of the filtering distributions; the theory in Subsection 3.3.2 adopts such performance metrics. If the moments of the filtering

Algorithm 1 Ensemble Kalman Filter (EnKF)

1: **Input**: $\Psi, H, \Xi, \Gamma, \mu^{(0)}, \Sigma^{(0)}, N$. Sequentially acquired data $\{y^{(j)}\}_{j\geq 1}$.
2: **Initialization**: $u_n^{(0)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu^{(0)}, \Sigma^{(0)}), \quad 1 \leq n \leq N$.
3: For $j = 1, 2, \dots$ do the following forecast and analysis steps:
4: **Forecast**:

$$\begin{aligned} \widehat{u}_n^{(j)} &= \Psi(u_n^{(j-1)}) + \xi_n^{(j)}, \quad \xi_n^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Xi), \quad 1 \leq n \leq N, \\ \widehat{m}^{(j)} &= \frac{1}{N}\sum_{n=1}^{N} \widehat{u}_n^{(j)}, \qquad \widehat{C}^{(j)} = \frac{1}{N-1}\sum_{n=1}^{N} \big(\widehat{u}_n^{(j)} - \widehat{m}_j\big)\big(\widehat{u}_n^{(j)} - \widehat{m}_j\big)^\top. \end{aligned} \tag{3.6}$$

5: **Analysis**:

$$\begin{aligned} K^{(j)} &= \widehat{C}^{(j)} H^\top (H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}, \\ y_n^{(j)} &= y^{(j)} + \eta_n^{(j)}, \quad \eta_n^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Gamma), \quad 1 \leq n \leq N, \\ u_n^{(j)} &= (I - K^{(j)}H)\widehat{u}_n^{(j)} + K^{(j)} y_n^{(j)}, \quad 1 \leq n \leq N, \\ \widehat{\mu}^{(j)} &= \frac{1}{N}\sum_{n=1}^{N} u_n^{(j)}, \qquad \widehat{\Sigma}^{(j)} = \frac{1}{N-1}\sum_{n=1}^{N} (u_n^{(j)} - \widehat{\mu}^{(j)})(u_n^{(j)} - \widehat{\mu}^{(j)})^\top. \end{aligned} \tag{3.7}$$

6: **Output**: Analysis mean $\widehat{\mu}^{(j)}$ and covariance $\widehat{\Sigma}^{(j)}$ for $j = 1, 2, \dots$

distributions are not available, performance metrics such as root mean squared error and coverage of confidence intervals can be employed Law and Stuart [2012], and we do so in the numerical experiments in Section 3.4.

3.3 Ensemble Kalman Filters with Resampling

In this section, we first introduce and motivate our main algorithm, EnKF with resampling (REnKF). We then present the non-asymptotic theoretical analysis of REnKF in a linear model dynamics setting.

3.3.1 Main Algorithm

The idea underlying REnKF, which is outlined in Algorithm 2, is to employ a resampling step at each filtering cycle to break the correlations between ensemble members described in Section 3.2. We consider here a particularly simple parametric resampling scheme in which at the beginning of each filtering cycle, ensembles are independently sampled from a Gaussian distribution whose mean and covariance match those of the analysis ensemble at the previous time step. Although the resampling mechanism can be made to be more sophisticated —for example, one may consider nonparametric resampling schemes in which the empirical distribution of the ensemble is used instead— we note that such complications may be difficult to justify given the simplicity, theoretical guarantees (Subsection 3.3.2), as well as the computational scalability and empirical performance (Section 3.4) of the proposed resampling strategy. Other than the resampling step, the forecast and analysis steps of REnKF agree with those of EnKF, and consequently any of the stochastic or deterministic implementations of EnKF can be adopted. Our focus here is on the stochastic implementation of EnKF in Algorithm 1. As discussed in the next subsection —see Remarks 3.3.1 and 3.3.3— non-asymptotic theory for deterministic implementations can be obtained as a by-product of the theory that we develop.

Algorithm 2 Ensemble Kalman Filter with Resampling (REnKF)

1: **Input**: $\Psi, H, \Xi, \Gamma, \mu^{(0)}, \Sigma^{(0)}, N$. Sequentially acquired data $\{y^{(j)}\}_{j\geq 1}$.
2: **Initialization**: Set $\widehat{\mu}^{(0)} = \mu^{(0)}$ and $\widehat{\Sigma}^{(0)} = \Sigma^{(0)}$.
3: For $j = 1, 2, \ldots$ do the following resampling, forecast, and analysis steps:
4: **Resampling**:

$$u_n^{(j-1)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\widehat{\mu}^{(j-1)}, \widehat{\Sigma}^{(j-1)}), \quad 1 \leq n \leq N. \tag{3.8}$$

5: **Forecast**: Do (3.6).
6: **Analysis**: Do (3.7).
7: **Output**: Analysis mean $\widehat{\mu}^{(j)}$ and covariance $\widehat{\Sigma}^{(j)}$ for $j = 1, 2, \ldots$

Notice from Algorithm 2 that correlations between particles could alternately be broken by resampling between the forecast and analysis steps. While such an approach would be amenable to a non-asymptotic analysis akin to the one we develop, we empirically found that resampling after the forecast step significantly deteriorates the performance of the filter in nonlinear settings. A heuristic explanation is that resampling tacitly introduces a Gaussian approximation, and the filtering distribution is better approximated by a Gaussian than the forecast distribution when the dynamics are nonlinear and the observations are Gaussian.

3.3.2 Non-asymptotic Error Bounds

Here we present theoretical guarantees for REnKF in a linear dynamics setting. We introduce the setting and necessary background in Subsection 3.3.2. Then, the main result is stated and discussed in Subsection 3.3.2.

Setting and Preliminaries

We consider REnKF in the following linear version of the hidden Markov model governing the relationship between the state and observation processes:

$$\text{(Initialization)} \qquad u^{(0)} \sim \mathcal{N}(\mu^{(0)}, \Sigma^{(0)}), \tag{3.9}$$

$$\text{(Dynamics)} \qquad u^{(j)} = Au^{(j-1)} + \xi^{(j)}, \qquad \xi^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Xi), \qquad j = 1, 2, \dots \tag{3.10}$$

$$\text{(Observation)} \qquad y^{(j)} = Hu^{(j)} + \eta^{(j)}, \qquad \eta^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Gamma), \qquad j = 1, 2, \dots \tag{3.11}$$

with $u^{(0)}$ independent of the i.i.d. sequences $\{\xi^{(j)}\}$ and $\{\eta^{(j)}\}$. Thus, we assume that the dynamics map Ψ in (3.2) is linear and represented by a given matrix $A \in \mathbb{R}^{d\times d}$. In this case, it is well known that the forecast distributions $p\big(u^{(j)}|Y^{(j-1)}\big) = \mathcal{N}(u^{(j)}; m^{(j)}, C^{(j-1)})$ and the filtering distributions $p\big(u^{(j)}|Y^{(j)}\big) = \mathcal{N}(u^{(j)}; \mu^{(j)}, \Sigma^{(j)})$ are both Gaussian, and the means and covariances of these distributions are given by the Kalman filter Sanz-Alonso et al. [2023b]. We aim to derive non-asymptotic bounds between the output $\widehat{\mu}^{(j)}$ and $\widehat{\Sigma}^{(j)}$ of REnKF and the output $\mu^{(j)}$ and $\Sigma^{(j)}$ of the Kalman filter.

We follow the exposition in Kwiatkowski and Mandel [2015] and introduce three operators that are central to the theory: the *Kalman gain* operator $\mathscr{K}$, the *mean-update* operator $\mathscr{M}$,

	Kalman Filter	REnKF
Forecast Mean	$m^{(j)} = A\mu^{(j-1)}$	$\widehat{m}^{(j)} = A\bar{u}^{(j-1)} + \bar{\xi}^{(j)}$
Forecast Cov.	$C^{(j)} = A\Sigma^{(j-1)}A^\top + \Xi$	$\widehat{C}^{(j)} = AS^{(j-1)}A^\top + \hat{\Xi}^{(j)} + A\widehat{C}^{(j)}_{u\xi} + \widehat{C}^{(j)}_{\xi u}A^\top$
Analysis Mean	$\mu^{(j)} = \mathscr{M}(m^{(j)}, C^{(j)}; y^{(j)})$	$\widehat{\mu}^{(j)} = \mathscr{M}(\widehat{m}^{(j)}, \widehat{C}^{(j)}; y^{(j)}) + \mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}$
Analysis Cov.	$\Sigma^{(j)} = \mathscr{C}(C^{(j)})$	$\widehat{\Sigma}^{(j)} = \mathscr{C}(\widehat{C}^{(j)}) + \hat{O}^{(j)}$

Table 3.1: Kalman filter and REnKF updates in terms of the operators (3.12), (3.13), and (3.14).

and the *covariance-update* operator $\mathscr{C}$, defined respectively by

$$\mathscr{K} : \mathcal{S}^d_+ \to \mathbb{R}^{d\times k}, \qquad \mathscr{K}(C) = CH^\top(HCH^\top + \Gamma)^{-1}, \tag{3.12}$$

$$\mathscr{M} : \mathbb{R}^d \times \mathcal{S}^d_+ \to \mathbb{R}^d, \qquad \mathscr{M}(m, C; y) = m + \mathscr{K}(C)(y - Hm), \tag{3.13}$$

$$\mathscr{C} : \mathcal{S}^d_+ \to \mathcal{S}^d_+, \qquad \mathscr{C}(C) = \big(I - \mathscr{K}(C)H\big)C. \tag{3.14}$$

With this notation, the mean and covariance updates from time $j-1$ to time j given by the Kalman filter are summarized in Table 3.1. The table also shows the corresponding updates for REnKF, where $\bar{u}^{(j-1)}$, $\bar{\xi}^{(j)}$ and $\bar{\eta}^{(j)}$ respectively denote the sample means of $\{u_n^{(j-1)}\}_{n=1}^N$, $\{\xi_n^{(j)}\}_{n=1}^N$, and $\{\eta_n^{(j)}\}_{n=1}^N$; $S^{(j-1)}$ denotes the empirical covariance of $\{u_n^{(j-1)}\}_{n=1}^N$; and $\widehat{C}^{(j)}_{u\xi} = (\widehat{C}^{(j)}_{\xi u})^\top$ denotes the empirical cross-covariance of $\{u_n^{(j-1)}\}_{n=1}^N$ and $\{\xi_n^{(j)}\}_{n=1}^N$. Finally, following Furrer and Bengtsson [2007], Al-Ghattas and Sanz-Alonso [2024b], we refer to

$$\begin{aligned}\widehat{O}^{(j)} :=\ & \mathscr{K}(\widehat{C}^{(j)})(\widehat{\Gamma}^{(j)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(j)}) \\ & + \big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)\widehat{C}^{(j)}_{u\eta}\mathscr{K}^\top(\widehat{C}^{(j)}) + \mathscr{K}(\widehat{C}^{(j)})(\widehat{C}^{(j)}_{u\eta})^\top\big(I - H^\top\mathscr{K}^\top(\widehat{C}^{(j)})\big).\end{aligned}$$

as the *offset*, where $\widehat{\Gamma}^{(j)}$ denotes the empirical covariance of $\{\eta_n^{(j)}\}_{n=1}^N$, and $\widehat{C}^{(j)}_{u\eta}$ denotes the empirical cross-covariance of $\{u_n^{(j-1)}\}_{n=1}^N$ and $\{\eta_n^{(j)}\}_{n=1}^N$.

Remark 3.3.1 (Deterministic Implementations). *As noted earlier, our presentation and analysis will focus on the stochastic (perturbed observation) implementation of* EnKF *described in Algorithm 1, and which is used within* REnKF*, see Algorithm 2. We claim that this approach is sufficient to cover both deterministic and stochastic updates. Indeed, Al-Ghattas and Sanz-Alonso [2024b] shows that deterministic and stochastic updates at time j can be succinctly written as*

$$\begin{aligned} \widehat{\mu}^{(j)} &= \mathscr{M}(\widehat{m}^{(j)}, \widehat{C}^{(j)}) + \varphi \mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}, \\ \widehat{\Sigma}^{(j)} &= \mathscr{C}(\widehat{C}^{(j)}) + \varphi \widehat{O}^{(j)}, \end{aligned} \tag{3.15}$$

where $\varphi = 1$ for the stochastic update and $\varphi = 0$ for the deterministic update. Therefore, relative to the deterministic update, theory for the stochastic update is additionally complicated by the need to consider the term $\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}$ in the mean update and the offset term $\widehat{O}^{(j)}$ in the covariance update. Accordingly, we are able to provide a result for the resampled version of the deterministic (square-root) EnKF *as a by-product of our more general theory, and we refer to Remark 3.3.3 for further discussion.*

Main Result

We define the *effective dimension* Tropp [2015] of a matrix $Q \in \mathcal{S}^d_+$ by

$$r_2(Q) := \frac{\text{Tr}(Q)}{|Q|}, \tag{3.16}$$

where $\text{Tr}(Q)$ and $|Q|$ denote the trace and operator norm of Q. The effective dimension quantifies the number of directions where Q has significant spectral content and may be significantly smaller than the ambient dimension d when the eigenvalues of Q decay quickly. As such, it is a more refined measure of complexity in high-dimensional problems with underlying low-dimensional structure. The monographs Tropp [2015], Vershynin [2018] refer to $r_2(Q)$ as the intrinsic dimension, while Koltchinskii and Lounici [2017] uses the term effective rank. This terminology is motivated by the observation that $1 \leq r_2(Q) \leq \text{rank}(Q) \leq d$ and that $r_2(Q)$ is insensitive to changes in the scale of Q, see Tropp [2015]. We now state our main result, Theorem 3.3.2, which provides non-asymptotic bounds on the deviation of REnKF from the Kalman filter for any time j.

Theorem 3.3.2. *Consider* REnKF, *Algorithm 2, with linear dynamics* $\Psi(\cdot) = A\cdot$. *Suppose that* $N \geq r_2(\Sigma^{(0)}) \vee r_2(\Gamma) \vee r_2(\Xi)$. *For any* $j = 1, 2, \dots,$ *and* $q \geq 1$

$$|||\widehat{\mu}^{(j)} - \mu^{(j)}|||_q \leq c_1 \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \right), \tag{3.17}$$

$$|||\widehat{\Sigma}^{(j)} - \Sigma^{(j)}|||_q \leq c_2 \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \right), \tag{3.18}$$

where $\mu^{(j)}$ *and* $\Sigma^{(j)}$ *are the mean and covariance of the filtering distributions, and* c_1, c_2 *are potentially different universal constants depending on*

$$|\Sigma^{(0)}|, |A|, |H|, |\Gamma^{-1}|, |\Gamma|, |\Xi|, q, j,$$

and c_1 *additionally depends on* $\{|y^{(\ell)} - Hm^{(\ell)}|\}_{\ell \leq j}$.

With the exception of [Majda and Tong, 2018, Theorem 3.4], which relies on covariance inflation and an additional projection step, Theorem 3.3.2 seems to be the first result in the literature that provides non-asymptotic guarantees on the performance of a stochastic EnKF over multiple assimilation cycles. We note that the assumption $N \geq r_2(\Sigma^{(0)}) \vee r_2(\Gamma) \vee r_2(\Xi)$ is merely for convenience and can be removed at the expense of a more cumbersome statement of the result. Importantly, the bounds (3.17) and (3.18) are non-asymptotic, in that they hold for a fixed ensemble size N. Further, the bounds are dimension-free as they do not exhibit any dependence on the state-space dimension d, implying that the ensemble need not scale with d in order for the algorithm to perform well, as has been observed empirically in the literature and confirmed in our numerical results in Section 3.4. Finally, similar to previous accuracy analyses for square-root ensemble Kalman filters Mandel et al. [2011], Al-Ghattas and Sanz-Alonso [2024b], variational data assimilation algorithms Sanz-Alonso and Stuart [2015], Law et al. [2016a], and particle filters [Sanz-Alonso et al., 2023b, Chapters 11 and 12], our proof relies on induction over the discrete time index j and does not account for potential dissipation of errors due to filter ergodicity. As a result, the constants c_1 and c_2 grow with j and our bounds (3.17) and (3.18) do not hold uniformly in time without, for instance, stability requirements on A.

Remark 3.3.3 (Resampled Square-Root Filter)**.** *While the result in Theorem 3.3.2 is specific to the stochastic* REnKF *in Algorithm 2, using the observation made in Remark 3.3.1 it is possible to show that for a deterministic variant, namely the square-root* REnKF*, and under the same assumptions on the ensemble size made in Theorem 3.3.2, we have that*

$$\begin{aligned} \| |\widehat{\mu}^{(j)} - \mu^{(j)}| \|_q &\le c_1 \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \right), \\ \| |\widehat{\Sigma}^{(j)} - \Sigma^{(j)}| \|_q &\le c_2 \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \right), \end{aligned} \tag{3.19}$$

where c_1, c_2 are potentially different universal constants depending on $|\Sigma^{(0)}|$, $|A|$, $|H|$, $|\Gamma^{-1}|$, $|\Xi|$, q, j, and c_1 additionally depends on $\{|y^{(\ell)} - Hm^{(\ell)}|\}_{\ell \le j}$. In contrast to (3.17) *and* (3.18)*, the bounds in* (3.19) *do not depend on the effective dimension of the noise covariance, $r_2(\Gamma)$, nor do the associated constants depend on $|\Gamma|$. The statistical price to pay for utilizing stochastic rather than deterministic updates is captured by these terms. We further note that [Al-Ghattas and Sanz-Alonso, 2024b, Corollary A.12], gives a non-asymptotic and multi-step analysis of a simplified version of the square-root filter (without resampling) with deterministic dynamics (that is, $\Xi = O_{d\times d}$). In such a setting, [Al-Ghattas and Sanz-Alonso,* *2024b, Corollary A.12] implies the following bounds*

$$|||\widehat{\mu}^{(j)} - \mu^{(j)}|||_q \le c_3\sqrt{\frac{r_2(\Sigma^{(0)})}{N}}, \qquad |||\widehat{\Sigma}^{(j)} - \Sigma^{(j)}|||_q \le c_4\sqrt{\frac{r_2(\Sigma^{(0)})}{N}},$$

where c_3, c_4 *are potentially different universal constants depending on* $|\Sigma^{(0)}|$, $|A|$, $|H|$, $|\Gamma^{-1}|$, q, j, *and* c_3 *additionally depends on* $\{|y^{(\ell)} - Hm^{(\ell)}|\}_{\ell \le j}$. *Theorem 3.3.2 should further be compared to [Kwiatkowski and Mandel, 2015, Theorem 6.1], which is also limited to the case* $\Xi = O_{d\times d}$ *and shows that* $|||\widehat{\mu}^{(j)} - \mu^{(j)}|||_q \le c_3' N^{-1/2}$ *and* $|||\widehat{\Sigma}^{(j)} - \Sigma^{(j)}|||_q \le c_4' N^{-1/2}$, *where* c_3', c_4' *are universal constants with the same dependencies as* c_3 *and* c_4. *Importantly, the bounds in [Kwiatkowski and Mandel, 2015, Theorem 6.1] do not capture the dependence of the algorithm on the prior covariance and also cannot be easily extended to handle stochastic dynamics* $\Xi \succ 0$ *as accomplished in Theorem 3.3.2.*

3.4 Numerical Results

In this section, we investigate the empirical performance of REnKF (Algorithm 2) and provide detailed comparisons to the stochastic EnKF (Algorithm 1). In Subsection 3.4.1, we consider a linear dynamics map, $\Psi(\cdot) = A\cdot$, with the primary goal of demonstrating the bounds of Theorem 3.3.2 in simulated settings. In Subsection 3.4.2, we study a nonlinear setting where Ψ represents the Δt-flow of the Lorenz 96 system, and Δt is the (constant) time-span between observations. The aim of this subsection is to show that REnKF achieves comparable performance to EnKF even in challenging nonlinear regimes, further motivating the study of resampling in the context of ensemble algorithms. In both Subsections 3.4.1 and 3.4.2, we examine the performance of REnKF and EnKF under varying noise levels, ensemble sizes, and state dimensions. Additionally, in Subsection 3.4.2, we consider cases in which we have access to either fully observed or partially observed dynamics. These scenarios offer a comprehensive perspective on the adaptability of REnKF to varying observational conditions, thereby highlighting its potential for wide applicability in real-world situations where data are often limited or incomplete. For all experiments, we generate a ground-truth state process $\{u^{(j)}\}_{j=0}^{J}$ for a time-window of length $J = 200$ using the initialization (3.1) and dynamics model (3.2). For each set of system parameters we examine, a unique set of observations $\{y^{(j)}\}_{j=1}^{J}$ is generated from the ground-truth state process utilizing the observation model (3.3). Python code to reproduce all numerical experiments is publicly available at https://github.com/Jiajun-Bao/EnKF-with-Resampling.

3.4.1 Linear Dynamics

In this subsection, we numerically investigate the performance of REnKF for the linear-Gaussian hidden Markov model (3.9)-(3.11) analyzed in Subsection 3.3.2. We will consider a variety of choices for the initial distribution, the dynamics noise covariance, and the observation noise covariance. Throughout, we take identity dynamics $A = I_d$ and full observations $H = I_d$. To compare the performance of EnKF and REnKF, we will consider the following metrics:

$$\text{(Mean Error)} \qquad \mathsf{E}_{\text{Linear}} = \frac{1}{J}\sum_{j=1}^{J} |\widehat{\mu}^{(j)} - \mu^{(j)}|_2, \tag{3.20}$$

$$\text{(CI Width)} \qquad \mathsf{W} = \frac{1}{J}\sum_{j=1}^{J}\frac{1}{d}\sum_{i=1}^{d} 2\times 1.96\sqrt{\widehat{\Sigma}_{ii}^{(j)}}\,, \tag{3.21}$$

$$\text{(CI Coverage)} \qquad \mathsf{V} = \frac{1}{J}\sum_{j=1}^{J}\frac{1}{d}\sum_{i=1}^{d} \mathbf{1}\big\{u^{(j)}(i) \in (\widehat{\mu}^{(j)}(i) \pm 1.96\sqrt{\widehat{\Sigma}_{ii}^{(j)}})\big\}. \tag{3.22}$$

The mean error (3.20) quantifies the approximation of the EnKF/REnKF analysis mean to the mean $\mu^{(j)}$ of the Kalman filter. Our theory for REnKF provides non-asymptotic bounds for this error, and our numerical results will show that this error is similar to that of EnKF in a variety of settings. The confidence interval (CI) width and coverage in (3.21)-(3.22)

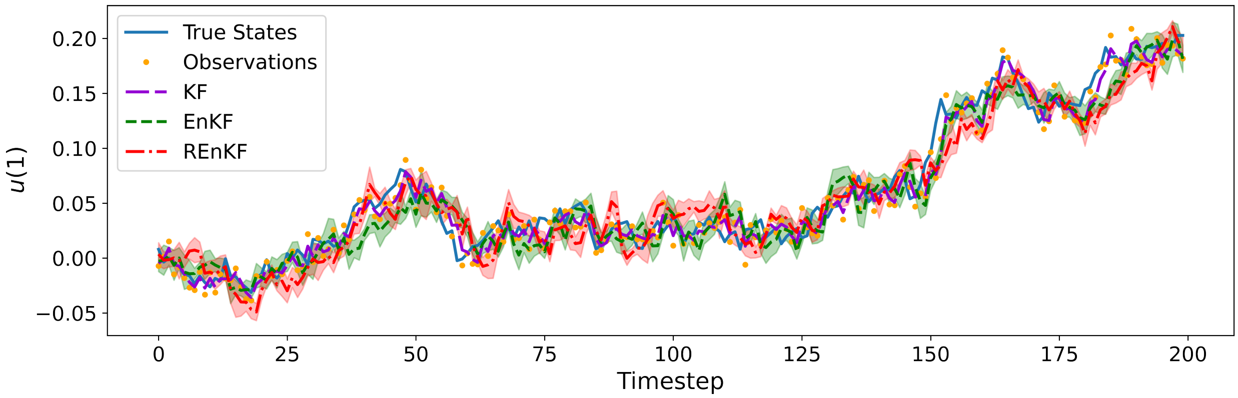

Figure 3.1: State estimation and uncertainty quantification for coordinate $u(1)$ in the linear setting with ensemble size $N = 10$ and small noise $\alpha = 10^{-4}$. Note that the Kaman Filter (KF) is optimal in the linear setting.

assess the ability of the filter to provide reliable uncertainty quantification: a short interval with high coverage would be preferable, but an overconfident short width interval with low coverage can lead to a misleading and potentially dangerous assessment of uncertainty. We illustrate these three metrics in Figure 3.1, which corresponds to a setup outlined in Table 3.2. This setup will be further explored in Subsection 3.4.1. As depicted in the plot, the indicator in (3.22) corresponds to whether the solid blue line (representing the true states) fall within the shaded confidence intervals. We point out that the ability of ensemble Kalman methods to provide reliable uncertainty quantification, especially in nonlinear settings, has often been questioned Ernst et al. [2015], Law and Stuart [2012]. Our results will show that the CIs obtained with REnKF have similar width and coverage as those obtained by EnKF, but that coverage for both algorithms is not reliable when the ensemble size is small (Subsections 3.4.1 and 3.4.2) or the dynamics are highly nonlinear (Subsection 3.4.2).

Since the outputs $\{\widehat{\mu}^{(j)}, \widehat{\Sigma}^{(j)}\}_{j=1}^{J}$ of EnKF and REnKF are random, for each experiment we run both algorithms M times and we report the average value of the metrics (3.20), (3.21), and (3.22) as well as the value of M. More details can be found in Appendix 3.7.1.

Effects of Noise Level and Ensemble Size

We perform two distinct analyses to assess the impact of different variables on the performance of EnKF and REnKF. The first, which we term the *noise-level* analysis, investigates the relationship between mean error, $\mathsf{E}_{\text{Linear}}$, and the noise level, α. The second analysis, referred to as the *ensemble-size analysis*, explores how the mean error varies with the ensemble size, N. Both analyses are carried out using a fixed state dimension $d = 20$.

In the noise-level analysis, α is varied over a grid of 15 evenly spaced values between 10^{-16} and 1, allowing us to investigate a range of scenarios beginning with those with virtually no noise to those with substantial noise. In order to isolate the influence of α, we maintain the initial distribution with a fixed zero mean and covariance $\Sigma^{(0)} = 10^{-8} \times I_{20}$, as well as a fixed ensemble size of $N = 20$. In the ensemble-size analysis, N is varied between 10 and 100, in increments of 10. To isolate the effects of N, we fix $\alpha = 10^{-1}$ and maintain the initial distribution to have a fixed zero mean and covariance $\Sigma^{(0)} = 1.1\alpha \times I_{20}$. The covariance is adjusted to represent a higher initial uncertainty level compared to the noise-level analysis. The factor 1.1 was introduced to ensure that the initial states possess a slightly different level of uncertainty relative to the noise in the dynamics and observations. Both analyses are averaged over $M = 10$ runs of the algorithms. The results of both analyses are depicted in Figure 3.2. In addition to $\mathsf{E}_{\text{Linear}}$, in Table 3.2 we consider the effect of varying α and N on CI widths, W, and CI coverage, V. Here, we categorize the levels of noise as being either small, moderate, or large, which correspond to α values of 10^{-4}, 10^{-2}, or 10^{-1} respectively, as described under Case A in Table 3.3. Further, we repeat the experiments with ensembles of size $N = 10$ and $N = 40$. For the experimental settings summarized in Table 3.2, the state dimension and initial distribution are taken as in the ensemble-size analysis described earlier. These metrics are calculated based on averages over $M = 100$ runs of the algorithms.

The results in Figure 3.2 and in Table 3.2 confirm that across a wide variety of linear experimental settings, REnKF exhibits similar performance to EnKF as measured by the mean

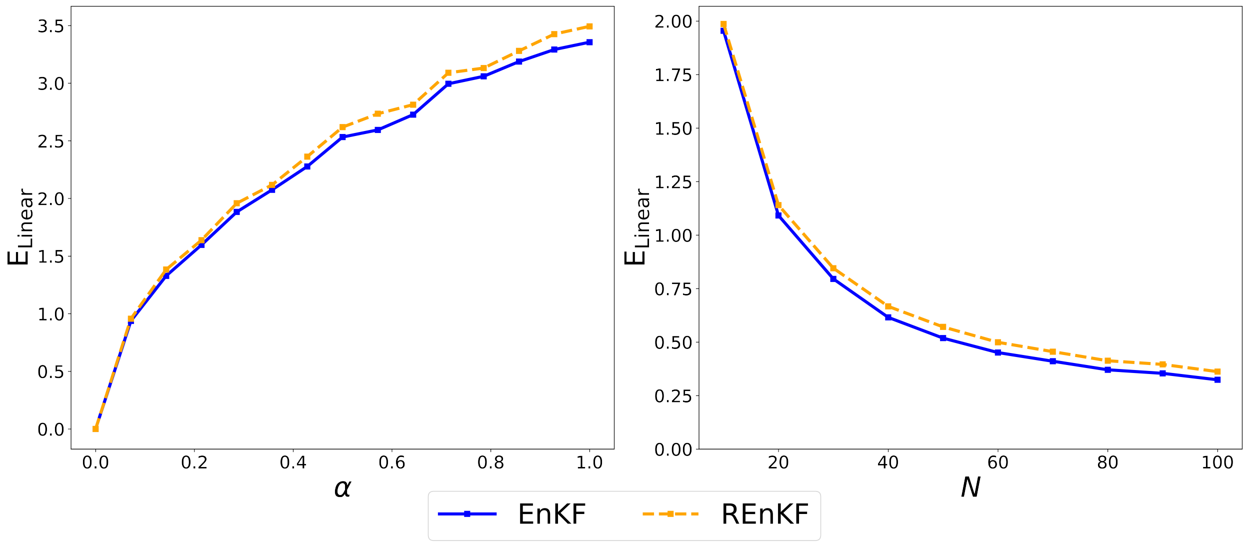

Figure 3.2: Effects of α and N in the linear setting with $d = 20$.

error, CI width, and CI coverage.

Effects of State Dimension and Spectrum Decay

We now study the sensitivity of EnKF and REnKF to changes in the state dimension, d. Recall that our main result, Theorem 3.3.2, implies that REnKF performs well whenever the ensemble size scales with the largest of the effective dimensions of the noise covariances: $\Sigma^{(0)}$, Γ, and Ξ. This motivates our study of covariance matrices with structure summarized in Case A and Case B of Table 3.3. In Case A, the effective dimension of the covariance matrix is proportional to the state dimension, d, and so the theory suggests that REnKF will do well only if the ensemble size also scales with d. In Case B, we consider covariance matrices that are diagonal, with i-th diagonal element proportional to $i^{-\beta}$ where $\beta > 0$ is a rate parameter controlling the speed of decay. Table 3.4 demonstrates that two matrices of this form that are equal in dimension may differ drastically in their effective dimension for different choices of β. Here, then, the theory suggests that REnKF will do well so long as the ensemble size scales with the effective dimension, which may be much smaller than d. To test our theory, we run REnKF under both cases A and B in Table 3.3 where d is varied over the set

Ensemble	Metric	Small Noise $\alpha = 10^{-4}$	Moderate Noise $\alpha = 10^{-2}$	Large Noise $\alpha = 10^{-1}$
$N = 10$	EnKF Mean Error	0.0608	0.6133	1.9931
	REnKF Mean Error	0.0616	0.6199	2.0310
	EnKF CI Width	0.0194	0.1940	0.6134
	REnKF CI Width	0.0188	0.1875	0.5930
	EnKF CI Coverage (%)	39.57	38.90	38.35
	REnKF CI Coverage (%)	37.83	37.14	36.58
$N = 40$	EnKF Mean Error	0.0193	0.1930	0.6243
	REnKF Mean Error	0.0209	0.2091	0.6739
	EnKF CI Width	0.0278	0.2780	0.8790
	REnKF CI Width	0.0274	0.2739	0.8663
	EnKF CI Coverage (%)	69.94	69.26	68.90
	REnKF CI Coverage (%)	68.65	67.76	67.43

Table 3.2: Performance metrics in the linear setting with $d = 20$.

Noise	Case A	Case B ($i = 1, \ldots, d$)	Case C
Dynamics (Ξ)	$\Xi^A = \alpha \times I_d$	$\Xi^B_{ii} = \alpha \times i^{-\beta}$	$\Xi^C = \alpha \times I_d$
Observation (Γ)	$\Gamma^A = \alpha \times I_d$	$\Gamma^B_{ii} = \alpha \times i^{-\beta}$	$\Gamma^C = \alpha \times I_{\frac{2d}{3}}$
Prior ($\Sigma^{(0)}$)	$(\Sigma^{(0)})^A = 1.1 \times \Xi^A$	$(\Sigma^{(0)})^B_{ii} = 1.1 \times \Xi^B_{ii}$	$(\Sigma^{(0)})^C = 1.1 \times \Xi^C$

Table 3.3: Covariance matrix settings explored numerically in Subsections 3.4.1 and 3.4.2.

$\{2^1, 2^2, \ldots, 2^8\}$ and where the ensemble size is fixed at $N = 10$ throughout. For both cases we fix $\alpha = 10^{-4}$ and for case B we consider $\beta \in \{0.1, 1, 1.5\}$. Figure 3.3 presents the results of averaging $\mathsf{E}_{\text{Linear}}$ over $M = 10$ runs of the algorithm in each of the experimental set-ups. We see that for all choices of β, EnKF and REnKF exhibit near-identical performance. For Case A, the performance deteriorates as d increases and this behavior is identical across all three displays. For Case B, when $\beta = 0.1$ (first display) so that the effective dimension increases significantly with dimension as described in the first row of Table 3.4, the performance

State dimension (d)	2	4	8	16	32	64	128	256
$\beta = 0.1$	1.93	3.70	7.02	13.25	24.89	46.64	87.25	163.05
$\beta = 1.0$	1.50	2.08	2.72	3.38	4.06	4.74	5.43	6.12
$\beta = 1.5$	1.35	1.67	1.93	2.12	2.26	2.36	2.44	2.49

Table 3.4: Effective dimension of initialization and noise covariances used in Figure 3.3.

deteriorates significantly as d increases. As β is increased to 1 in the second display, so that the effective dimension grows slowly with d, performance deteriorates at a much slower rate. This is further pronounced in the final display with $\beta = 1.5$. These numerical results demonstrate the key role played by the effective dimension in determining the performance of EnKF and REnKF, and are in agreement with Theorem 3.3.2 for REnKF.

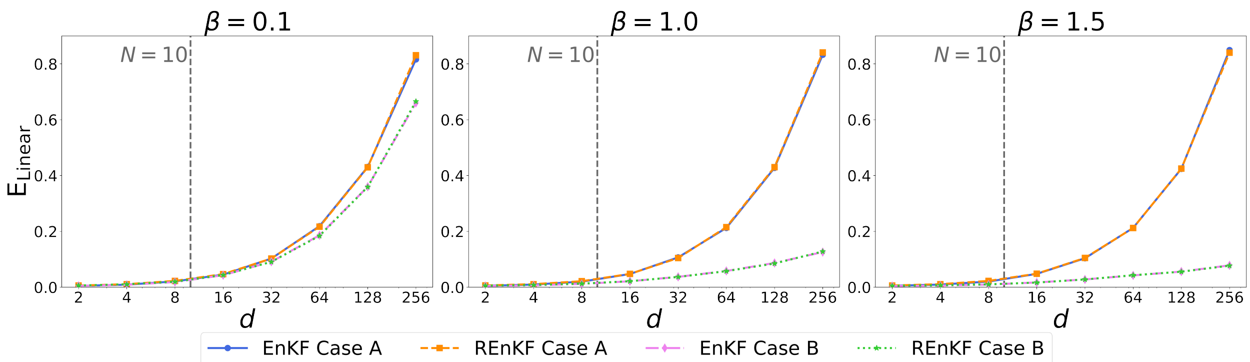

Figure 3.3: Effect of spectrum decay in the linear setting.

3.4.2 Lorenz 96 Dynamics

In this subsection, we extend our numerical investigation of REnKF to the nonlinear setting by taking Ψ in (3.2) to be the Δt-flow of the Lorenz 96 equations. Here Δt represents the time-span between observations, which is assumed to be constant. Assuming the following cyclic boundary conditions $u(-1) = u(d-1)$, $u(0) = u(d)$, and $u(d+1) = u(1)$ with $d \geq 4$, the system is governed by:

$$\frac{du(i)}{dt} = \Big(u(i+1) - u(i-2)\Big)u(i-1) - u(i) + F, \quad i = 1, \dots, d. \tag{3.23}$$

In our experiments, we set $\Delta t = 0.01$, $F = 8$, and the state dimension d is subject to variation. The choice $F = 8$ leads to strongly chaotic turbulence, which hinders predictability in the absence of observations Majda and Harlim [2012]. For the observation process (3.3), we consider both full observations in which $H = I_d$, and partial observations in which only two out of every three state components are observed. The latter setting results in a modified $H \in \mathbb{R}^{\frac{2d}{3} \times d}$ which corresponds to I_d with every third row removed. This observation set-up is motivated by Sanz-Alonso and Stuart [2015], Law et al. [2015], which prove that observing two-out-of-three coordinates of the Lorenz 96 system suffices in order to tame the unpredictability of the system and achieve long-time filter accuracy in a small noise regime. As in Subsection 3.4.1, we examine various choices of initial distribution, dynamics noise covariance, and observation noise covariance. To compare EnKF and REnKF, we make use of the same CI width (3.21) and CI coverage (3.22) metrics as in Subsection 3.4.1. However, since in the nonlinear setting the mean of the filtering distribution is not available in closed form, we replace the metric $\mathsf{E}_{\text{Linear}}$ with

$$\mathsf{E}_{\text{L96}} = \frac{1}{J} \sum_{j=1}^{J} |\widehat{\mu}^{(j)} - u^{(j)}|_2, \tag{3.24}$$

which quantifies the accuracy of the filter as an estimator of the ground-truth state process $\{u^{(j)}\}_{j=1}^{J}$. As before, the metrics we report are averaged over M runs of the algorithms.

		Full Observation			Partial Observation		
Ensemble	Metric	Small Noise $\alpha = 10^{-4}$	Moderate Noise $\alpha = 10^{-2}$	Large Noise $\alpha = 10^{-1}$	Small Noise $\alpha = 10^{-4}$	Moderate Noise $\alpha = 10^{-2}$	Large Noise $\alpha = 10^{-1}$
$N = 21$	EnKF Mean Error	0.1011	0.9573	3.0231	0.4064	3.3882	10.5921
	REnKF Mean Error	0.1016	0.9616	3.0335	0.4071	3.3565	10.6379
	EnKF CI Width	0.0208	0.2083	0.6586	0.0266	0.2660	0.8412
	REnKF CI Width	0.0205	0.2047	0.6475	0.0258	0.2584	0.8167
	EnKF CI Coverage (%)	50.24	51.55	51.61	39.62	43.25	43.26
	REnKF CI Coverage (%)	49.07	50.34	50.44	38.25	42.04	41.87
$N = 84$	EnKF Mean Error	0.0582	0.5682	1.7971	0.2919	2.4181	7.6282
	REnKF Mean Error	0.0590	0.5760	1.8218	0.2977	2.5004	7.9011
	EnKF CI Width	0.0281	0.2813	0.8895	0.0438	0.4383	1.3861
	REnKF CI Width	0.0279	0.2785	0.8806	0.0412	0.4120	1.3033
	EnKF CI Coverage (%)	87.96	88.61	88.61	71.47	75.31	75.30
	REnKF CI Coverage (%)	86.80	87.52	87.52	69.25	72.54	72.61

Table 3.5: Performance metrics for the Lorenz 96 model with $d = 42$.

In Table 3.5, we compare the performance of REnKF and EnKF. In the case of full observations, the covariance configuration is outlined in Case A of Table 3.3, and in the case of partial observations it is outlined in Case C of Table 3.3. We repeat the experiments with ensembles of size $N = 21$ and $N = 84$, and the metrics are computed over $M = 100$ runs of the algorithms. In Figure 3.4, we present a single representative simulation of the first component $u(1)$ —which is observed— and the third component $u(3)$ —which is unobserved— corresponding to a particular choice of parameters in Table 3.5. Additional experiments in the accompanying Github repository show that, as the noise level α increases, state estimation remains effective for observed variables but deteriorates for unobserved ones. This behavior explains the larger error for moderate and large noise levels in the partial observation set-up in Table 3.5.

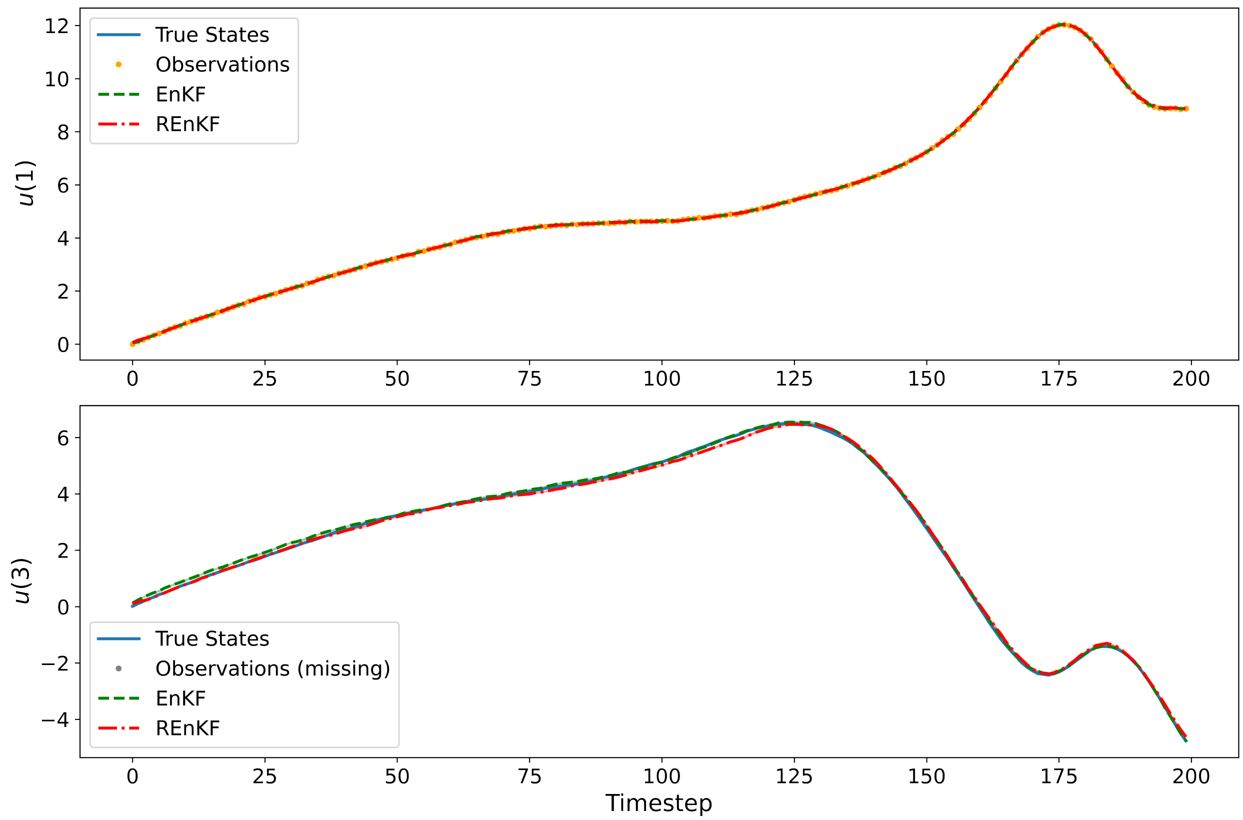

Figure 3.4: State estimation of coordinates $u(1)$ (observed) and $u(3)$ (unobserved) in a partially observed Lorenz 96 system with ensemble size $N = 21$ and small noise $\alpha = 10^{-4}$. REnKF accurately recovers observed and unobserved coordinates of the state.

In Figure 3.5, we further analyze the effects of varying α (column 1), N (column 2), and d (column 3) on $\mathsf{E}_{\mathrm{L96}}$ in both the full observation (row 1) and partial observation (row 2) settings. More precisely, in the first column of Figure 3.5, α is varied over a grid of 15

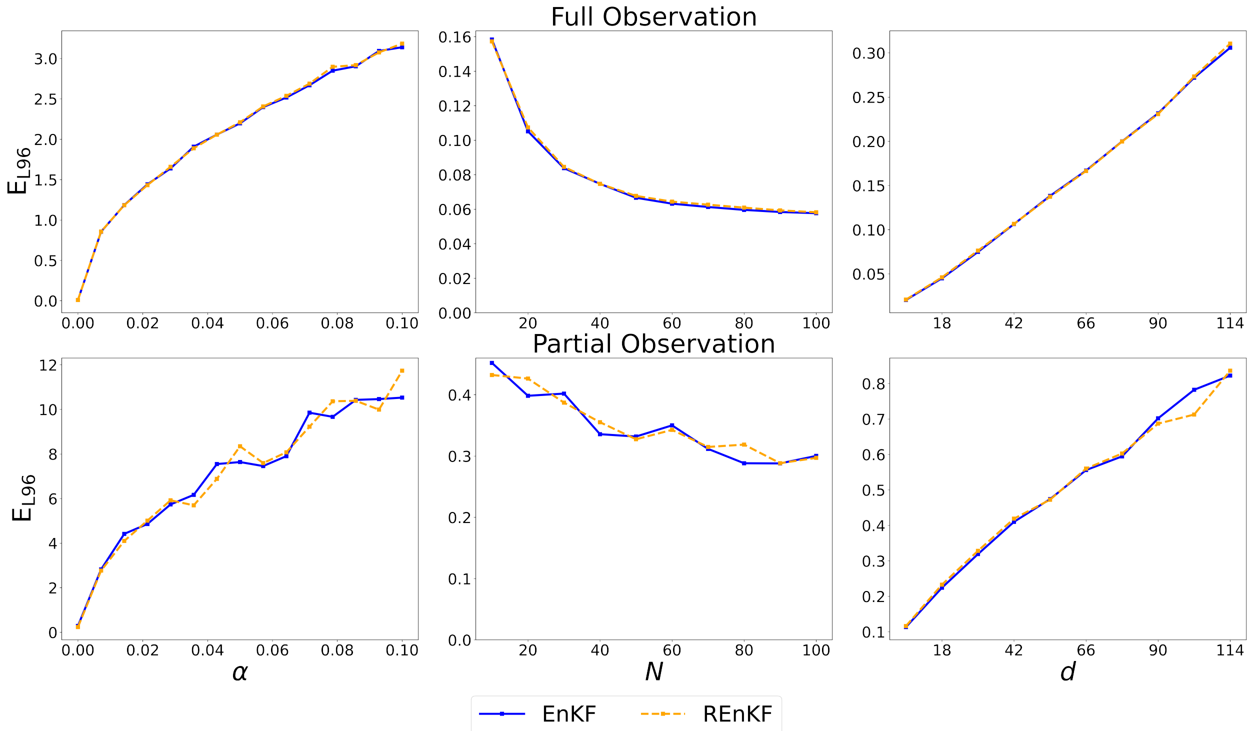

Figure 3.5: Effects of α, N, and d in the Lorenz 96 example.

evenly spaced values between 10^{-16} and 1 while holding fixed $N = 20$ and $d = 42$. In both full and partial observation settings, we take the initial distribution to have zero mean and covariance $\Sigma^{(0)} = 10^{-8} \times I_{42}$. In the second column of Figure 3.5, the ensemble size N ranges from 10 to 100, increasing in steps of 10, while fixing $\alpha = 10^{-4}$ and $d = 42$. In both full and partial observation settings, we take the initial distribution to have zero mean and covariance $\Sigma^{(0)} = 1.1\alpha \times I_{42}$. In the third column of Figure 3.5, the dimension d is varied over the values in $\{6, 18, 30, 42, 54, 66, 78, 90, 102\}$ which are all multiples of 3 to facilitate convenient calculations in the partially observed setting. We fix $N = 20$ and $\alpha = 10^{-4}$ and in both full and partial observation settings, we take the initial distribution to have zero mean and covariance $\Sigma^{(0)} = 1.1\alpha \times I_d$, respectively.

Our findings, as illustrated in Table 3.5 and Figure 3.5, demonstrate that REnKF achieves performance comparable to that of EnKF, even in challenging nonlinear regimes. Notably, for both algorithms we observe a slightly inferior performance with partial observations compared to full observations under identical conditions. Moreover, a consistent trend is noticed in the dependency of $\mathsf{E}_{\mathsf{L96}}$ on the noise level, state dimension, and ensemble size. Notice, however, that the performance of REnKF deteriorates further in non-Gaussian settings with partial observations, large N, and large noise. Such worsened performance may be partly explained by the additional Gaussian assumption tacitly imposed in the resampling step, which further destroys the non-Gaussian structure of the problem for nonlinear forward models. Table 3.5 further demonstrates that REnKF is as effective as EnKF in the task of uncertainty quantification. Nevertheless, both EnKF and REnKF encounter difficulties in delivering reliable uncertainty quantification, especially in scenarios with partial observation and small ensemble size.

3.5 Proof of Theorem 3.3.2

The result will be established by strong induction on the *mean bound* (3.17) and the *covariance bound* (3.18) along with induction on two additional bounds: for any $j = 1, 2, \dots$ and $q \geq 1$

$$|||\mathrm{Tr}(\widehat{\Sigma}^{(j-1)})|||_q \leq c_3 r_2(\Sigma^{(0)}), \tag{3.25}$$

$$|||\widehat{C}^{(j)} - C^{(j)}|||_q \leq c_4 \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \right), \tag{3.26}$$

where c_3 and c_4 are again potentially different universal constants that depend on the same parameters as c_2 in the statement of Theorem 3.3.2. We will refer to (3.25) as the *covariance trace bound* and to (3.26) as the *forecast covariance bound*. In this section, we require the following additional notation: given two positive sequences $\{a_n\}$ and $\{b_n\}$, the relation $a_n \lesssim b_n$ denotes that $a_n \leq cb_n$ for some constant $c > 0$. If the constant c depends on some quantity τ, then we write $a \lesssim_\tau b$. Throughout, we denote positive universal constants by c, c_1, c_2, c_3, c_4, and the value of a universal constant may differ from line to line. In some cases, the explicit dependence of a universal constant on the parameter τ is indicated by writing $c(\tau)$.

This section is organized as follows. Subsection 3.5.1 contains preliminary results. We then prove the base case $j = 1$ in Subsection 3.5.2. Finally, in Subsection 3.5.3 we show that the bounds (3.17), (3.18), (3.25), and (3.26) hold for j assuming they hold for all $\ell \leq j-1$.

3.5.1 Preliminary Results

Lemma 3.5.1 (Operator Norm of Covariance). *For any $j \geq 0$, let $\Sigma^{(j)}$ be the analysis covariance at iteration j. Then,*

$$|\Sigma^{(j)}| \leq |A|^{2j}|\Sigma^{(0)}| + |\Xi| \sum_{\ell=0}^{j-1} |A|^{2\ell} \leq c(|A|, |\Xi|, |\Sigma^{(0)}|, j).$$

Proof. By Lemma 3.7.6, $|\mathscr{C}(C)| \leq |C|$, and so

$$\begin{aligned} |\Sigma^{(j)}| = |\mathscr{C}(C^{(j)})| \leq |C^{(j)}| = |A\Sigma^{(j-1)}A^\top + \Xi| \leq |A|^2|\Sigma^{(j-1)}| + |\Xi| \\ \leq |A|^4|\Sigma^{(j-2)}| + |A|^2|\Xi| + |\Xi| \leq \cdots \leq |A|^{2j}|\Sigma^{(0)}| + |\Xi| \sum_{\ell=0}^{j-1} |A|^{2\ell}. \end{aligned}$$

□

Lemma 3.5.2 (Trace of Offset). *For any $j \geq 1$, we have that*

$$\begin{aligned} Tr(\widehat{O}^{(j)}) \leq |H|^2|\hat{\Gamma}^{(j)} - \Gamma||\Gamma^{-1}|^2|\widehat{C}^{(j)}|\, Tr(\widehat{C}^{(j)}) \\ + 2(1 + |\widehat{C}^{(j)}||H|^2|\Gamma^{-1}|)|\Gamma^{-1}||\widehat{C}^{(j)}_{u\eta}||H|\, Tr(\widehat{C}^{(j)}). \end{aligned}$$

Proof. Write $\widehat{O}^{(j)} = \sum_{\ell=1}^{3} \widehat{O}_\ell^{(j)}$ with

$$
\begin{aligned}
\widehat{O}_1^{(j)} &:= \mathscr{K}(\widehat{C}^{(j)})(\widehat{\Gamma}^{(j)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(j)}), \\
\widehat{O}_2^{(j)} &:= \big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)\widehat{C}_{u\eta}^{(j)}\mathscr{K}^\top(\widehat{C}^{(j)}), \\
\widehat{O}_3^{(j)} &:= \mathscr{K}(\widehat{C}^{(j)})(\widehat{C}_{u\eta}^{(j)})^\top\big(I - H^\top\mathscr{K}^\top(\widehat{C}^{(j)})\big).
\end{aligned}
$$

By linearity of the trace, $\mathrm{Tr}(\widehat{O}^{(j)}) = \sum_{\ell=1}^{3} \mathrm{Tr}(\widehat{O}_\ell^{(j)})$. Note first that by Lemma 3.7.6 applied four times

$$
\begin{aligned}
\mathrm{Tr}(\hat{O}_1^{(j)}) &\le |\hat{\Gamma}^{(j)} - \Gamma|\mathrm{Tr}(\mathscr{K}^\top(\widehat{C}^{(j)})\mathscr{K}(\widehat{C}^{(j)})) \\
&= |\hat{\Gamma}^{(j)} - \Gamma|\mathrm{Tr}((H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}H(\widehat{C}^{(j)})^\top\widehat{C}^{(j)}H^\top(H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}) \\
&\le |\hat{\Gamma}^{(j)} - \Gamma||(H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}|^2\mathrm{Tr}((\widehat{C}^{(j)})^\top\widehat{C}^{(j)}H^\top H) \\
&\le |H|^2|\hat{\Gamma}^{(j)} - \Gamma||(H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}|^2|\widehat{C}^{(j)}|\mathrm{Tr}(\widehat{C}^{(j)}) \\
&\le |H|^2|\hat{\Gamma}^{(j)} - \Gamma||\Gamma^{-1}|^2|\widehat{C}^{(j)}|\mathrm{Tr}(\widehat{C}^{(j)}),
\end{aligned}
$$

where the final inequality holds since $H\widehat{C}^{(j)}H^\top + \Gamma \succeq \Gamma$ implies that $\Gamma^{-1} \succeq (H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}$. Invoking once more Lemma 3.7.6 repeatedly, we get that

$$
\begin{aligned}
\mathrm{Tr}(\hat{O}_2^{(j)}) &= \mathrm{Tr}\big(\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)\widehat{C}_{u\eta}^{(j)}\mathscr{K}^\top(\widehat{C}^{(j)})\big) \\
&\le |\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)|\mathrm{Tr}\big(\widehat{C}_{u\eta}^{(j)}\mathscr{K}^\top(\widehat{C}^{(j)})\big) \\
&\le |\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)||(H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}|\mathrm{Tr}(H\widehat{C}^{(j)}\widehat{C}_{u\eta}^{(j)}) \\
&\le |\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)||(H\widehat{C}^{(j)}H^\top + \Gamma)^{-1}||\widehat{C}_{u\eta}^{(j)}||H|\mathrm{Tr}(\widehat{C}^{(j)}) \\
&\le (1 + |\widehat{C}^{(j)}||H|^2|\Gamma^{-1}|)|\Gamma^{-1}||\widehat{C}_{u\eta}^{(j)}||H|\mathrm{Tr}(\widehat{C}^{(j)}),
\end{aligned}
$$

where the last inequality holds since, by Lemma 3.7.1,

$$|\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)| \leq 1 + |\mathscr{K}(\widehat{C}^{(j)})||H| \leq 1 + |\widehat{C}^{(j)}||H|^2|\Gamma^{-1}|.$$

Finally, note that since $\hat{O}_3^{(j)} = (\hat{O}_2^{(j)})^\top$, the analysis of $\hat{O}_3^{(j)}$ follows in similar fashion. □

3.5.2 Base Case

In the next four subsections we establish the covariance trace bound (3.25), the forecast covariance bound (3.26), the mean bound (3.17), and the covariance bound (3.18) in the base case $j = 1$.

Covariance Trace Bound

Since $\widehat{\Sigma}^{(0)} = \Sigma^{(0)}$, we directly obtain that

$$\||\mathrm{Tr}(\widehat{\Sigma}^{(0)})|\|_q = \mathrm{Tr}(\Sigma^{(0)}) = |\Sigma^{(0)}|r_2(\Sigma^{(0)}).$$

Forecast Covariance Bound

Let $\mathfrak{q} \in \{q, 2q, 4q\}$. It follows by the triangle inequality, Theorem 3.7.5, and Lemma 3.7.7 that

$$\begin{aligned}
|||\widehat{C}^{(1)} - C^{(1)}|||_{\mathfrak{q}} &\le |A|^2 |||S^{(0)} - \Sigma^{(0)}|||_{\mathfrak{q}} + |||\hat{\Xi}^{(1)} - \Xi|||_{\mathfrak{q}} + 2|A||||\widehat{C}^{(1)}_{u\xi}|||_{\mathfrak{q}} \\
&\lesssim_q |A|^2 |\Sigma^{(0)}| \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} + |\Xi| \sqrt{\frac{r_2(\Xi)}{N}} \\
&+ 2|A|(|\Sigma^{(0)}| \vee |\Xi|) \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \right) \\
&\le c(|A|, |\Sigma^{(0)}|, |\Xi|, q) \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \right). \qquad (3.27)
\end{aligned}$$

Mean Bound

By Lemma 3.7.2,

$$\begin{aligned}
|||\widehat{\mu}^{(1)} - \mu^{(1)}|||_q &= |||\mathscr{M}(\widehat{m}^{(1)}, \widehat{C}^{(1)}; y^{(1)}) - \mathscr{M}(m^{(1)}, C^{(1)}; y^{(1)})|||_q + |||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q \\
&\le \Big\| |\widehat{m}^{(1)} - m^{(1)}| \Big\|_q + |H|^2 |\Gamma^{-1}| |||\widehat{C}^{(1)}|||_{2q} |||\widehat{m}^{(1)} - m^{(1)}|||_{2q} \\
&+ |||\widehat{C}^{(1)} - C^{(1)}|||_q |H||\Gamma^{-1}| \big(1 + |H|^2 |\Gamma^{-1}||C^{(1)}|\big) |y^{(1)} - Hm^{(1)}| \\
&+ |||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q.
\end{aligned}$$

The mean bound (3.17) with $j = 1$ is then a direct consequence of the bounds that we now establish on $|||\widehat{m}^{(1)} - m^{(1)}|||_{\mathfrak{q}}$ for $\mathfrak{q} \in \{q, 2q\}$ and on $|||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q$.

Controlling $|||\widehat{m}^{(1)} - m^{(1)}|||_{\mathfrak{q}}$ for $\mathfrak{q} \in \{q, 2q\}$ It follows by the triangle inequality and Lemma 3.7.9 applied twice that

$$\begin{aligned}|||\widehat{m}^{(1)} - m^{(1)}|||_{\mathfrak{q}} &\leq |A| |||\bar{u}^{(0)} - \mu^{(0)}|||_{\mathfrak{q}} + |||\bar{\xi}^{(1)}|||_{\mathfrak{q}} \lesssim_q |A||\Sigma^{(0)}|\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} + |\Xi|\sqrt{\frac{r_2(\Xi)}{N}} \\ &\leq c(|A|, |\Sigma^{(0)}|, |\Xi|, q)\left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}}\right).\end{aligned}$$

Controlling $|||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q$ By Cauchy-Schwarz

$$|||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q \leq |||\mathscr{K}(\widehat{C}^{(1)})|||_{2q} |||\bar{\eta}^{(1)}|||_{2q}.$$

We bound each term in turn. By Lemma 3.7.9, $|||\bar{\eta}^{(1)}|||_{2q} \lesssim_q \sqrt{\mathrm{Tr}(\Gamma)/N} = \sqrt{|\Gamma| r_2(\Gamma)/N}$, and by Lemma 3.7.1 and the forecast covariance bound (3.27),

$$\begin{aligned}|||\mathscr{K}(\widehat{C}^{(1)})|||_{2q} &\leq |H||\Gamma^{-1}| |||\widehat{C}^{(1)}|||_{2q} \leq |H||\Gamma^{-1}|\left(|||\widehat{C}^{(1)} - C^{(1)}|||_{2q} + |C^{(1)}|\right) \\ &\lesssim |H||\Gamma^{-1}||C^{(1)}| c(|A|, |\Sigma^{(0)}|, |\Xi|, q)\left(1 \vee \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}}\right).\end{aligned}$$

Therefore,

$$|||\mathscr{K}(\widehat{C}^{(1)})\bar{\eta}^{(1)}|||_q \leq c(|H|, |\Sigma^{(0)}|, |\Xi|, |H|, |\Gamma^{-1}|, |\Gamma|, q)\left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right).$$

Covariance Bound

From Lemma 3.7.3 and the forecast covariance bound derived in Subsection 3.5.2, we have

$$\begin{aligned}
|||\widehat{\Sigma}^{(1)} - \Sigma^{(1)}|||_q &\le |||\mathscr{C}(\widehat{C}^{(1)}) - \mathscr{C}(C^{(1)})|||_q + |||\hat{O}|||_q \\
&\le |||\widehat{C}^{(1)} - C^{(1)}|||_q (1 + |H|^2 |\Gamma^{-1}||C^{(1)}|) \\
&+ (|H|^2|\Gamma^{-1}| + |H|^4 |\Gamma^{-1}|^2 |C^{(1)}|) |||\widehat{C}^{(1)}|||_{2q} |||\widehat{C}^{(1)} - C^{(1)}|||_{2q} + |||\hat{O}^{(1)}|||_q \\
&\le c(|A|, |H|, |\Gamma^{-1}|, |\Sigma^{(0)}|, |\Xi|, q) \left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \right) + |||\hat{O}^{(1)}|||_q.
\end{aligned}$$

To derive the covariance bound (3.18), we need to control the offset term $|||\widehat{O}^{(1)}|||_q$. First, using the triangle inequality we write

$$\begin{aligned}
|||\widehat{O}^{(1)}|||_q &\le |||\mathscr{K}(\widehat{C}^{(1)})(\widehat{\Gamma}^{(1)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(1)})|||_q + |||\big(I - \mathscr{K}(\widehat{C}^{(1)})H\big)\widehat{C}^{(1)}_{u\eta}\mathscr{K}^\top(\widehat{C}^{(1)})|||_q \\
&+ |||\mathscr{K}(\widehat{C}^{(1)})(\widehat{C}^{(1)}_{u\eta})^\top\big(I - H^\top \mathscr{K}^\top(\widehat{C}^{(1)})\big)|||_q \\
&=: |||\hat{O}^{(1)}_1|||_q + |||\hat{O}^{(1)}_2|||_q + |||\hat{O}^{(1)}_3|||_q.
\end{aligned}$$

We next bound each term in turn.

Controlling $|||\hat{O}^{(1)}_1|||_q$ By Lemma 3.7.1, Theorem 3.7.5, and the forecast covariance bound (3.27), it holds that

$$\begin{aligned}
|||\mathscr{K}(\widehat{C}^{(1)})(\widehat{\Gamma}^{(1)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(1)})|||_q &\le |||\mathscr{K}(\widehat{C}^{(1)})|||_{4q}^2 |||\widehat{\Gamma}^{(1)} - \Gamma|||_{2q} \\
&\le |H|^2 |\Gamma^{-1}|^2 |||\widehat{C}^{(1)}|||_{4q}^2 |||\widehat{\Gamma}^{(1)} - \Gamma|||_{2q} \\
&\lesssim_q |H|^2 |\Gamma^{-1}|^2 |\Gamma| \sqrt{\frac{r_2(\Gamma)}{N}} \left(1 \vee \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \right)^2 \\
&\le c(|H|, |\Gamma|, |\Gamma^{-1}|, |\Xi|, |\Sigma^{(0)}|, q) \sqrt{\frac{r_2(\Gamma)}{N}},
\end{aligned}$$

where the last inequality uses the fact that $N \geq r_2(\Sigma^{(0)}) \vee r_2(\Xi) \vee r_2(\Gamma)$.

Controlling $|||\hat{O}_2^{(1)}|||_q$ By Lemma 3.7.1, Lemma 3.7.7, and the forecast covariance bound (3.27), we get

$$
\begin{aligned}
|||\big(I - \mathscr{K}(\widehat{C}^{(1)})H\big)\widehat{C}^{(1)}_{u\eta}\mathscr{K}^\top(\widehat{C}^{(1)})|||_q &\leq |||\mathscr{K}(\widehat{C}^{(1)})||I - \mathscr{K}(\widehat{C}^{(1)})H||\widehat{C}^{(1)}_{u\eta}|||_q \\
&\leq |||\mathscr{K}(\widehat{C}^{(1)})|(1 + |\mathscr{K}(\widehat{C}^{(1)})||H|)|\widehat{C}^{(1)}_{u\eta}|||_q \\
&\leq |||\widehat{C}^{(1)}_{u\eta}|||_{2q}\left(|H||\Gamma^{-1}||||\widehat{C}^{(1)}|||_{2q} + |H|^3|\Gamma^{-1}|^2|||\widehat{C}^{(1)}|^2\|_{2q}\right) \\
&\leq c(|A|, |\Sigma^{(0)}|, |\Xi|, q)\left(1 \vee \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right) \\
&\quad \times \left(|H|^3|\Gamma^{-1}| + |H|^7|\Gamma^{-1}|^2\right)(|\Sigma^{(0)}| \vee |\Gamma|)\left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right) \\
&\leq c(|A|, |H|, |\Gamma^{-1}|, |\Gamma|, |\Sigma^{(0)}|, |\Xi|, q)\left(\sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right).
\end{aligned}
$$

Controlling $|||\hat{O}_3^{(1)}|||_q$ Note that $|||\hat{O}_3^{(1)}|||_q = |||\hat{O}_2^{(1)}|||_q$.

3.5.3 Induction Step

In this subsection, to reduce notation we write $\Omega = \sqrt{\frac{r_2(\Sigma^{(0)})}{N}} \vee \sqrt{\frac{r_2(\Xi)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}$. Throughout, we work under the *inductive hypothesis* that, for all $\ell \leq j-1$, it holds that

$$
\begin{aligned}
|||\mathrm{Tr}(\widehat{\Sigma}^{(\ell-1)})|||_q &\leq c_1 r_2(\Sigma^{(0)}), \qquad & |||\widehat{C}^{(\ell)} - C^{(\ell)}|||_{\mathfrak{q}} &\leq c_2\Omega, \\
|||\widehat{\mu}^{(\ell)} - \mu^{(\ell)}|||_q &\leq c_3\Omega, & |||\widehat{\Sigma}^{(\ell)} - \Sigma^{(\ell)}|||_q &\leq c_4\Omega,
\end{aligned}
\tag{3.28}
$$

where c_1, c_2, c_3, and c_4 are constants depending on $|\Sigma^{(0)}|, |A|, |H|, |\Gamma^{-1}|, |\Gamma|, |\Xi|, q$ and j, and c_3 additionally depends on $\{|y^{(i)} - Hm^{(i)}|\}_{i \leq \ell-1}$. For the remainder of the proof, c and c' denote constants that depend on $|\Sigma^{(0)}|, |A|, |H|, |\Gamma^{-1}|, |\Gamma|, |\Xi|, q$ and j, and c' additionally depends on $\{|y^{(i)} - Hm^{(i)}|\}_{i \le \ell-1}$ and are potentially different from line to line. In the next four subsections we show that, under the inductive hypothesis, the four bounds in (3.28) also hold for $\ell = j$. Throughout, we use without further notice that $|\Sigma^{(\ell)}| \lesssim c(|A|, |\Xi|, |\Sigma^{(0)}|, \ell)$, which was proved in Lemma 3.5.1.

Covariance Trace Bound

By Lemma 3.7.3, $\mathrm{Tr}(\mathscr{C}(\widehat{C}^{(j-1)})) \le \mathrm{Tr}(\widehat{C}^{(j-1)})$ follows from the fact that $\mathscr{C}(\widehat{C}^{(j-1)}) \preceq \widehat{C}^{(j-1)}$, and so

$$\|\mathrm{Tr}(\widehat{\Sigma}^{(j-1)})\|_q \le \|\mathrm{Tr}(\mathscr{C}(\widehat{C}^{(j-1)}))\|_q + \|\mathrm{Tr}(\hat{O}^{(j-1)})\|_q \le \|\mathrm{Tr}(\widehat{C}^{(j-1)})\|_q + \|\mathrm{Tr}(\hat{O}^{(j-1)})\|_q,$$

We will show that both of the terms on the right-hand side are bounded above by a constant times $r_2(\Sigma^{(0)})$.

Controlling $\|\mathbf{Tr}(\widehat{C}^{(j-1)})\|_q$ Noting first that

$$\begin{aligned}\mathbb{E}\left[\widehat{C}^{(j-1)}\middle|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right] &= \mathbb{E}\left[AS^{(j-2)}A^\top + \hat{\Xi}^{(j-1)} + A\widehat{C}_{u\xi}^{(j-1)} + \widehat{C}_{\xi u}^{(j-1)}A^\top\middle|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right] \\ &= \mathbb{E}\left[AS^{(j-2)}A^\top\middle|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right] = A\widehat{\Sigma}^{(j-2)}A^\top,\end{aligned}$$

and by Lemma 3.7.6, it holds almost surely that

$$\mathrm{Tr}\big(A\widehat{\Sigma}^{(j-2)}A^\top\big) \le |A|^2\mathrm{Tr}(\widehat{\Sigma}^{(j-2)}) = |A|^2|\widehat{\Sigma}^{(j-2)}|r_2(\widehat{\Sigma}^{(j-2)}).$$

Then, by iterated expectations and Lemma 3.7.8, we have

$$
\begin{aligned}
\mathbb{E}\left[\mathrm{Tr}(\widehat{C}^{(j-1)})\right]^q &= \mathbb{E}\left[\mathbb{E}\left[\left(\mathrm{Tr}(\widehat{C}^{(j-1)})\right)^q \,\Big|\, \widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right]\right] \\
&\lesssim_q \mathbb{E}\left[\mathbb{E}\left[\left(\mathrm{Tr}\left(\widehat{C}^{(j-1)} - \mathbb{E}\big[\widehat{C}^{(j-1)}|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\big]\right)\right)^q \,\Big|\, \widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right]\right] \\
&\qquad\qquad + \mathbb{E}\left[\left(\mathrm{Tr}\left(\mathbb{E}\big[\widehat{C}^{(j-1)}|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\big]\right)\right)^q\right] \\
&\lesssim \mathbb{E}\left[\left(\frac{\mathrm{Tr}\big(\mathbb{E}\big[\widehat{C}^{(j-1)}|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\big]\big)}{\sqrt{N}}\right)^q\right] + \mathbb{E}\left[\left(\mathrm{Tr}\left(\mathbb{E}\left[\widehat{C}^{(j-1)}|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right]\right)\right)^q\right] \\
&\le \frac{|A|^{2q}}{N^{q/2}}\mathbb{E}\left[\left(\mathrm{Tr}(\widehat{\Sigma}^{(j-2)})\right)^q\right] + |A|^{2q}\mathbb{E}\left[\left(\mathrm{Tr}(\widehat{\Sigma}^{(j-2)})\right)^q\right] \\
&\lesssim \frac{|A|^{2q}}{N^{q/2}} c r_2(\Sigma^{(0)})^q + |A|^{2q} c r_2(\Sigma^{(0)})^q \le c r_2(\Sigma^{(0)})^q,
\end{aligned}
$$

where the second to last inequality holds by the inductive hypothesis (3.28).

Controlling $\|\hat{O}^{(j-1)}\|_q$ By definition, we have

$$
\begin{aligned}
\widehat{O}^{(j-1)} &= \mathscr{K}(\widehat{C}^{(j-1)})(\widehat{\Gamma}^{(j-1)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(j-1)}) + \big(I - \mathscr{K}(\widehat{C}^{(j-1)})H\big)\widehat{C}_{u\eta}^{(j-1)}\mathscr{K}^\top(\widehat{C}^{(j-1)}) \\
&\quad + \mathscr{K}(\widehat{C}^{(j-1)})(\widehat{C}_{u\eta}^{(j-1)})^\top\big(I - H^\top\mathscr{K}^\top(\widehat{C}^{(j-1)})\big) \\
&=: \hat{O}_1^{(j-1)} + \hat{O}_2^{(j-1)} + \hat{O}_3^{(j-1)}.
\end{aligned}
$$

Therefore, $\|\mathrm{Tr}(\hat{O}^{(j-1)})\|_q \le \|\mathrm{Tr}(\hat{O}_1^{(j-1)})\|_q + \|\mathrm{Tr}(\hat{O}_2^{(j-1)})\|_q + \|\mathrm{Tr}(\hat{O}_3^{(j-1)})\|_q$.

Controlling $\|\mathbf{Tr}(\hat{O}_1^{(j-1)})\|_q$ By Lemma 3.5.2,

$$
\mathrm{Tr}(\hat{O}_1^{(j-1)}) \le |H|^2|\widehat{\Gamma}^{(j-1)} - \Gamma||\Gamma^{-1}|^2|\widehat{C}^{(j-1)}|\mathrm{Tr}(\widehat{C}^{(j-1)}),
$$

and so

$$\begin{aligned}\|\mathrm{Tr}(\hat{O}_1^{(j-1)})\|_q &\leq |H|^2|\Gamma^{-1}|^2\||\hat{\Gamma}^{(j-1)} - \Gamma||\widehat{C}^{(j-1)}|\mathrm{Tr}(\widehat{C}^{(j-1)})\|_q \\ &\leq |H|^2|\Gamma^{-1}|^2\|\|\hat{\Gamma}^{(j-1)} - \Gamma|\|_{2q}\||\widehat{C}^{(j-1)}|\|\|_{4q}\|\mathrm{Tr}(\widehat{C}^{(j-1)})\|_{4q}.\end{aligned}$$

By Theorem 3.7.5, $\||\hat{\Gamma}^{(j-1)} - \Gamma|\|_{2q} \lesssim_q |\Gamma|\sqrt{\frac{r_2(\Gamma)}{N}}$, and by the inductive hypothesis (3.28) and the fact that $|C^{(j-1)}| \leq |A|^2|\Sigma^{(j-2)}| + |\Xi|$, we have

$$\||\widehat{C}^{(j-1)}|\|_{4q} \leq |C^{(j-1)}| + \||\widehat{C}^{(j-1)} - C^{(j-1)}|\|_{4q} \leq c\,(1 \vee \Omega)\,.$$

We have also previously shown that $\|\mathrm{Tr}(\widehat{C}^{(j-1)})\|_{4q} \lesssim r_2(\Sigma^{(0)})$. Noting that $(1 \vee \Omega)\, r_2(\Sigma^{(0)}) \lesssim r_2(\Sigma^{(0)})$, we get that $\|\mathrm{Tr}(\hat{O}_1^{(j-1)})\|_q \leq c r_2(\Sigma^{(0)})$.

Controlling $\||\mathbf{Tr}(\hat{O}_2^{(j-1)})|\|_q$ By Lemma 3.5.2,

$$\mathrm{Tr}(\hat{O}_2^{(j-1)}) \leq (1 + |\widehat{C}^{(j-1)}||H|^2|\Gamma^{-1}|)|\Gamma^{-1}||H||\widehat{C}_{u\eta}^{(j-1)}|\mathrm{Tr}(\widehat{C}^{(j-1)}).$$

Therefore,

$$\||\mathrm{Tr}(\hat{O}_2^{(j-1)})|\|_q \leq (1 + \||\widehat{C}^{(j-1)}|\|_{2q}|H|^2|\Gamma^{-1}|)|\Gamma^{-1}||H|\||\widehat{C}_{u\eta}^{(j-1)}|\|_{4q}\|\mathrm{Tr}(\widehat{C}^{(j-1)})\|_{4q}.$$

By iterated expectation, Lemma 3.7.7, and Lemma 3.7.8 we have, for $\mathfrak{q} \in \{q, 2q, 4q\}$,

$$
\begin{aligned}
\||\widehat{C}_{u\eta}^{(j-1)}|\|_{\mathfrak{q}} &= \mathbb{E}\left[\mathbb{E}\left[|\widehat{C}_{u\eta}^{(j-1)}|^{\mathfrak{q}}|\widehat{\mu}^{(j-2)}, \widehat{\Sigma}^{(j-2)}\right]\right]^{1/\mathfrak{q}} \\
&\lesssim_q \left\|(|\widehat{\Sigma}^{(j-2)}| \vee |\Gamma|)\left(\sqrt{\frac{r_2(\widehat{\Sigma}^{(j-2)})}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right)\right\|_{\mathfrak{q}} \\
&\le (\||\widehat{\Sigma}^{(j-2)}|\|_{2\mathfrak{q}} \vee |\Gamma|)\left(\left\|\sqrt{\frac{r_2(\widehat{\Sigma}^{(j-2)})}{N}}\right\|_{2\mathfrak{q}} \vee \sqrt{\frac{r_2(\Gamma)}{N}}\right). \qquad (3.29)
\end{aligned}
$$

By the triangle inequality and the inductive hypothesis (3.28), it follows that

$$
\||\widehat{\Sigma}^{(j-2)}|\|_{2\mathfrak{q}} \le \||\widehat{\Sigma}^{(j-2)} - \Sigma^{(j-2)}|\|_{2\mathfrak{q}} + |\Sigma^{(j-2)}| \le c\,(1 \vee \Omega)\,,
$$

and also that

$$
\left\|\sqrt{\frac{r_2(\widehat{\Sigma}^{(j-2)})}{N}}\right\|_{2\mathfrak{q}}^{2\mathfrak{q}} = \mathbb{E}\left[\left(\frac{r_2(\widehat{\Sigma}^{(j-2)})}{N}\right)^{\mathfrak{q}}\right] \lesssim \mathbb{E}\left[\left(\frac{\mathrm{Tr}(\widehat{\Sigma}^{(j-2)})}{N}\right)^{\mathfrak{q}}\right] \le cN^{-\mathfrak{q}} r_2(\Sigma^{(0)})^{\mathfrak{q}}.
$$

Using identical arguments to those used to control $\||\mathrm{Tr}(\hat{O}_1^{(j-1)})|\|_q$, we have that

$$
\|\mathrm{Tr}(\hat{O}_2^{(j-1)})\|_q \le c r_2(\Sigma^{(0)}).
$$

Controlling $\|\mathbf{Tr}(\hat{O}_3^{(j-1)})\|_q$ Note that $\|\mathrm{Tr}(\hat{O}_2^{(j-1)})\|_q = \|\mathrm{Tr}(\hat{O}_2^{(j-1)})\|_q$.

Forecast Covariance Bound

Let $\mathfrak{q} \in \{q, 2q, 4q\}$. By the triangle inequality, the inductive hypothesis (3.28), and Theorem 3.7.5, we have

$$
\begin{aligned}
|||\widehat{C}^{(j)} - C^{(j)}|||_{\mathfrak{q}} &\le |A|^2 |||S^{(j-1)} - \Sigma^{(j-1)}|||_{\mathfrak{q}} + |||\hat{\Xi}^{(j)} - \Xi|||_{\mathfrak{q}} + 2|A||||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}} \\
&\le |A|^2 \left(|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}} + |||\widehat{\Sigma}^{(j-1)} - \Sigma^{(j-1)}|||_{\mathfrak{q}} \right) + |||\hat{\Xi}^{(j)} - \Xi|||_{\mathfrak{q}} + 2|A||||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}} \\
&\le c|A|^2 \left(|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}} + \Omega \right) + |\Xi| \left(1 \vee \sqrt{\frac{r_2(\Xi)}{N}} \right) + 2|A||||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}}.
\end{aligned}
$$

The forecast covariance bound (3.26) is then a direct consequence of the bounds that we now establish on $|||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}}$ and $|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}}$.

Controlling $|||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}}$ By an identical analysis to the one used in bounding $|||\widehat{C}^{(j)}_{u\eta}|||_{\mathfrak{q}}$ in (3.29), we have that

$$
|||\widehat{C}^{(j)}_{u\xi}|||_{\mathfrak{q}} \lesssim c\Omega. \tag{3.30}
$$

Controlling $|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}}$ By iterated expectations and Theorem 3.7.5, we have

$$
\begin{aligned}
|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}}^{\mathfrak{q}} &= \mathbb{E}\left[|S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|^{\mathfrak{q}} \right] = \mathbb{E}\left[\mathbb{E}\left[|S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|^{\mathfrak{q}} \middle| \widehat{\mu}^{(j-1)}, \widehat{\Sigma}^{(j-1)} \right] \right] \\
&\lesssim_q \mathbb{E}\left[|\widehat{\Sigma}^{(j-1)}|^{\mathfrak{q}} \left(\frac{r_2(\widehat{\Sigma}^{(j-1)})}{N} \right)^{\mathfrak{q}/2} \right] = \mathbb{E}\left[|\widehat{\Sigma}^{(j-1)}|^{\mathfrak{q}/2} \left(\frac{\operatorname{Tr}(\widehat{\Sigma}^{(j-1)})}{N} \right)^{\mathfrak{q}/2} \right] \\
&\le \sqrt{\mathbb{E}|\widehat{\Sigma}^{(j-1)}|^{\mathfrak{q}}} \sqrt{\mathbb{E}\left[\frac{\operatorname{Tr}(\widehat{\Sigma}^{(j-1)})}{N} \right]^{\mathfrak{q}}}.
\end{aligned}
$$

By the covariance trace bound proved in Subsection 3.5.3 and the inductive hypothesis (3.28), we then have that $|||S^{(j-1)} - \widehat{\Sigma}^{(j-1)}|||_{\mathfrak{q}} \lesssim c\Omega$.

Mean Bound

By Lemma 3.7.2, we have

$$
\begin{aligned}
|||\widehat{\mu}^{(j)} - \mu^{(j)}|||_q &= |||\mathscr{M}(\widehat{m}^{(j)}, \widehat{C}^{(j)}; y^{(j)}) - \mathscr{M}(m^{(j)}, C^{(j)}; y^{(j)})|||_q + |||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q \\
&\le \Big\| |\widehat{m}^{(j)} - m^{(j)}| \Big\|_q + |H|^2|\Gamma^{-1}||||\widehat{C}^{(j)}|||_{2q}|||\widehat{m}^{(j)} - m^{(j)}|||_{2q} \\
&+ |||\widehat{C}^{(j)} - C^{(j)}|||_q|H||\Gamma^{-1}|\big(1 + |H|^2|\Gamma^{-1}||C^{(j)}|\big)|y^{(j)} - Hm^{(j)}| \\
&+ |||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q.
\end{aligned}
$$

The induction step for the mean bound (3.17) is then a direct consequence of the bounds that we now establish on $|||\widehat{m}^{(j)} - m^{(j)}|||_{\mathfrak{q}}$ for $\mathfrak{q} \in \{q, 2q\}$ and on $|||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q$.

Controlling $|||\widehat{m}^{(j)} - m^{(j)}|||_{\mathfrak{q}}$ for $\mathfrak{q} \in \{q, 2q\}$ It follows by the triangle inequality and the inductive hypothesis (3.28) that

$$
\begin{aligned}
|||\widehat{m}^{(j)} - m^{(j)}|||_{\mathfrak{q}} &\le |A||||\bar{u}^{(j-1)} - \mu^{(j-1)}|||_{\mathfrak{q}} + |||\bar{\xi}^{(j)}|||_{\mathfrak{q}} \\
&\le |A| \left(|||\bar{u}^{(j-1)} - \widehat{\mu}^{(j-1)}|||_{\mathfrak{q}} + |||\widehat{\mu}^{(j-1)} - \mu^{(j-1)}|||_{\mathfrak{q}} \right) + |||\bar{\xi}^{(j)}|||_{\mathfrak{q}} \\
&\le c'|A| \left(|||\bar{u}^{(j-1)} - \widehat{\mu}^{(j-1)}|||_{\mathfrak{q}} + \Omega \right) + c(q)|\Xi|\sqrt{\frac{r_2(\Xi)}{N}}.
\end{aligned}
$$

By iterated expectations, Lemma 3.7.9, and the covariance trace bound proved in Subsection 3.5.3, it follows that

$$
\begin{aligned}
|||\bar{u}^{(j-1)} - \widehat{\mu}^{(j-1)}|||_{\mathfrak{q}}^{\mathfrak{q}} &= \mathbb{E}\big[|\bar{u}^{(j-1)} - \widehat{\mu}^{(j-1)}|^{\mathfrak{q}}\big] = \mathbb{E}\left[\mathbb{E}\left[|\bar{u}^{(j-1)} - \widehat{\mu}^{(j-1)}|^{\mathfrak{q}} \,\middle|\, \widehat{\mu}^{(j-1)}, \widehat{\Sigma}^{(j-1)}\right]\right] \\
&\lesssim_q \mathbb{E}\left[\left(\frac{\operatorname{Tr}(\widehat{\Sigma}^{(j-1)})}{N}\right)^{\mathfrak{q}/2}\right] \le c \left(\frac{r_2(\Sigma^{(0)})}{N}\right)^{\mathfrak{q}/2},
\end{aligned}
$$

and so $|||\widehat{m}^{(j)} - m^{(j)}|||_{\mathfrak{q}} \lesssim c'\Omega$.

Controlling $|||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q$ Note first that $|||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q \le |||\mathscr{K}(\widehat{C}^{(j)})|||_{2q}|||\bar{\eta}^{(j)}|||_{2q}$. We then have by Lemma 3.7.9, $|||\bar{\eta}^{(j)}|||_{2q} \lesssim_q \sqrt{\frac{\mathrm{Tr}(\Gamma)}{N}} = \sqrt{|\Gamma|\frac{r_2(\Gamma)}{N}}$, and by Lemma 3.7.1 and the forecast covariance bound established in Subsection 3.5.3, we have

$$
\begin{aligned}
|||\mathscr{K}(\widehat{C}^{(j)})|||_{2q} \le |H||\Gamma^{-1}||||\widehat{C}^{(j)}|||_{2q} &\le |H||\Gamma^{-1}|\left(|||\widehat{C}^{(j)} - C^{(j)}|||_{2q} + |C^{(j)}|\right) \\
&\le |H||\Gamma^{-1}|c\,(1 \vee \Omega)\,.
\end{aligned}
$$

Therefore, $|||\mathscr{K}(\widehat{C}^{(j)})\bar{\eta}^{(j)}|||_q \le c\Omega$.

Covariance Bound

By Lemma 3.7.3, we have

$$
\begin{aligned}
|||\widehat{\Sigma}^{(j)} - \Sigma^{(j)}|||_q &\le |||\mathscr{C}(\widehat{C}^{(j)}) - \mathscr{C}(C^{(j)})|||_q + |||\hat{O}^{(j)}|||_q \\
&\le |||\widehat{C}^{(j)} - C^{(j)}|||_q(1 + |A|^2|\Gamma^{-1}||C^{(j)}|) \\
&+ (|A|^2|\Gamma^{-1}| + |A|^4|\Gamma^{-1}|^2|C^{(j)}|)|||\widehat{C}^{(j)}|||_{2q}|||\widehat{C}^{(j)} - C^{(j)}|||_{2q} + |||\hat{O}^{(j)}|||_q.
\end{aligned}
$$

The induction step for the forecast covariance has been proved in Subsection 3.5.3, and so in order to show the induction step for the covariance bound we only need to control the offset term. First, using the triangle inequality, we write

$$
\begin{aligned}
|||\widehat{O}^{(j)}|||_q \le |||\mathscr{K}(\widehat{C}^{(j)})(\widehat{\Gamma}^{(j)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(j)})|||_q + |||\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)\widehat{C}^{(j)}_{u\eta}\mathscr{K}^\top(\widehat{C}^{(j)})|||_q \\
+ |||\mathscr{K}(\widehat{C}^{(j)})(\widehat{C}^{(j)}_{u\eta})^\top\big(I - H^\top\mathscr{K}^\top(\widehat{C}^{(j)})\big)|||_q = \|\hat{O}^{(j)}_1\|_q + \|\hat{O}^{(j)}_2\|_q + \|\hat{O}^{(j)}_3\|_q.
\end{aligned}
$$

We next bound each term in turn.

Controlling $\|\hat{O}_1^{(j)}\|_q$ By Lemma 3.7.1, the forecast covariance bound established in Subsection 3.5.3, and Theorem 3.7.5, it holds that

$$
\begin{aligned}
\|\hat{O}_1^{(j)}\|_q &= \||\mathscr{K}(\widehat{C}^{(j)})(\widehat{\Gamma}^{(j)} - \Gamma)\mathscr{K}^\top(\widehat{C}^{(j)})|\|_q \\
&\le \||\mathscr{K}(\widehat{C}^{(j)})|\|_{4q}^2 \||\widehat{\Gamma}^{(j)} - \Gamma|\|_{2q} \le |H|^2|\Gamma^{-1}|^2 \||\widehat{C}^{(j)}|\|_{4q}^2 \||\widehat{\Gamma}^{(j)} - \Gamma|\|_{2q} \\
&\le c\,(1 \vee \Omega)^2 \sqrt{\frac{r_2(\Gamma)}{N}} \le c\sqrt{\frac{r_2(\Gamma)}{N}},
\end{aligned}
$$

where the last inequality uses that by assumption $N \ge r_2(\Sigma^{(0)}) \vee r_2(\Xi) \vee r_2(\Gamma)$.

Controlling $\|\hat{O}_2^{(j)}\|_q$ By Lemma 3.7.1, inequality (3.30), and the forecast covariance bound established in Subsection 3.5.3, we get

$$
\begin{aligned}
\|\hat{O}_2^{(j)}\|_q &= \||\big(I - \mathscr{K}(\widehat{C}^{(j)})H\big)\widehat{C}_{u\eta}^{(j)}\mathscr{K}^\top(\widehat{C}^{(j)})|\|_q \le \||\mathscr{K}(\widehat{C}^{(j)})||I - \mathscr{K}(\widehat{C}^{(j)})H||\widehat{C}_{u\eta}^{(j)}|\|_q \\
&\le \||\mathscr{K}(\widehat{C}^{(j)})|(1 + |\mathscr{K}(\widehat{C}^{(j)})||H|)|\widehat{C}_{u\eta}^{(j)}|\|_q \\
&\le \||\widehat{C}_{u\eta}^{(j)}|\|_{2q} \left(|H||\Gamma^{-1}|\||\widehat{C}^{(j)}|\|_{2q} + |H|^3|\Gamma^{-1}|^2\||\widehat{C}^{(j)}|^2\|_{2q}\right) \le c\Omega.
\end{aligned}
$$

Controlling $\|\hat{O}_3^{(j)}\|_q$ Note that $\|\hat{O}_3^{(j)}\|_q = \|\hat{O}_2^{(j)}\|_q$.

3.6 Conclusions

This chapter has investigated REnKF, a modification of EnKF with improved theoretical guarantees. Theorem 3.3.2 gives non-asymptotic error bounds for a stochastic EnKF over multiple assimilation cycles. Numerical experiments demonstrate that the benefits of introducing resampling for theory purposes do not come at the price of a deterioration in state estimation or uncertainty quantification tasks.

Resampling techniques for ensemble Kalman algorithms deserve further research. From a theory viewpoint, resampling offers a promising path to develop long-time filter accuracy theory, blending our inductive analysis with existing results that ensure long-time stability of the filtering distributions Sanz-Alonso and Stuart [2015]. From a methodological viewpoint, other resampling schemes can be considered Naesseth et al. [2018]. Finally, while our numerical investigation has focused on settings where the standard EnKF algorithm is effective, an important open problem is to identify dynamical systems and/or observation models for which resampling may offer an empirical advantage.

3.7 Additional Results

3.7.1 Metrics for Numerical Results

In this appendix, we give a more extensive description of the Monte Carlo procedure utilized to calculate the metrics referred to in Section 3.4. We summarize the approach in Algorithm 0. We require the following additional notation: We write $\text{diag}(A) = (A_{11}, A_{22}, \dots, A_{dd})^\top$. For a function $g : \mathbb{R} \to \mathbb{R}$, $g(u) = \big(g(u(1)), \dots, g(u(d))\big)^\top$ is the element-wise application of g to u.

3.7.2 Technical Results

Additional Notation

Given a non-decreasing, non-zero convex function $\psi : [0, \infty] \to [0, \infty]$ with $\psi(0) = 0$, the Orlicz norm of a real random variable X is $\|X\|_\psi = \inf\{t > 0 : \mathbb{E}[\psi(t^{-1}|X|)] \leq 1\}$. In particular, for the choice $\psi_p(x) = e^{x^p} - 1$ for $p \geq 1$, real random variables that satisfy $\|X\|_{\psi_2} < \infty$ are referred to as sub-Gaussian, and those that satisfy $\|X\|_{\psi_1} < \infty$ are sub-Exponential. The random vector Y is sub-Gaussian (sub-Exponential) if $\|v^\top Y\|_{\psi_2} < \infty$ ($\|v^\top Y\|_{\psi_1} < \infty$) for any vector v satisfying $|v|_2 = 1$.

Algorithm 3 Metrics Calculation for Numerical Results

1: **Fixed Quantities**: Ground-truth state $\{u^{(j)}\}_{j=0}^{J}$, observations $\{y^{(j)}\}_{j=1}^{J}$, and Kalman filter means $\{\mu^{(j)}\}_{j=1}^{J}$.

2: **Monte Carlo Trials**: For $m = 1, 2, \ldots, M$ run algorithm $\mathsf{A} \in \{\mathsf{EnKF}(1), \mathsf{REnKF}(2)\}$ and obtain $\{\widehat{\mu}_m^{(j),\mathsf{A}}, \widehat{\Sigma}_m^{(j),\mathsf{A}}\}_{j,m}^{J,M}$.

3: *Mean Error:*

$$\mathsf{E}^{\mathsf{A}}_{m,\text{Linear}} = \frac{1}{J}\sum_{j=1}^{J} |\widehat{\mu}_m^{(j),\mathsf{A}} - \mu^{(j)}|_2, \qquad \mathsf{E}^{\mathsf{A}}_{m,\text{L96}} = \frac{1}{J}\sum_{j=1}^{J} |\widehat{\mu}_m^{(j),\mathsf{A}} - u^{(j)}|_2. \tag{3.31}$$

Confidence Interval: Let $\hat{\sigma}_m^{(j),\mathsf{A}} = \sqrt{\text{diag}(\widehat{\Sigma}_m^{(j),\mathsf{A}})}$, then compute

$$\begin{aligned}
\mathsf{I}_m^{(j),\mathsf{A}} &= \widehat{\mu}_m^{(j),\mathsf{A}} \pm 1.96 \times \hat{\sigma}_m^{(j),\mathsf{A}}, && \text{(Interval)}\\
\mathsf{W}_m^{\mathsf{A}} &= \frac{2 \times 1.96}{dJ}\sum_{j=1}^{J} |\hat{\sigma}_m^{(j),\mathsf{A}}|_1, && \text{(Average Width)}\\
\mathsf{V}_m^{\mathsf{A}} &= \frac{1}{dJ}\sum_{j=1}^{J}\sum_{i=1}^{d} \mathbf{1}\{u^{(j)}(i) \in \mathsf{I}_m^{(j),\mathsf{A}}(i)\}. && \text{(Average Coverage)}
\end{aligned} \tag{3.32}$$

4: **Output**:

$$\begin{aligned}
\mathsf{E}^{\mathsf{A}}_{\text{Linear}} &= \frac{1}{M}\sum_{m=1}^{M} \mathsf{E}^{\mathsf{A}}_{m,\text{Linear}}, \qquad \mathsf{E}^{\mathsf{A}}_{\text{L96}} = \frac{1}{M}\sum_{m=1}^{M} \mathsf{E}^{\mathsf{A}}_{m,\text{L96}},\\
\mathsf{W}^{\mathsf{A}} &= \frac{1}{M}\sum_{m=1}^{M} \mathsf{W}^{\mathsf{A}}_m, \qquad \mathsf{V}^{\mathsf{A}} = \frac{1}{M}\sum_{m=1}^{M} \mathsf{V}^{\mathsf{A}}_m.
\end{aligned} \tag{3.33}$$

Lemma 3.7.1 (Properties of the Kalman Gain Operator [Kwiatkowski and Mandel, 2015, Lemma 4.1 & Corollary 4.2]). *Let $\mathscr{K}$ be the Kalman gain operator defined in* (3.12). *Let $P, Q \in \mathcal{S}_+^d$, $\Gamma \in \mathcal{S}_{++}^k$, and $H \in \mathbb{R}^{k\times d}$. The following hold:*

$$\begin{aligned}
|\mathscr{K}(Q) - \mathscr{K}(P)| &\le |Q-P||H||\Gamma^{-1}|\Big(1 + \min\left(|P|, |Q|\right)|H|^2|\Gamma^{-1}|\Big),\\
|\mathscr{K}(Q)| &\le |Q||H||\Gamma^{-1}|,\\
|I - \mathscr{K}(Q)H| &\le 1 + |Q||H|^2|\Gamma^{-1}|.
\end{aligned}$$

Lemma 3.7.2 (Properties of the Mean-Update Operator [Kwiatkowski and Mandel, 2015, Lemma 4.10]). *Let $\mathscr{M}$ be the mean-update operator defined in* (3.13). *Let $m \in \mathbb{R}^d$ be a random vector and Q be a random matrix such that $Q \in \mathcal{S}_+^d$ almost surely. Let $P \in \mathcal{S}_+^d$, $\Gamma \in \mathcal{S}_{++}^k$, $H \in \mathbb{R}^{k\times d}$, $y \in \mathbb{R}^k$, and $m' \in \mathbb{R}^d$ be deterministic. Then, for any $1 \le q < \infty$ and $y \in \mathbb{R}^k$ it holds that*

$$\begin{aligned}
\||\mathscr{M}(m,Q;y) - \mathscr{M}(m',P;y)|\|_q \le \big\||m-m'|\big\|_q &+ |H|^2|\Gamma^{-1}|\||Q|\|_{2q}\big\||m-m'|\big\|_{2q}\\
&+ \||Q-P|\|_q|H||\Gamma^{-1}|\big(1+|H|^2|\Gamma^{-1}||P|\big)|y-Hm'|.
\end{aligned}$$

Lemma 3.7.3 (Properties of the Covariance-Update Operator [Kwiatkowski and Mandel, 2015, Lemmas 4.6 & 4.8]). *Let $\mathscr{C}$ be the covariance-update operator defined in* (3.14). *Let P, Q, m, m', y, H and Γ all be defined as in Lemma 3.7.2. Then, for any $1 \le q < \infty$, it holds that*

$$\begin{aligned}
0 \preceq \mathscr{C}(Q) \preceq Q, &\qquad |\mathscr{C}(Q)| \le |Q|,\\
\||\mathscr{C}(Q) - \mathscr{C}(P)|\|_q \le \||Q-P|\|_q&(1+|H|^2|\Gamma^{-1}||P|)\\
&+ (|H|^2|\Gamma^{-1}| + |H|^4|\Gamma^{-1}|^2|P|)\||Q|\|_{2q}\||Q-P|\|_{2q}.
\end{aligned}$$

Theorem 3.7.4 (Gaussian Norm Concentration, [Vershynin, 2018, Exercise 6.3.5]). *Let*

$X \in \mathbb{R}^d$ *be a Gaussian random vector with* $\mathbb{E}[X] = \mu^X$, *var*$[X] = \Sigma^X$. *Then, for any* $t \geq 1$, *with probability at least* $1 - ce^{-t}$ *it holds that*

$$|X - \mu^X|_2 \lesssim \sqrt{Tr(\Sigma^X)} + \sqrt{t|\Sigma^X|} \lesssim \sqrt{|\Sigma^X|(r_2(\Sigma^X) \vee t)}\,.$$

Theorem 3.7.5 (Covariance Bound, [Koltchinskii and Lounici, 2017, Corollary 2])**.** *Let* $X_1, \ldots, X_n$ *be i.i.d. copies of a d-dimensional Gaussian vector* X *with* $\mathbb{E}[X] = 0$ *and var*$[X] = \Sigma$. *Let* $\widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^n X_i X_i^\top$ *be the sample covariance estimator. For any* $q \geq 1$, *it holds that*

$$|||\widehat{\Sigma} - \Sigma|||_q \lesssim_q |\Sigma| \left(\sqrt{\frac{r_2(\Sigma)}{n}} \vee \frac{r_2(\Sigma)}{n} \right).$$

Lemma 3.7.6. *Let* $A, B \in \mathcal{S}_+^d$. *It holds that*

$$Tr(AB) \leq |A|\,Tr(B).$$

Lemma 3.7.7 (Cross-Covariance Estimation —Unstructured Case)**.** *Let* $u_1, \ldots, u_N \in \mathbb{R}^d$ *be i.i.d. Gaussian random vectors with* $\mathbb{E}[u_1] = m$ *and var*$[u_1] = C$. *Let* $\eta_1, \ldots, \eta_N \in \mathbb{R}^k$ *be i.i.d. Gaussian random vectors with* $\mathbb{E}[\eta_1] = 0$ *and var*$[\eta_1] = \Gamma$, *and assume that the two sequences are independent. Let*

$$\widehat{C}_{u\eta} = \frac{1}{N-1}\sum_{n=1}^{N}(u_n - \widehat{m})(\eta_n - \bar{\eta})^\top,$$

and assume that $N \geq r_2(C) \vee r_2(\Gamma)$. *Then,*

$$|||\widehat{C}_{u\eta}|||_q \lesssim_q (|C| \vee |\Gamma|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \right).$$

Proof. By [Al-Ghattas and Sanz-Alonso, 2024b, Lemma A.3], there exists a constant c such that, for all $t \geq 1$, it holds with probability at least $1 - ce^{-t}$ that

$$|\widehat{C}_{u\eta}| \lesssim (|C| \vee |\Gamma|) \left(\sqrt{\frac{r_2(C)}{N}} \vee \sqrt{\frac{r_2(\Gamma)}{N}} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N} \right).$$

Integrating the tail bound then yields the result. □

Lemma 3.7.8. *Let $X_1, \ldots, X_n$ be i.i.d. copies of a d-dimensional Gaussian vector X with $\mathbb{E}[X] = 0$ and $var[X] = \Sigma$. Let $\widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^n X_i X_i^\top$ be the sample covariance estimator. Then, for any $\delta \geq 1$, it holds with probability at least $1 - 2e^{-\delta}$ that*

$$|\mathit{Tr}(\widehat{\Sigma}) - \mathit{Tr}(\Sigma)| \leq c\,\mathit{Tr}(\Sigma) \left(\sqrt{\frac{\delta}{n}} \vee \frac{\delta}{n} \right).$$

Further, for any $q \geq 1$,

$$\| \, |\mathit{Tr}(\widehat{\Sigma}) - \mathit{Tr}(\Sigma)| \, \|_q \lesssim_q \frac{\mathit{Tr}(\Sigma)}{\sqrt{n}}.$$

Proof. Let $Z_{ij} = \Sigma_{jj}^{-1/2} X_{ij}$ and note that, for any $t > 0$,

$$\begin{aligned} \mathbb{P}(|\mathrm{Tr}(\widehat{\Sigma}) - \mathrm{Tr}(\Sigma)| > t) = \mathbb{P}(|\mathrm{Tr}(\widehat{\Sigma} - \Sigma)| > t) &= \mathbb{P}\left(\left| \sum_{i=1}^n \left(\sum_{j=1}^d (X_{ij}^2 - \mathbb{E}X_{ij}^2) \right) \right| > nt \right) \\ &= \mathbb{P}\left(\left| \sum_{i=1}^n \left(\sum_{j=1}^d \Sigma_{jj} (Z_{ij}^2 - \mathbb{E}Z_{ij}^2) \right) \right| > nt \right). \end{aligned}$$

Note that the random variables $\sum_{j=1}^d \Sigma_{jj}(Z_{ij}^2 - \mathbb{E}Z_{ij}^2)$ for $i = 1, \ldots, n$ are independent, mean-zero and sub-exponential with ψ_1 norm at most $C\mathrm{Tr}(\Sigma)$, since

$$\begin{aligned}\left\|\sum_{j=1}^{d}\Sigma_{jj}(Z_{ij}^2-\mathbb{E}Z_{ij}^2)\right\|_{\psi_1} &\leq \sum_{j=1}^{d}\Sigma_{jj}\left\|Z_{ij}^2-\mathbb{E}Z_{ij}^2\right\|_{\psi_1}\leq C\sum_{j=1}^{d}\Sigma_{jj}\left\|Z_{ij}^2\right\|_{\psi_1}\\ &= C\sum_{j=1}^{d}\Sigma_{jj}\left\|Z_{ij}\right\|_{\psi_2}^2\leq C\sum_{j=1}^{d}\Sigma_{jj}=C\mathrm{Tr}(\Sigma).\end{aligned}$$

The second inequality holds due to the Centering Lemma, [Vershynin, 2018, Lemma 2.6.8]. Therefore, by Bernstein's inequality we have

$$\mathbb{P}\left(\left|\sum_{i=1}^{n}\left(\sum_{j=1}^{d}\Sigma_{jj}(Z_{ij}^2-\mathbb{E}Z_{ij}^2)\right)\right|>nt\right)\leq 2\exp\left(-c\min\left(\frac{nt^2}{(\mathrm{Tr}(\Sigma))^2},\frac{nt}{\mathrm{Tr}(\Sigma)}\right)\right).$$

For the expectation bound, we note that

$$\begin{aligned}\||\mathrm{Tr}(\widehat{\Sigma})-\mathrm{Tr}(\Sigma)|\|_q^q &= \int_0^\infty \mathbb{P}(|\mathrm{Tr}(\widehat{\Sigma})-\mathrm{Tr}(\Sigma)|^q>t)\,dt\\ &\leq \zeta^q+q\int_C^\infty t^{q-1}\mathbb{P}(|\mathrm{Tr}(\widehat{\Sigma})-\mathrm{Tr}(\Sigma)|>t)\,dt\\ &\leq \zeta^q+2q\int_0^\infty t^{q-1}\exp\left(-c\min\left(\frac{nt^2}{(\mathrm{Tr}(\Sigma))^2},\frac{nt}{\mathrm{Tr}(\Sigma)}\right)\right)dt\\ &= \zeta^q+2qc\max\left(\frac{\Gamma(q/2)(\mathrm{Tr}(\Sigma))^q}{n^{q/2}},\frac{\Gamma(q)(\mathrm{Tr}(\Sigma))^q}{n^q}\right).\end{aligned}$$

Taking $\zeta=\mathrm{Tr}(\Sigma)/n$, it then follows that

$$\||\mathrm{Tr}(\widehat{\Sigma})-\mathrm{Tr}(\Sigma)|\|_q\lesssim\zeta+c\mathrm{Tr}(\Sigma)\max\left(\frac{1}{\sqrt{n}},\frac{1}{n}\right)\lesssim\frac{\mathrm{Tr}(\Sigma)}{\sqrt{n}}.$$

□

Lemma 3.7.9. *Let $X_1,\ldots,X_n$ be i.i.d. copies of a d-dimensional Gaussian vector X with*

$\mathbb{E}[X] = \mu$ *and* $var[X] = \Sigma$*. Let* $\overline{X} = \frac{1}{N}\sum_{n=1}^{N} X_n$*. Then, for any* $q \geq 1$*,*

$$|||\bar{X} - \mu|||_q \lesssim_q \sqrt{\frac{Tr(\Sigma)}{N}}.$$

Proof. Let $c_2 := c_1\sqrt{\mathrm{Tr}(\Sigma)/N}$ where c_1 is a sufficiently large positive constant, then

$$\begin{aligned}
\mathbb{E}\left[|\bar{X} - \mu|^q\right] &= \int_0^\infty \mathbb{P}(|\bar{X} - \mu|^q > y)\,dy \leq c_2^q + \int_{c_2}^\infty \mathbb{P}(|\bar{X} - \mu|^q > y)\,dy \\
&= c_2^q + \int_{c_2}^\infty q y^{q-1}\mathbb{P}(|\bar{X} - \mu| > y)\,dy \\
&= c_2^q + \int_{c_2 - c\mathrm{Tr}(\Sigma/N)}^\infty q\left(c\sqrt{\frac{\mathrm{Tr}(\Sigma)}{N}} + t\right)^{q-1}\mathbb{P}\left(|\bar{X} - \mu| > c\sqrt{\frac{\mathrm{Tr}(\Sigma)}{N}} + t\right)dt
\end{aligned}$$

where the last equality holds by a change of variable. By Theorem 3.7.4 it follows that $\mathbb{P}(|\overline{X} - \mu| \geq c\sqrt{\mathrm{Tr}(\Sigma)} + t) \leq \exp(-ct^2/|\Sigma|)$, and so the expression in the above display is bounded above by

$$\begin{aligned}
&c_2^q + \int_{c_2 - c\mathrm{Tr}(\Sigma/N)}^\infty q\left(c\sqrt{\frac{\mathrm{Tr}(\Sigma)}{N}} + t\right)^{q-1}\exp\left(-\frac{c_2 n t^2}{|\Sigma|}\right)dt \\
&\lesssim c_2^q + \int_0^\infty q\left(\left(\frac{c\mathrm{Tr}(\Sigma)}{N}\right)^{(q-1)/2} + t^{q-1}\right)\exp\left(-\frac{c_2 n t^2}{|\Sigma|}\right)dt \\
&= c_2^q + q\left(\frac{1}{2}\Gamma(q/2)\left(\frac{|\Sigma|}{N}\right)^{q/2} + \frac{1}{2}\left(\frac{c\mathrm{Tr}(\Sigma)}{N}\right)^{(q-1)/2}\sqrt{\frac{\pi|\Sigma|}{N}}\right) \\
&\lesssim c_2^q + q\left(\frac{1}{2}\Gamma(q/2)\left(\frac{|\Sigma|}{N}\right)^{q/2} + \frac{1}{2}\left(\frac{c\mathrm{Tr}(\Sigma)}{N}\right)^{q/2}\right) \\
&\lesssim \left(\frac{\mathrm{Tr}(\Sigma)}{N}\right)^{q/2}.
\end{aligned}$$

Therefore,

$$|||\bar{X} - \mu||| \lesssim_q c_2 + \sqrt{\frac{|\Sigma|}{N}} + c\sqrt{\frac{\mathrm{Tr}(\Sigma)}{N}} \lesssim \sqrt{\frac{\mathrm{Tr}(\Sigma)}{N}},$$

where the last inequality holds since $\mathrm{Tr}(\Sigma) \geq |\Sigma|$ and the choice of c_2. $\square$

CHAPTER 4

COVARIANCE OPERATOR ESTIMATION: SPARSITY, LENGTHSCALE, AND ENSEMBLE KALMAN FILTERS

This chapter is adapted from the publication listed below and is used with permission of the publisher.

O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek, *Covariance operator estimation: sparsity, lengthscale, and ensemble Kalman filters*, *Bernoulli, 31(3), 2377-2402*, 2025

4.1 Introduction

This chapter studies thresholded estimation of the covariance operator of a Gaussian random field. Under a sparsity assumption on the covariance model, we bound the estimation error in terms of the sparsity level and the expected supremum of the field. Using this bound, we then analyze covariance operator estimation in the interesting regime where the correlation lengthscale is small, and show that the thresholded covariance estimator achieves an exponential improvement in sample complexity compared with the standard sample covariance estimator. As an application of the theory, we demonstrate the advantage of using thresholded covariance estimators within ensemble Kalman filters.

The first contribution of this chapter is to lift the theory of covariance estimation from finite to infinite dimension. In the finite-dimensional setting, a rich body of work Wu and Pourahmadi [2003], Bickel and Levina [2008b], El Karoui [2008], Cai and Yuan [2012], Cai and Zhou [2012a,b], Chen et al. [2012], Wainwright [2019], Al-Ghattas and Sanz-Alonso [2024b] shows that, exploiting various forms of sparsity, it is possible to consistently estimate the covariance matrix of a vector $u \in \mathbb{R}^{d_u}$ with $N \sim \log(d_u)$ samples. The sparsity of the covariance matrix —along with the use of thresholded, tapered, or banded estimators that exploit this structure— facilitates an exponential improvement in sample complexity relative to the unstructured case, where $N \sim d_u$ samples are needed Bai and Yin [2008], Gordon [1985], Vershynin [2010]. In this work we investigate the setting in which u is an infinite-dimensional random field with an approximately sparse covariance model. Specifically, we generalize notions of approximate sparsity often employed in the finite-dimensional covariance estimation literature Bickel and Levina [2008a], Cai and Zhou [2012b]. We show that the statistical error of thresholded estimators can be bounded in terms of the expected supremum of the field and the sparsity level, the latter of which quantifies the rate of spatial decay of correlations of the random field. Our analysis not only lifts existing theory from finite to infinite dimension, but also provides non-asymptotic moment bounds not yet available in finite dimension.

The second contribution of this chapter is to showcase the benefit of thresholding in the challenging regime where the correlation lengthscale of the field is small relative to the size of the physical domain. While a vast literature in nonparametric statistics Ghosal and van der Vaart [2017] and approximation theory Wendland [2004] highlights the key role of smoothness in determining optimal convergence rates for many nonparametric estimation tasks, our non-asymptotic theory emphasizes that the lengthscale rather than the smoothness of the covariance function drives the difficulty of the estimation problem and the advantage of thresholded estimators.

Fields with small correlation lengthscale are ubiquitous in applications. For instance, they arise naturally in climate science and numerical weather forecasting, where global forecasts need to account for the effect of local processes with a small correlation lengthscale, such as cloud formation or propagation of gravitational waves. We show that thresholded estimators achieve an exponential improvement in sample complexity: For a field with lengthscale λ in d-dimensional physical space, the standard sample covariance requires $N \sim \lambda^{-d}$ samples, while thresholded estimators only require $N \sim \log(\lambda^{-d})$. Therefore, our theory suggests that the parameter λ^{-d} plays the same role in infinite dimension as d_u in the classical finite-dimensional setting. To analyze thresholded estimators in the small lengthscale regime, we use our general non-asymptotic moment bounds and the sharp scaling of sparsity level and expected supremum with lengthscale.

The third contribution of this chapter is to demonstrate the advantage of using thresholded covariance estimators within ensemble Kalman filters Evensen [2009]. Our interest in covariance operator estimation was motivated by the widespread use of localization techniques within ensemble Kalman methods in inverse problems and data assimilation, see e.g. Houtekamer and Mitchell [2001], Houtekamer and Zhang [2016], Farchi and Bocquet [2019], Tong and Morzfeld [2023], Chen et al. [2022]. Many inverse problems in medical imaging and the geophysical sciences are most naturally formulated in function space Stuart [2010], Bui-Thanh et al. [2013], Bigoni et al. [2020]; likewise, data assimilation is primarily concerned with sequential estimation of spatial fields, e.g. temperature or precipitation Kalnay [2003], Carrassi et al. [2018]. Theoretical insight for these applications calls for sparse covariance estimation theory in function space, which has not been the focus in the literature. Perhaps partly for this reason, the empirical success of localization techniques in ensemble Kalman methods is poorly understood, with few exceptions that study localization in finite dimension Tong [2018], Al-Ghattas and Sanz-Alonso [2024b]. The work Sanz-Alonso and Waniorek [2024] studies the behavior of ensemble Kalman methods under mesh discretization, but it does not consider localization. In this chapter, we use our novel non-asymptotic covariance estimation theory to obtain a sufficient sample size to approximate an idealized mean-field ensemble Kalman filter using a localized ensemble Kalman update. In finite dimension, Al-Ghattas and Sanz-Alonso [2024b] studies the ensemble approximation of mean-field algorithms for inverse problems and Al-Ghattas et al. [2024a] conducts a multi-step analysis of ensemble Kalman filters without localization.

The chapter is organized as follows. We first state and discuss our three main theorems in the following section. Then, the next three sections contain the proof of these theorems, along with further auxiliary results of independent interest. We close with conclusions, discussion, and future directions.

Notation Given two positive sequences $\{a_n\}$ and $\{b_n\}$, the relation $a_n \lesssim b_n$ denotes that $a_n \le cb_n$ for some constant $c > 0$. If the constant c depends on some quantity τ, then we write $a \lesssim_\tau b$. If both $a_n \lesssim b_n$ and $b_n \lesssim a_n$ hold simultaneously, then we write $a_n \asymp b_n$. For a finite-dimensional vector a, $|a|$ denotes its Euclidean norm. For an operator $\mathcal{A}$, $\|\mathcal{A}\|$ denotes its operator norm, $\mathcal{A}^*$ its adjoint, and $\mathrm{Tr}(\mathcal{A})$ its trace.

4.2 Main Results

This section states and discusses the main results of the chapter. In Subsection 4.2.1 we analyze the thresholded sample covariance estimator in a general setting, and establish moment bounds in Theorem 4.2.2. In Subsection 4.2.2 we consider a small lengthscale regime, and show in Theorem 4.2.8 that the thresholded estimator significantly improves upon the standard sample covariance estimator. Finally, in Subsection 4.2.3 we apply our new covariance estimation theory to demonstrate the advantage of using thresholded covariance estimators within ensemble Kalman filters.

4.2.1 Thresholded Estimation of Covariance Operators

Let $u, u_1, u_2, \dots, u_N$ be *i.i.d.* centered almost surely continuous Gaussian random functions on $D = [0,1]^d$ taking values in $\mathbb{R}$ with covariance function (kernel) $k : D \times D \to \mathbb{R}$ and covariance operator $\mathcal{C} : L^2(D) \to L^2(D)$, so that, for $x, x' \in D$ and $\psi \in L^2(D)$,

$$k(x,x') := \mathbb{E}\Big[u(x)u(x')\Big], \qquad (\mathcal{C}\psi)(\cdot) := \int_D k(\cdot, x')\psi(x')\, dx'.$$

The sample covariance function $\widehat{k}(x, x')$ and sample covariance operator $\widehat{\mathcal{C}}$ are defined analogously by

$$\widehat{k}(x, x') := \frac{1}{N} \sum_{n=1}^{N} u_n(x) u_n(x'), \qquad (\widehat{\mathcal{C}} \, \psi)(\cdot) := \int_D \widehat{k}(\cdot, x') \psi(x') \, dx'.$$

We introduce the thresholded sample covariance estimators with thresholding parameter ρ_N

$$\widehat{k}_{\rho_N}(x, x') := \widehat{k}(x, x') \mathbf{1}_{\{|\widehat{k}(x,x')| \geq \rho_N\}}(x, x'), \qquad (\widehat{\mathcal{C}}_{\rho_N} \, \psi)(\cdot) := \int_D \widehat{k}_{\rho_N}(\cdot, x') \psi(x') \, dx',$$

where $\mathbf{1}_A$ denotes the indicator function of the set A. Our first main result, Theorem 4.2.2 below, relies on the following general assumption:

Assumption 4.2.1. $u, u_1, u_2, \ldots, u_N$ *are i.i.d. centered almost surely continuous Gaussian random functions on* $D = [0, 1]^d$ *taking values in* $\mathbb{R}$ *with covariance function* k. *Moreover, the following holds:*

(i) $\sup_{x \in D} \mathbb{E}\big[u(x)^2\big] = 1.$

(ii) For some $q \in (0, 1)$ *and* $R_q > 0$, $\sup_{x \in D} \left(\int_D |k(x, x')|^q \, dx' \right)^{\frac{1}{q}} \leq R_q$.

We assume fully observed functional data and defer extensions to partially observed data James et al. [2000], James and Sugar [2003], Yao et al. [2005a,b], Qiao et al. [2020], Fang et al. [2023] to future work. Assumption 4.2.1 (i) normalizes the fields to have unit maximum marginal variance over D. Assumption 4.2.1 (ii) generalizes standard notions of sparsity in finite dimension to our infinite-dimensional setting —refer e.g. to Bickel and Levina [2008a], Cai and Zhou [2012b], Wainwright [2019], which study estimation of a covariance matrix $\Sigma = (\sigma_{ij}) \in \mathbb{R}^{d_u \times d_u}$ under the row-wise approximate sparsity assumption that $\max_i \sum_{j=1}^{d_u} |\sigma_{ij}|^q \leq \widetilde{R}_q^q$.

Our first main result establishes moment bounds on the deviation of the thresholded covariance estimator from its target in terms of the approximate sparsity level R_q and the expected supremum of the field, the latter of which determines the scaling of ρ_N. We prove Theorem 4.2.2 and several auxiliary results of independent interest in Section 4.3.

Theorem 4.2.2. *Suppose that Assumption 4.2.1 holds. Let* $1 \leq c_0 \leq \sqrt{N}$ *and set*

$$\rho_N := c_0 \left[\frac{1}{N} \vee \frac{1}{\sqrt{N}} \mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \vee \frac{1}{N} \Big(\mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \Big)^2 \right], \tag{4.1}$$

$$\widehat{\rho}_N := c_0 \left[\frac{1}{N} \vee \frac{1}{\sqrt{N}} \Big(\frac{1}{N} \sum_{n=1}^{N} \sup_{x \in D} u_n(x) \Big) \vee \frac{1}{N} \Big(\frac{1}{N} \sum_{n=1}^{N} \sup_{x \in D} u_n(x) \Big)^2 \right]. \tag{4.2}$$

Then, for any $p \geq 1$,

$$\big[\mathbb{E} \| \widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C} \|^p \big]^{\frac{1}{p}} \lesssim_p R_q^q \rho_N^{1-q} + \rho_N e^{-\frac{c}{p} N \left(\rho_N \wedge \rho_N^2 \right)}, \tag{4.3}$$

where c *is a universal constant.*

An appealing feature of Theorem 4.2.2 is that it holds for *any* sample size $N \geq 1$. The following immediate corollary provides a simplified statement which holds for sufficiently large sample size.

Corollary 4.2.3. *Suppose that Assumption 4.2.1 holds and that* $\sqrt{N} \geq \mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \geq \frac{1}{\sqrt{N}}$. *Set*

$$\rho_N := \frac{1}{\sqrt{N}} \mathbb{E}\Big[\sup_{x \in D} u(x) \Big], \qquad \widehat{\rho}_N := \frac{1}{\sqrt{N}} \Big(\frac{1}{N} \sum_{n=1}^{N} \sup_{x \in D} u_n(x) \Big).$$

Then, for any $p \geq 1$,

$$\big[\mathbb{E} \| \widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C} \|^p \big]^{\frac{1}{p}} \lesssim_p R_q^q \rho_N^{1-q} + \rho_N e^{-\frac{c}{p} N \rho_N^2},$$

where c *is a universal constant.*

To the best of our knowledge, Theorem 4.2.2 and Corollary 4.2.3 are the first results in the literature to consider covariance operator estimation under the natural sparsity Assumption 4.2.1 (ii). As will be discussed next, the first of the two terms in the right-hand side of (4.3) is reminiscent of existing results for covariance matrix estimation. The second term in (4.3) depends only on the expected supremum of the field, and, as we will show in Subsection 4.2.2, it is negligible in the small lengthscale regime.

For covariance matrix estimation under ℓ_q-sparsity, [Wainwright, 2019, Theorem 6.27] proves that if the sample covariance matrix satisfies $|\widehat{\Sigma}_{ij} - \Sigma_{ij}| \lesssim \widetilde{\rho}_N$ for all $1 \leq i, j \leq d_u$, then the error of an estimator with thresholding parameter $\widetilde{\rho}_N$ can be bounded by $\widetilde{R}_q^q \widetilde{\rho}_N^{1-q}$, where $\widetilde{R}_q$ is a quantity analogous to our R_q that controls the row-wise ℓ_q-sparsity of Σ. This explains the choice of thresholding parameter

$$\widetilde{\rho}_N \asymp \frac{1}{\sqrt{N}} \sqrt{\log d_u} \asymp \frac{1}{\sqrt{N}} \mathbb{E}\left[\max_{1 \leq i \leq d_u} u_i\right]$$

in finite dimension, which ensures an entry-wise control on the sample covariance matrix with high probability. Analogously, our infinite-dimensional theory relies on sup-norm bounds for the sample covariance function $\widehat{k}(x, x')$; we obtain these bounds in Subsection 4.3.1 using tools from empirical process theory. For instance, Proposition 4.3.3 shows that with our choice of thresholding parameter ρ_N, we have $\sup_{x \in D} |\widehat{k}(x, x') - k(x, x')| \lesssim \rho_N$ with high probability. Therefore, Theorem 4.2.2 and Corollary 4.2.3 reveal that the expected supremum is the key dimension-free quantity that determines the choice of thresholding parameter and the error of estimation in both finite and infinite-dimensional settings. Since in practice the expected supremum of the field (and hence ρ_N) is unknown, we replace it with $\widehat{\rho}_N$ to define a computable thresholded estimator $\widehat{\mathcal{C}}_{\widehat{\rho}_N}$. The concentration of $\widehat{\rho}_N$ around ρ_N is established in Lemma 4.3.4.

Remark 4.2.4. *In contrast to existing results in the finite-dimensional setting (see e.g. Bickel and Levina [2008a], Cai and Zhou [2012b], Wainwright [2019]) that provide in-probability* *bounds or moment bounds of order up to* $p = 2$*, Theorem 4.2.2 provides moment bounds for all* $p \geq 1$*. For example, [Wainwright, 2019, Theorem 6.27] shows a high-probability statement where* $\widetilde{\rho}_N$ *necessarily depends on the desired confidence level. Consequently, [Wainwright, 2019, Theorem 6.27] cannot be used to derive moment bounds of arbitrary order. In contrast, Theorem 4.2.2 shows that the tuning parameter of the covariance operator estimator need not be tied to the confidence level. The proof technique therefore contributes to the literature on confidence parameter independent estimators; see e.g. Bellec et al. [2018] for an analogous finding that, contrary to standard practice Bickel et al. [2009], the Lasso tuning parameter need not depend on the confidence level.*

Remark 4.2.5. *The proof of the small lengthscale results in Subsections 4.2.2 and 4.2.3 utilizes Theorem 4.2.2 with a careful choice of thresholding parameter prefactor* c_0*. However, the exponential improvement in sample complexity established in Theorems 4.2.8 and 4.2.10 holds for any fixed value* $c_0 \gtrsim 1$*. As noted in [Bickel and Levina, 2008a, Section 3] and [Cai and Liu, 2011, Section 4], establishing an optimal choice of prefactor* c_0 *is challenging even in the simpler setting of covariance matrix estimation, where* c_0 *is often taken as a fixed constant or chosen empirically through cross-validation Bickel and Levina [2008a], Cai and Liu [2011], Cai and Yuan [2012]. We will numerically showcase in Subsection 4.2.2 the exponential improvement of a thresholded estimator with the choice* $c_0 = 5$*.*

Remark 4.2.6. *As in the finite-dimensional setting Cai and Zhou [2012b], El Karoui [2008], our thresholded estimator* $\widehat{\mathcal{C}}_{\widehat{\rho}_N}$ *is positive semi-definite with high probability, but it is not guaranteed to be positive semi-definite. Fortunately, a simple modification ensures positive semi-definiteness while maintaining the same order of estimation error achieved by the original estimator. Notice that* $\widehat{\mathcal{C}}_{\widehat{\rho}_N}$ *is a self-adjoint and Hilbert-Schmidt operator since* $\int_{D\times D} \big|\widehat{k}_{\rho_N}(x,x')\big|^2 dx dx' < \infty$*, see [Hunter and Nachtergaele, 2001, Example 9.23]. Therefore, there is an orthonormal basis* $\{\varphi_i\}_{i=1}^{\infty}$ *of* $L^2(D)$ *consisting of eigenfunctions of* $\widehat{\mathcal{C}}_{\widehat{\rho}_N}$ *such that* $\widehat{k}_{\rho_N}(x,x') = \sum_{i=1}^{\infty} \widehat{\lambda}_i \varphi_i(x)\varphi_i(x')$*, where* $\widehat{\lambda}_i$ *is the* i*-th eigenvalue of* $\widehat{\mathcal{C}}_{\widehat{\rho}_N}$*. Let* $\widehat{\lambda}_i^+ = \widehat{\lambda}_i \vee 0$

be the positive part of $\widehat{\lambda}_i$ *and define*

$$\widehat{k}^+_{\rho_N}(x,x') := \sum_{i=1}^{\infty} \widehat{\lambda}_i^+ \varphi_i(x)\varphi_i(x'), \qquad (\widehat{\mathcal{C}}^+_{\rho_N}\,\psi)(\cdot) := \int_D \widehat{k}^+_{\rho_N}(\cdot,x')\psi(x')\,dx'.$$

Then, $\widehat{\mathcal{C}}^+_{\rho_N}$ *is positive semi-definite and further*

$$\begin{aligned}\|\widehat{\mathcal{C}}^+_{\rho_N} - \mathcal{C}\| &\le \|\widehat{\mathcal{C}}^+_{\rho_N} - \widehat{\mathcal{C}}_{\rho_N}\| + \|\widehat{\mathcal{C}}_{\rho_N} - \mathcal{C}\| \le \max_{i:\widehat{\lambda}_i \le 0} |\widehat{\lambda}_i| + \|\widehat{\mathcal{C}}_{\rho_N} - \mathcal{C}\| \\ &\le \max_{i:\widehat{\lambda}_i \le 0} |\widehat{\lambda}_i - \lambda_i| + \|\widehat{\mathcal{C}}_{\rho_N} - \mathcal{C}\| \le 2\|\widehat{\mathcal{C}}_{\rho_N} - \mathcal{C}\|,\end{aligned}$$

where λ_i *is the i-th eigenvalue of* $\mathcal{C}$*. Thus,* $\widehat{\mathcal{C}}^+_{\rho_N}$ *is positive semi-definite and attains the same estimation error as the original thresholded estimator* $\widehat{\mathcal{C}}_{\rho_N}$*. In light of this fact, we will henceforth assume that* $\widehat{\mathcal{C}}_{\rho_N}$ *is positive semi-definite wherever needed.*

4.2.2 Small Lengthscale Regime

Our second main result, Theorem 4.2.8, shows that in the small lengthscale regime thresholded estimators enjoy an exponential improvement in sample complexity relative to the sample covariance estimator. To formalize this regime, we introduce the following additional assumption:

Assumption 4.2.7. *The following holds:*

(i) k *depends on a correlation lengthscale parameter* $\lambda > 0$, *so that* $k(x,x') = \mathsf{K}(|x-x'|/\lambda)$ *for an isotropic base kernel* $\mathsf{k} : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ *with* $\mathsf{k}(x,x') = \mathsf{K}(|x-x'|)$.

(ii) The base kernel k *is positive, so that* $\mathsf{k}(x,x') = \mathsf{K}(|x-x'|) > 0$. *Further,* $\mathsf{K}(r)$ *is differentiable, strictly decreasing on* $[0,\infty)$, *and satisfies* $\lim_{r\to\infty}\mathsf{K}(r) = 0$.

Assumption 4.2.7 makes explicit the dependence of the kernel on the correlation lengthscale parameter λ. While restrictive, the requirement of isotropy is often invoked in applications Williams and Rasmussen [2006], Stein [2012]. As discussed later, the nonparametric Assumption 4.2.7 is satisfied by important parametric covariance functions, such as squared exponential and Matérn models. The *small lengthscale regime* holds whenever Assumption 4.2.7 is satisfied and λ is sufficiently small. In the scientific applications that motivate our work, the dimension of the physical space is small ($d = 1, 2, 3$). Hence, we will treat d as a constant in our analysis of the small lenghtscale regime. Theorem 4.2.8 compares the errors of sample and thresholded covariance estimators. The proof can be found in Section 4.4.

Theorem 4.2.8. *Suppose that Assumptions 4.2.1 and 4.2.7 hold. Let $c_0 \gtrsim 1$ be an absolute constant and set*

$$\widehat{\rho}_N := \frac{c_0}{\sqrt{N}}\Big(\frac{1}{N}\sum_{n=1}^{N} \sup_{x\in D} u_n(x)\Big).$$

There is a universal constant $\lambda_0 > 0$ such that for $\lambda < \lambda_0$ and $N \gtrsim \log(\lambda^{-d})$, the sample covariance estimator and the thresholded covariance estimator satisfy

$$\frac{\mathbb{E}\|\widehat{\mathcal{C}} - \mathcal{C}\|}{\|\mathcal{C}\|} \asymp \sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N}, \tag{4.4}$$

$$\frac{\mathbb{E}\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\|}{\|\mathcal{C}\|} \le c(q)\Big(\frac{\log(\lambda^{-d})}{N}\Big)^{\frac{1-q}{2}}, \tag{4.5}$$

where $c(q)$ is a constant that depends only on q.

Remark 4.2.9. *The term $c(q)$ in (4.5) admits a form*

$$c(q) \asymp \frac{\int_0^\infty \mathsf{K}(r)^q r^{d-1}dr}{\int_0^\infty \mathsf{K}(r) r^{d-1}dr}.$$

As an explicit example, for the squared exponential kernel defined in (4.7), we have $\mathsf{K}^{SE}(r) = e^{-r^2/2}$ and a straightforward calculation shows that $c(q) \asymp q^{-d/2}$.

Theorem 4.2.8 shows that, for sufficiently small λ, we need $N \gtrsim \lambda^{-d}$ samples to control the relative error of the sample covariance estimator, while $N \gtrsim \log(\lambda^{-d})$ samples suffice to control the relative error of the thresholded estimator. The error bound in (4.5) is reminiscent of the convergence rate $s_0\big(\frac{\log d_u}{N}\big)^{(1-q)/2}$ of thresholded estimators for ℓ_q-sparse matrices $\Sigma \in \mathbb{R}^{d_u \times d_u}$ with sparsity level s_0 Bickel and Levina [2008a], Cai and Zhou [2012b]. Therefore, Theorem 4.2.8 indicates that, in our infinite-dimensional setting, the parameter λ^{-d} plays an analogous role to d_u and $c(q)$ plays an analogous role to s_0. However, we remark that the estimation error in Theorem 4.2.8 is *relative error*, whereas in the finite-dimensional covariance matrix estimation literature Bickel and Levina [2008a], Cai and Zhou [2012b], Cai and Liu [2011], the estimation error is often *absolute error*. While in the finite-dimensional setting the sparsity parameter s_0 may increase with d_u, the constant $c(q)$ in our bound (4.5) is independent of the lengthscale parameter λ. Moreover, inspired by the minimax optimality of thresholded estimators for ℓ_q-sparse covariance matrix estimation Cai and Zhou [2012b], we conjecture that the convergence rate (4.5) is also minimax optimal, and we intend to investigate this question in future work.

The bound (4.4) for the sample covariance estimator relies on the seminal work Koltchinskii and Lounici [2017], which shows that, for any sample size N,

$$\frac{\mathbb{E}\|\widehat{\mathcal{C}} - \mathcal{C}\|}{\|\mathcal{C}\|} \asymp \sqrt{\frac{r(\mathcal{C})}{N}} \vee \frac{r(\mathcal{C})}{N}, \qquad r(\mathcal{C}) := \frac{\mathrm{Tr}(\mathcal{C})}{\|\mathcal{C}\|}. \tag{4.6}$$

Consequently, (4.4) follows by a sharp characterization of the operator norm and the trace of $\mathcal{C}$ in terms of λ. In contrast, the bound (4.5) for the thresholded estimator relies on our new Theorem 4.2.2, and requires an analogous characterization of the thresholding parameter ρ_N and approximate sparsity level R_q in terms of λ.

In the remainder of this subsection, we illustrate Theorem 4.2.8 with a simple numerical experiment where we consider the estimation of covariance operators for squared exponential (SE) and Matérn (Ma) models in dimensions $d = 1$ and $d = 2$ at small lengthscales. We emphasize that our theory is developed under mild nonparametric assumptions on the covariance kernel as outlined in Assumption 4.2.7; however, for simplicity here we focus on two important parametric models. For $x, x' \in D$, define the corresponding covariance functions

$$k_\lambda^{\mathrm{SE}}(x,x') := \exp\left(-\frac{|x-x'|^2}{2\lambda^2}\right), \tag{4.7}$$

$$k_\lambda^{\mathrm{Ma}}(x,x') := \frac{2^{1-\nu}}{\Gamma(\nu)}\left(\frac{\sqrt{2\nu}}{\lambda}|x-x'|\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}}{\lambda}|x-x'|\right), \tag{4.8}$$

where Γ denotes the Gamma function and K_ν denotes the modified Bessel function of the second kind. In both cases, the parameter λ is interpreted as the correlation lengthscale of the field and Assumption 4.2.7 is satisfied. Moreover, Assumption 4.2.1 is satisfied by the squared exponential model, and it is satisfied by the Matérn model provided that the smoothness parameter ν satisfies $\nu > (\frac{d-1}{2} \vee \frac{1}{2})$. We refer to [Sanz-Alonso and Yang, 2022b, Lemma 4.2] for the almost sure continuity of random samples and to [Nobile and Tesei, 2015, Appendix 3, Lemma 11] for the Hölder continuity of the Matérn covariance function $\mathsf{K}^{\mathrm{Ma}}(r)$. For the Matérn model, we take the smoothness parameter to be $\nu = 3/2$ in our experiments.

We will report results in physical dimension $d = 1$ and $d = 2$. To respectively resolve small lengthscales up to order $\lambda \asymp 10^{-3}$ and $\lambda \asymp 10^{-2}$, we discretize the domain $D = [0,1]$ with a mesh of $L = 1250$ uniformly spaced points and the domain $D = [0,1]^2$ with $L = 10,000$ points. In the $d = 1$ case we consider a total of 30 lengthscales arranged uniformly in log-space and ranging from 10^{-3} to $10^{-0.1}$, and in the $d = 2$ case we consider a total of 10 lengthscales arranged in log-space and ranging from $10^{-2.3}$ to $10^{-0.1}$. For each lengthscale λ, with corresponding covariance operator $\mathcal{C}$, the discretized covariance operators are given by the $L \times L$ covariance matrices

$$\mathcal{C}^{ij} := k(x_i, x_j), \qquad 1 \le i, j \le L,$$

and we sample $N = 5\log(\lambda^{-1})$ realizations of a Gaussian process on the mesh, denoted $u_1, \ldots, u_N \sim \mathcal{N}(0, \mathcal{C})$. We then compute the empirical and thresholded sample covariance matrices

$$\widehat{\mathcal{C}}^{ij} := \frac{1}{N}\sum_{n=1}^{N} u_n(x_i)u_n(x_j), \qquad \widehat{\mathcal{C}}^{ij}_{\widehat{\rho}_N} := \widehat{\mathcal{C}}^{ij}\mathbf{1}_{\{|\widehat{\mathcal{C}}^{ij}|\geq \widehat{\rho}_N\}}, \qquad 1 \leq i,j \leq L,$$

scaling the thresholding level $\widehat{\rho}_N$ as described in Theorem 4.2.2.

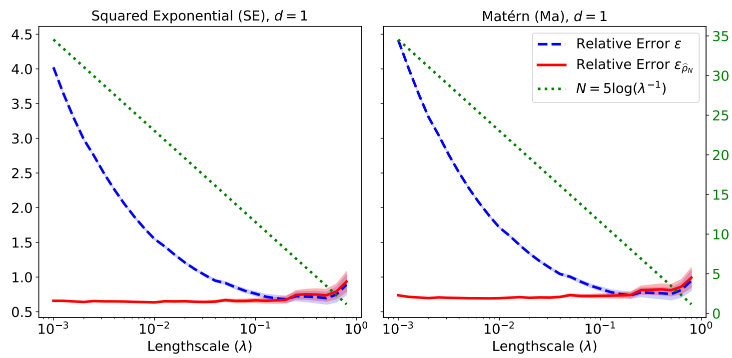

Figure 4.1: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue) and thresholded ($\varepsilon_{\widehat{\rho}_N}$, solid red) covariance estimators based on sample size (N, dotted green) for the squared exponential kernel (left) and Matérn kernel (right) in $d = 1$ over 100 trials.

To quantify the performance of each of the estimators, we compute their relative errors

$$\varepsilon := \frac{\|\widehat{\mathcal{C}} - \mathcal{C}\|}{\|\mathcal{C}\|}, \qquad \varepsilon_{\widehat{\rho}_N} := \frac{\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\|}{\|\mathcal{C}\|}.$$

The experiment is repeated a total of 100 times for each lengthscale in the case $d = 1$ and 30 times for each lengthscale in the case $d = 2$. In Figure 4.1, we plot average relative errors as well as 95% confidence intervals over the 100 trials for both squared exponential and Matérn models in $d = 1$, along with the sample size for each lengthscale setting. In Figure 4.2, we present the $d = 2$ analog of Figure 4.1. Our theoretical results are clearly illustrated: taking only $N = 5\log(\lambda^{-d})$ samples, the relative error in the thresholded estimator remains constant as the lengthscale decreases, whereas the relative error in the sample covariance operator diverges. Notice that Figures 4.1 and 4.2 also show that thresholding can increase the relative error for fields with large correlation lengthscale.

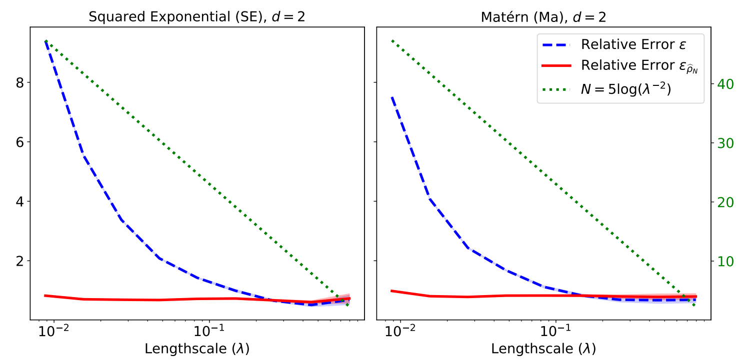

Figure 4.2: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue) and thresholded ($\varepsilon_{\widehat{\rho}_N}$, solid red) covariance estimators based on sample size (N, dotted green) for the squared exponential kernel (left) and Matérn kernel (right) in $d = 2$ over 30 trials.

4.2.3 Application in Ensemble Kalman Filters

Nonlinear filtering is concerned with online estimation of the state of a dynamical system from partial and noisy observations. Filtering algorithms blend the dynamics and observations by sequentially solving inverse problems of the form

$$y = \mathcal{A}u + \eta, \tag{4.9}$$

where $y \in \mathbb{R}^{d_y}$ denotes the observation, $u \in L^2(D)$ denotes the state, $\mathcal{A} : L^2(D) \to \mathbb{R}^{d_y}$ is a linear observation operator, and $\eta \sim \mathcal{N}(0, \Gamma)$ is the observation error with positive definite covariance matrix Γ. In Bayesian filtering Sanz-Alonso et al. [2023a], the model dynamics define a prior or *forecast* distribution on the state, which is combined with the data likelihood implied by the observation model (4.9) to obtain a posterior or *analysis* distribution. In most applications, the update from forecast to analysis distribution must be implemented through an approximate filtering algorithm. For instance, in operational numerical weather forecasting where the state may represent a temperature field along the surface of the Earth, discretizations of size 10^9 are routinely used to capture small lengthscales on the order of kilometers. In this setting, computing exactly the Kalman formulas that define the forecast-to-analysis update would be unfeasible.

Ensemble Kalman filters (EnKFs) are a rich family of algorithms scalable to highly complex data assimilation tasks Evensen [2009], including operational numerical weather forecasting Houtekamer and Zhang [2016]. The key idea behind these methods is to represent forecast and analysis distributions using an ensemble of N particles, so that the computational cost is controlled by the number of particles, which is typically small, rather than by the level of discretization. For instance, in operational weather forecasting $N \sim 10^2 \ll 10^9$; we refer to Tippett et al. [2003] for a summary of the computational and memory costs of different EnKFs in terms of the discretization level and the number of particles. Taking as input a forecast ensemble $\{u_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathcal{C})$ and observed data y generated according to (4.9), EnKFs produce an analysis ensemble $\{\upsilon_n\}_{n=1}^N$. Each analysis particle υ_n is obtained by nudging a forecast particle u_n towards the observed data y. The amount of nudging is controlled by a *Kalman gain* operator to be estimated using the first two moments of the forecast ensemble. Vanilla implementations of EnKFs rely on the sample covariance, see e.g. [Sanz-Alonso et al., 2023a, Algorithm 10.2]. However, some form of covariance localization is required for EnKFs to scale to operational settings Houtekamer and Mitchell [2001]. While the use of localization within EnKFs is standard, few works have demonstrated its statistical benefit Tong [2018], Al-Ghattas and Sanz-Alonso [2024b], and none in the functional setting that is most relevant in applications. In this subsection we show that thresholded covariance operator estimators within the EnKF analysis step can dramatically reduce the ensemble size required to approximate an idealized, non-implementable, *mean-field* EnKF that uses the population moments of the forecast distribution. Consequently, we identify an ensemble size which suffices for each EnKF particle to be updated similarly as in the limit of infinite number of particles. We refer to Herty and Visconti [2019], Calvello et al. [2022] for recent works that study the behavior of EnKFs in the mean-field limit.

Define the mean-field EnKF analysis update by

$$\upsilon_n^\star := u_n + \mathscr{K}(\mathcal{C})\big(y - \mathcal{A}u_n - \eta_n\big), \qquad 1 \le n \le N, \tag{4.10}$$

where $\{\eta_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,\Gamma)$ and

$$\mathscr{K}(\mathcal{C}) := \mathcal{C}\mathcal{A}^*(\mathcal{A}\mathcal{C}\mathcal{A}^* + \Gamma)^{-1} \tag{4.11}$$

denotes the Kalman gain. Practical algorithms do not have access to the forecast distribution, and rely instead on the forecast ensemble to estimate both $\mathcal{C}$ and $\mathscr{K}$. We will investigate two popular analysis steps, given by

$$\upsilon_n := u_n + \mathscr{K}(\widehat{\mathcal{C}})\big(y - \mathcal{A}u_n - \eta_n\big), \qquad 1 \le n \le N, \tag{4.12}$$

$$\upsilon_n^\rho := u_n + \mathscr{K}(\widehat{\mathcal{C}}_{\rho_N})\big(y - \mathcal{A}u_n - \eta_n\big), \qquad 1 \le n \le N. \tag{4.13}$$

The analysis step in (4.12) is known as the *perturbed observation* or *stochastic* EnKF Burgers et al. [1998]. For simplicity of exposition, we will assume here that when updating u_n, this particle is not included in the sample covariance $\widehat{\mathcal{C}}$ used to define the Kalman gain. This slight modification of the sample covariance will facilitate a cleaner statement and proof of our main result, Theorem 4.2.10, without altering the qualitative behavior of the algorithm. The analysis step in (4.13) is based on a thresholded covariance operator estimator. Again, we assume that the thresholded estimator $\widehat{\mathcal{C}}_{\rho_N}$ is defined without using the particle u_n. The following result is a direct consequence of our theory on covariance operator estimation in the small lengthscale regime. The proof can be found in Section 4.5.

Theorem 4.2.10 (Approximation of Mean-Field EnKF)**.** *Suppose that Assumptions 4.2.1 and 4.2.7 hold. Let y be generated according to* (4.9) *with bounded observation operator $\mathcal{A}: L^2(D) \to \mathbb{R}^{d_y}$. Let $\upsilon_n^\star$ be the mean-field* EnKF *update in* (4.10)*, and let υ_n and υ_n^ρ be the*

EnKF *and localized* EnKF *updates in* (4.12) *and* (4.13). *Let* $c_0 \gtrsim 1$ *be an absolute constant and set*

$$\rho_N \asymp \frac{c_0}{\sqrt{N}}\Big(\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x)\Big).$$

Then, there is a universal constant $\lambda_0 > 0$ *such that for* $\lambda < \lambda_0$ *and* $N \gtrsim \log(\lambda^{-d})$,

$$\mathbb{E}\left[|\upsilon_n - \upsilon_n^{\star}| \mid u_n, \eta_n\right] \lesssim c\left(\sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N}\right),$$
$$\mathbb{E}\left[|\upsilon_n^{\rho} - \upsilon_n^{\star}| \mid u_n, \eta_n\right] \lesssim c\left[c(q)\left(\frac{\log(\lambda^{-d})}{N}\right)^{\frac{1-q}{2}}\right],$$

where $c = \|\mathcal{A}\|\|\Gamma^{-1}\|\|\mathcal{C}\||y - \mathcal{A}u_n - \eta_n|$.

4.3 Thresholded Estimation of Covariance Operators

This section studies thresholded estimation of covariance operators in the general setting of Assumption 4.2.1. In Subsection 4.3.1 we show uniform error bounds on the sample covariance function estimator $\widehat{k}(x, x')$. These results are used in Subsection 4.3.2 to prove our first main result, Theorem 4.2.2.

4.3.1 Covariance Function Estimation

In this subsection we establish uniform error bounds on the sample covariance function estimator. These bounds will play a central role in our analysis of thresholded estimation of covariance operators developed in the next subsection. We first establish a high-probability bound, which is uniform over both arguments of the covariance function.

Proposition 4.3.1. *Under Assumption 4.2.1, there exist positive absolute constants* c_1, c_2

such that, for all $t \geq 1$, it holds with probability at least $1 - c_1 e^{-c_2 t}$ that

$$\sup_{x,x' \in D} \left| \widehat{k}(x,x') - k(x,x') \right| \lesssim \left[\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \right] \vee \frac{\left(\mathbb{E}\left[\sup_{x\in D} u(x)\right]\right)^2}{N}.$$

Proof. We will apply the product empirical process bound in [Mendelson, 2016, Theorem 1.13]. To that end, define the evaluation functional at $x \in D$ by

$$\ell_x : u \longmapsto \ell_x(u) = u(x)$$

and write

$$\left| \widehat{k}(x,x') - k(x,x') \right| = \left| \frac{1}{N} \sum_{n=1}^{N} u_n(x) u_n(x') - \mathbb{E}\left[u(x) u(x') \right] \right| = \left| \frac{1}{N} \sum_{n=1}^{N} \ell_x(u_n) \ell_{x'}(u_n) - \mathbb{E}\left[\ell_x(u) \ell_{x'}(u) \right] \right|,$$

so that

$$\sup_{x,x' \in D} \left| \widehat{k}(x,x') - k(x,x') \right| = \sup_{f,g \in \mathcal{F}} \left| \frac{1}{N} \sum_{n=1}^{N} f(u_n) g(u_n) - \mathbb{E}\left[f(u) g(u) \right] \right|,$$

where $\mathcal{F} := \{\ell_x\}_{x \in D}$ denotes the family of evaluation functionals. Note that $\{\ell_x\}_{x \in D}$ are continuous linear functionals on $C(D)$, the space of continuous functions on D endowed with its usual topology. We can then apply [Mendelson, 2016, Theorem 1.13] (see also [Al-Ghattas and Sanz-Alonso, 2024b, Theorem B.11]) which implies that, with probability $1 - c_1 e^{-c_2 t}$,

$$\sup_{x,x' \in D} \left| \widehat{k}(x,x') - k(x,x') \right| \lesssim \left[\left(\frac{t}{N} \vee \sqrt{\frac{t}{N}} \right) \left(\sup_{f \in \mathcal{F}} \|f\|_{\psi_2} \gamma_2\left(\mathcal{F}, \psi_2\right) \right) \right] \vee \frac{\gamma_2^2\left(\mathcal{F}, \psi_2\right)}{N}, \tag{4.14}$$

where here and henceforth γ_2 denotes Talagrand's generic complexity [Talagrand, 2022, Definition 2.7.3] and ψ_2 denotes the Orlicz norm with Orlicz function $\psi(x) = e^{x^2} - 1$, see e.g. [Vershynin, 2018, Definition 2.5.6]. Since u is Gaussian, the ψ_2-norm of linear functionals is equivalent to the L^2-norm. Hence,

$$\sup_{f\in\mathcal{F}} \|f\|_{\psi_2} \lesssim \sup_{f\in\mathcal{F}} \|f\|_{L^2} = \sup_{f\in\mathcal{F}} \sqrt{\mathbb{E}\left[f^2(u)\right]} = \sup_{x\in D} \sqrt{\mathbb{E}\left[u^2(x)\right]} = \sup_{x\in D} \sqrt{k(x,x)} = 1, \tag{4.15}$$

where we used Assumption 4.2.1 (i) in the last step. Next, to control the complexity $\gamma_2\left(\mathcal{F},\psi_2\right)$, let

$$\mathsf{d}(x,x') := \sqrt{\mathbb{E}\left[(u(x)-u(x'))^2\right]} = \|\ell_x(\cdot)-\ell_{x'}(\cdot)\|_{L^2(P)}, \quad x,x'\in D,$$

where P is the distribution of the random function u. Then,

$$\gamma_2(\mathcal{F},\psi_2) \overset{\text{(i)}}{\lesssim} \gamma_2(\mathcal{F},L^2) = \gamma_2(D,\mathsf{d}) \overset{\text{(ii)}}{\asymp} \mathbb{E}\left[\sup_{x\in D} u(x)\right], \tag{4.16}$$

where (i) follows by the equivalence of ψ_2 and L^2 norms for linear functionals and (ii) follows by Talagrand's majorizing-measure theorem [Talagrand, 2022, Theorem 2.10.1]. Combining the inequalities (4.14), (4.15), and (4.16) gives the desired result. □

Corollary 4.3.2. *Under Assumption 4.2.1, it holds that, for any $p\geq 1$,*

$$\left(\mathbb{E}\left[\sup_{x,x'\in D}\left|\widehat{k}(x,x')-k(x,x')\right|^p\right]\right)^{\frac{1}{p}} \lesssim_p \frac{\mathbb{E}\left[\sup_{x\in D} u(x)\right]}{\sqrt{N}} \vee \frac{\left(\mathbb{E}\left[\sup_{x\in D} u(x)\right]\right)^2}{N}.$$

Proof. The result follows by integrating the tail bound in Proposition 4.3.1. □

In contrast to Proposition 4.3.1, the following result provides uniform control over the error when holding fixed one of the two covariance function inputs. For this easier estimation task, we obtain an improved exponential tail bound that we will use in the proof of Theorem 4.2.2.

Proposition 4.3.3. *Suppose that Assumption 4.2.1 holds. Let* $1 \le c_0 \le N$ *and set*

$$\rho_N := c_0 \left[\frac{1}{N} \vee \frac{1}{\sqrt{N}} \mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \vee \frac{1}{N} \Big(\mathbb{E}\Big[\sup_{x \in D} u(x) \Big] \Big)^2 \right].$$

Then, for every $x' \in D$*, it holds with probability at least* $1 - 4e^{-c_1 N(\rho_N \wedge \rho_N^2)}$ *that*

$$\sup_{x \in D} \left| \widehat{k}(x, x') - k(x, x') \right| \lesssim \rho_N.$$

Proof. We will apply the multiplier empirical process bound in [Mendelson, 2016, Theorem 4.4]. To that end, we write

$$\begin{aligned} \left| \widehat{k}(x, x') - k(x, x') \right| &= \left| \frac{1}{N} \sum_{n=1}^{N} u_n(x) u_n(x') - \mathbb{E}\left[u(x) u(x') \right] \right| \\ &= \left| \frac{1}{N} \sum_{n=1}^{N} \ell_x(u_n) \ell_{x'}(u_n) - \mathbb{E}\left[\ell_x(u) \ell_{x'}(u) \right] \right|, \end{aligned}$$

so that for the class $\mathcal{F} := \{\ell_x\}_{x \in D}$ of evaluation functionals and for a fixed $g \in \mathcal{F}$, we have

$$\begin{aligned} \sup_{x \in D} \left| \widehat{k}(x, x') - k(x, x') \right| &= \sup_{f \in \mathcal{F}} \left| \frac{1}{N} \sum_{n=1}^{N} f(u_n) g(u_n) - \mathbb{E}\left[f(u) g(u) \right] \right| \\ &= \frac{1}{N} \sup_{f \in \mathcal{F}} \left| \sum_{n=1}^{N} \big(f(u_n) \xi_n - \mathbb{E}\left[f(u) \xi \right] \big) \right|, \end{aligned}$$

where $\xi_n := g(u_n)$. Note that $\xi_1, \ldots, \xi_N$ are i.i.d. copies of $\xi \sim \mathcal{N}\big(0, k(x', x')\big)$, where $x' \in D$ is the point indexed by g. By [Mendelson, 2016, Theorem 4.4] we have that for any $s, t \ge 1$, it holds with probability at least $1 - 2e^{-c_1 s^2 (\mathbb{E}[\sup_{x \in D} u(x)])^2} - 2e^{-c_1 N t^2}$ that

$$\sup_{x \in D} |\widehat{k}(x, x') - k(x, x')| \lesssim \frac{st \|\xi\|_{\psi_2} \, \mathbb{E}[\sup_{x \in D} u(x)]}{\sqrt{N}} \le \frac{st \, \mathbb{E}[\sup_{x \in D} u(x)]}{\sqrt{N}}, \tag{4.17}$$

where the last inequality follows by the fact that $\|\xi\|_{\psi_2} \leq \sqrt{k(x', x')} \leq \sup_{x\in D} \sqrt{k(x,x)} = 1$. We consider three cases:

Case 1: If $\mathbb{E}[\sup_{x\in D} u(x)] < \frac{1}{\sqrt{N}}$, then $\rho_N = \frac{c_0}{N} < 1$. We take

$$s = \frac{c_0}{\sqrt{N}\,\mathbb{E}[\sup_{x\in D} u(x)]} > 1, \qquad t = 1,$$

and then (4.17) implies that it holds with probability at least $1 - 2e^{-c_1 c_0^2/N} - 2e^{-c_1 N} \overset{\text{(i)}}{\geq} 1 - 4e^{-c_1 c_0^2/N} = 1 - 4e^{-c_1 N \rho_N^2}$ that

$$\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \lesssim \frac{st\,\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} = \frac{c_0}{N} = \rho_N,$$

where (i) follows since $c_0 < N$ by assumption.

Case 2: If $\frac{1}{\sqrt{N}} \leq \mathbb{E}[\sup_{x\in D} u(x)] \leq \sqrt{N}$, then $\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)]$. In this case, if $\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)] > 1$, we take

$$s = \sqrt{\frac{c_0\sqrt{N}}{\mathbb{E}[\sup_{x\in D} u(x)]}} \geq 1, \qquad t = \sqrt{\frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)]} > 1,$$

and then (4.17) implies that it holds with probability at least $1 - 4e^{-c_1 c_0 \sqrt{N}\mathbb{E}[\sup_{x\in D} u(x)]} = 1 - 4e^{-c_1 N \rho_N}$ that

$$\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \lesssim \frac{st\,\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)] = \rho_N.$$

If $\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)] \leq 1$, then we take $s = c_0 \geq 1$ and $t = 1$, and (4.17) implies that, with probability at least

$$1 - 2e^{-c_1 c_0^2 (\mathbb{E}[\sup_{x\in D} u(x)])^2} - 2e^{-c_1 N} \geq 1 - 4e^{-c_1 c_0^2 (\mathbb{E}[\sup_{x\in D} u(x)])^2} = 1 - 4e^{-c_1 N \rho_N^2},$$

it holds that

$$\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \lesssim \frac{st\,\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)] = \rho_N.$$

Case 3: If $\mathbb{E}[\sup_{x\in D} u(x)] > \sqrt{N}$, then $\rho_N = \frac{c_0}{N}(\mathbb{E}[\sup_{x\in D} u(x)])^2 > 1$. We take

$$s = \sqrt{c_0} \geq 1, \qquad t = \sqrt{c_0}\frac{\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} > 1,$$

and (4.17) implies that it holds with probability at least $1 - 4e^{-c_1 c_0(\mathbb{E}[\sup_{x\in D} u(x)])^2} = 1 - 4e^{-c_1 N\rho_N}$ that

$$\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \lesssim \frac{st\,\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} = \frac{c_0}{N}(\mathbb{E}[\sup_{x\in D} u(x)])^2 = \rho_N.$$

Combining the three cases above gives the desired result. □

4.3.2 Proof of Theorem 4.2.2

Before proving Theorem 4.2.2, the next result establishes moment and concentration bounds for the estimator $\widehat{\rho}_N$ of the thresholding parameter ρ_N.

Lemma 4.3.4. *Under the setting of Theorem 4.2.2, it holds that*

(A) For any $p \geq 1$, $\mathbb{E}\big[\,\widehat{\rho}_N^{\,p}\big] \lesssim_p \rho_N^p$.

(B) For any $t \in (0,1)$,

$$\mathbb{P}\left[\widehat{\rho}_N < t\rho_N\right] \leq 2\,e^{-\frac{1}{2}(1-\sqrt{t})^2 N(\mathbb{E}[\sup_{x\in D} u(x)])^2}\mathbf{1}\big\{\mathbb{E}[\sup_{x\in D} u(x)] \geq 1/\sqrt{N}\big\} \tag{4.18}$$

$$\leq 2\,e^{-\frac{1}{2}(1-\sqrt{t})^2 N(\rho_N \wedge \rho_N^2)}. \tag{4.19}$$

The proof of Lemma 4.3.4 can be found in Appendix A in the Supplementary Material

Al-Ghattas et al. [2024c].

Proof of Theorem 4.2.2. As shown in Lemma B.1 in the Supplementary Material Al-Ghattas et al. [2024c], the operator norm can be upper bounded as

$$\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\| \leq \sup_{x \in D} \int_D \left| \widehat{k}_{\widehat{\rho}_N}(x, x') - k(x, x') \right| dx'.$$

Let $\Omega_x := \left\{ x' \in D : |k(x, x')| \geq \widehat{\rho}_N \right\}$ and let Ω_x^c be its complement. Then, we have

$$\begin{aligned}
&\mathbb{E}\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\|^p \leq \mathbb{E}\left[\left(\sup_{x \in D} \int_D \left| \widehat{k}_{\widehat{\rho}_N}(x, x') - k(x, x') \right| \, dx' \right)^p \right] \\
&\leq 2^{p-1} \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x} \left| \widehat{k}_{\widehat{\rho}_N}(x, x') - k(x, x') \right| dx' \right)^p \right] \\
&+ 2^{p-1} \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x^c} \left| \widehat{k}_{\widehat{\rho}_N}(x, x') - k(x, x') \right| dx' \right)^p \right] \\
&\lesssim_p \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x} \left| \widehat{k}_{\widehat{\rho}_N}(x, x') - k(x, x') \right| dx' \right)^p \right] \\
&+ \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x^c} |k(x, x')| \mathbf{1}\left\{ |\widehat{k}(x, x')| < \widehat{\rho}_N \right\} dx' \right)^p \right] \\
&\quad + \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x^c} |\widehat{k}(x, x') - k(x, x')| \mathbf{1}\left\{ |\widehat{k}(x, x')| \geq \widehat{\rho}_N \right\} \mathbf{1}\left\{ |\widehat{k}(x, x') - k(x, x')| < 4|k(x, x')| \right\} dx' \right)^p \right] \\
&\quad + \mathbb{E}\left[\left(\sup_{x \in D} \int_{\Omega_x^c} |\widehat{k}(x, x') - k(x, x')| \mathbf{1}\left\{ |\widehat{k}(x, x')| \geq \widehat{\rho}_N \right\} \mathbf{1}\left\{ |\widehat{k}(x, x') - k(x, x')| \geq 4|k(x, x')| \right\} dx' \right)^p \right] \\
&=: I_1 + I_2 + I_3 + I_4,
\end{aligned} \tag{4.20}$$

where in the second inequality we used that $|a+b|^p \leq 2^{p-1}(|a|^p + |b|^p)$, which follows directly from the convexity of $f(x) = |x|^p$ for $p \geq 1$. We next bound the four terms $\{I_i\}_{i=1}^4$. To ease notation, we define

$$\|\widehat{k} - k\|_{\max} := \sup_{x, x' \in D} |\widehat{k}(x, x') - k(x, x')|.$$

For I_1, using that

$$\left|\widehat{k}_{\widehat{\rho}_N}(x,x') - k(x,x')\right| \leq \left|\widehat{k}_{\widehat{\rho}_N}(x,x') - \widehat{k}(x,x')\right| + \left|\widehat{k}(x,x') - k(x,x')\right| \leq \widehat{\rho}_N + \|\widehat{k} - k\|_{\max},$$

we have

$$I_1 = \mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x} \left|\widehat{k}_{\widehat{\rho}_N}(x,x') - k(x,x')\right| dx'\right)^p\right] \leq \mathbb{E}\left[\Big(\sup_{x\in D}\mathrm{Vol}(\Omega_x)\Big)^p \left(\widehat{\rho}_N + \|\widehat{k} - k\|_{\max}\right)^p\right],$$

where $\mathrm{Vol}(\Omega_x)$ denotes the Lebesgue measure of Ω_x. Notice that

$$R_q^q \geq \sup_{x\in D}\int_D |k(x,x')|^q dx' \geq \sup_{x\in D}\int_{\Omega_x} |k(x,x')|^q dx' \geq \sup_{x\in D}\int_{\Omega_x} \widehat{\rho}_N^{\,q} dx' = \widehat{\rho}_N^{\,q}\sup_{x\in D}\mathrm{Vol}(\Omega_x).$$

Combining this bound with the trivial bound $\sup_x \mathrm{Vol}(\Omega_x) \leq \mathrm{Vol}(D) = 1$ gives

$$\sup_{x\in D}\mathrm{Vol}(\Omega_x) \leq R_q^q \widehat{\rho}_N^{\,-q} \wedge 1.$$

Therefore, by Cauchy-Schwarz, we have that

$$\begin{aligned} I_1 &\leq \mathbb{E}\left[\left(R_q^q\widehat{\rho}_N^{\,-q} \wedge 1\right)^p (\widehat{\rho}_N + \|\widehat{k} - k\|_{\max})^p\right] \\ &\leq \sqrt{\mathbb{E}\Big[\left(R_q^q\widehat{\rho}_N^{\,-q} \wedge 1\right)^{2p}\Big]\ \mathbb{E}\Big[\left(\widehat{\rho}_N + \|\widehat{k} - k\|_{\max}\right)^{2p}\Big]}. \end{aligned} \tag{4.21}$$

Using Lemma 4.3.4 and Corollary 4.3.2 yields that

$$\mathbb{E}\Big[\left(\widehat{\rho}_N + \|\widehat{k} - k\|_{\max}\right)^{2p}\Big] \lesssim_p \mathbb{E}\big[\left(\widehat{\rho}_N\right)^{2p}\big] + \mathbb{E}\Big[\|\widehat{k} - k\|_{\max}^{2p}\Big] \lesssim_p \rho_N^{2p}. \tag{4.22}$$

On the other hand,

$$\mathbb{E}\Big[\big(R_q^q\widehat{\rho}_N^{-q}\wedge 1\big)^{2p}\Big] = R_q^{2pq}\,\mathbb{E}\Big[\widehat{\rho}_N^{-2pq}\wedge R_q^{-2pq}\Big] = R_q^{2pq}\int_0^\infty \mathbb{P}\Big[\big(\widehat{\rho}_N^{-2pq}\wedge R_q^{-2pq}\big)>t\Big]dt$$
$$= R_q^{2pq}\int_0^{R_q^{-2pq}} \mathbb{P}\Big[\widehat{\rho}_N^{-2pq}>t\Big]\,dt = 2pqR_q^{2pq}\int_{R_q}^\infty \mathbb{P}\left[\,\widehat{\rho}_N<t\right]\ t^{-2pq-1}dt.$$

If $R_q>\rho_N$, then

$$\mathbb{E}\Big[\big(R_q^q\widehat{\rho}_N^{-q}\wedge 1\big)^{2p}\Big] \le 2pqR_q^{2pq}\int_{\rho_N}^\infty t^{-2pq-1}dt = R_q^{2pq}\rho_N^{-2pq}. \tag{4.23}$$

If $R_q<\rho_N$, then

$$\begin{aligned}
&\mathbb{E}\Big[\big(R_q^q\widehat{\rho}_N^{-q}\wedge 1\big)^{2p}\Big] = 2pqR_q^{2pq}\left(\int_{\rho_N}^\infty + \int_{R_q}^{\rho_N}\right)\mathbb{P}\left[\widehat{\rho}_N<t\right]\ t^{-2pq-1}dt\\
&\le 2pqR_q^{2pq}\int_{\rho_N}^\infty t^{-2pq-1}dt + 2pqR_q^{2pq}\int_{R_q}^{\rho_N}\mathbb{P}\left[\widehat{\rho}_N<t\right]\ t^{-2pq-1}dt\\
&= R_q^{2pq}\rho_N^{-2pq} + 2pqR_q^{2pq}\rho_N^{-2pq}\int_{R_q\rho_N^{-1}}^1 \mathbb{P}\left[\widehat{\rho}_N<t\rho_N\right]\ t^{-2pq-1}dt\\
&\overset{\text{(i)}}{\le} R_q^{2pq}\rho_N^{-2pq} + 2pqR_q^{2pq}\rho_N^{-2pq}\int_{R_q\rho_N^{-1}}^1 2\exp\Big(-\frac{1}{2}(1-\sqrt{t})^2N(\rho_N\wedge\,\rho_N^2)\Big)t^{-2pq-1}dt\\
&\overset{\text{(ii)}}{=} R_q^{2pq}\rho_N^{-2pq}\left[1+8pq\int_0^{\sqrt{N(\rho_N\wedge\rho_N^2)}(1-\sqrt{R_q\rho_N^{-1}})}\frac{\big(N(\rho_N\wedge\,\rho_N^2)\big)^{2pq}\exp(-\frac{1}{2}t^2)}{\big(\sqrt{N(\rho_N\wedge\,\rho_N^2)}-t\big)^{4pq+1}}\ dt\right]\\
&\overset{\text{(iii)}}{\lesssim} R_q^{2pq}\rho_N^{-2pq} + R_q^{2pq}\rho_N^{-2pq}\cdot 8pq\left(\frac{2R_q^{-2pq}\rho_N^{2pq}}{4pq}e^{-\frac{1}{8}N(\rho_N\wedge\rho_N^2)\left(1-\sqrt{R_q\rho_N^{-1}}\right)^2}+\frac{2^{4pq}}{4pq}\right)\\
&\lesssim_p R_q^{2pq}\rho_N^{-2pq} + e^{-\frac{1}{8}N(\rho_N\wedge\rho_N^2)\left(1-\sqrt{R_q\rho_N^{-1}}\right)^2} \overset{\text{(iv)}}{\lesssim}_p R_q^{2pq}\rho_N^{-2pq} + e^{-cN(\rho_N\wedge\rho_N^2)},
\end{aligned} \tag{4.24}$$

where (i) follows from Lemma 4.3.4, (ii) follows by a change of variable, and (iii) follows by applying Lemma C.1 in the Supplementary Material Al-Ghattas et al. [2024c] with $\alpha=\sqrt{N(\rho_N\wedge\,\rho_N^2)}$ and $\beta=\sqrt{N(\rho_N\wedge\,\rho_N^2)}\sqrt{R_q\rho_N^{-1}}$. To prove (iv), notice that if $R_q\le\frac{1}{4}\rho_N$, then $\left|1-\sqrt{R_q\rho_N^{-1}}\right|>\frac{1}{2}$ and (iv) holds; if $\frac{1}{4}\rho_N<R_q<\rho_N$, then

$$e^{-\frac{1}{8}N(\rho_N\wedge\rho_N^2)\left(1-\sqrt{R_q\rho_N^{-1}}\right)^2}\le 1<16^pR_q^{2p}\rho_N^{-2p}\le 16^pR_q^{2pq}\rho_N^{-2pq}.$$

Combining the inequalities (4.21), (4.22), (4.23), and (4.24) gives that

$$I_1\le\sqrt{\mathbb{E}\Big[\big(R_q^q\widehat{\rho}_N^{-q}\wedge 1\big)^{2p}\Big]\ \mathbb{E}\Big[\big(\widehat{\rho}_N+\|\widehat{k}-k\|_{\max}\big)^{2p}\Big]}\lesssim_p R_q^{pq}\rho_N^{p(1-q)}+\rho_N^pe^{-cN(\rho_N\wedge\rho_N^2)}.$$

For I_2 and I_3,

$$\begin{aligned}
I_2+I_3&=\mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x^c}\big|k(x,x')\big|\ \mathbf{1}\big\{|\widehat{k}(x,x')|<\widehat{\rho}_N\big\}\,dx'\right)^p\right]\\
&+\mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x^c}|\widehat{k}(x,x')-k(x,x')|\,\mathbf{1}\big\{|\widehat{k}(x,x')|\ge\widehat{\rho}_N\big\}\,\mathbf{1}\big\{|\widehat{k}(x,x')-k(x,x')|<4|k(x,x')|\big\}\,dx'\right)^p\right]\\
&\lesssim\mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x^c}\big|k(x,x')\big|\,dx'\right)^p\right]=\mathbb{E}\left[\left(\widehat{\rho}_N\sup_{x\in D}\int_{\Omega_x^c}\left(\frac{|k(x,x')|}{\widehat{\rho}_N}\right)dx'\right)^p\right]\\
&\overset{\text{(i)}}{\le}\mathbb{E}\left[\left(\widehat{\rho}_N\sup_{x\in D}\int_{\Omega_x^c}\left(\frac{|k(x,x')|}{\widehat{\rho}_N}\right)^q dx'\right)^p\right]\le\mathbb{E}\left[R_q^{pq}\widehat{\rho}_N^{p(1-q)}\right]\overset{\text{(ii)}}{\lesssim_p}R_q^{pq}\rho_N^{p(1-q)},
\end{aligned}$$

where (i) follows since $q\in(0,1)$ and $|k(x,x')|<\widehat{\rho}_N$ for $x'\in\Omega_x^c$. To prove (ii), we notice that if $p(1-q)\le 1$, then using Jensen's inequality and Lemma 4.3.4 yields that $\mathbb{E}[\widehat{\rho}_N^{p(1-q)}]\le(\mathbb{E}[\widehat{\rho}_N])^{p(1-q)}\lesssim_p\rho_N^{p(1-q)}$. If $p(1-q)>1$, Lemma 4.3.4 implies that $\mathbb{E}[\widehat{\rho}_N^{p(1-q)}]\lesssim_p\rho_N^{p(1-q)}$.

For I_4,

$$
\begin{aligned}
I_4 &= \mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x^c} |\widehat{k}(x,x')-k(x,x')|\mathbf{1}\{|\widehat{k}(x,x')|\geq\widehat{\rho}_N\}\mathbf{1}\{|\widehat{k}(x,x')-k(x,x')|\geq 4|k(x,x')|\}\,dx'\right)^p\right]\\
&\overset{\text{(i)}}{\leq} \mathbb{E}\left[\left(\sup_{x\in D}\int_{\Omega_x^c} |\widehat{k}(x,x')-k(x,x')|\mathbf{1}\{|\widehat{k}(x,x')|\geq\widehat{\rho}_N\}\mathbf{1}\{|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\right)^p\right]\\
&\leq \mathbb{E}\left[\left(\sup_{x\in D}\int_D \sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\,\mathbf{1}\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\right)^p\right]\\
&= \mathbb{E}\left[\left(\int_D \sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\,\mathbf{1}\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\right)^p\right]\\
&\leq \mathbb{E}\left[\left(\|\widehat{k}-k\|_{\max}\int_D \mathbf{1}\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\right)^p\right]\\
&\leq \left(\mathbb{E}\big[\|\widehat{k}-k\|_{\max}^{2p}\big]\right)^{1/2}\left(\mathbb{E}\Big[\Big(\int_D \mathbf{1}\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\Big)^{2p}\Big]\right)^{1/2}\\
&\overset{\text{(ii)}}{\leq} \left(\mathbb{E}\big[\|\widehat{k}-k\|_{\max}^{2p}\big]\right)^{1/2}\left(\mathbb{E}\Big[\int_D \mathbf{1}\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\}\,dx'\Big]\right)^{1/2}\\
&= \left(\mathbb{E}\big[\|\widehat{k}-k\|_{\max}^{2p}\big]\right)^{1/2}\left(\int_D \mathbb{P}\left[\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq \frac{2}{3}\widehat{\rho}_N\right]dx'\right)^{1/2},
\end{aligned}
$$

where (i) follows since $|\widehat{k}(x,x')-k(x,x')|\geq 4|k(x,x')|$ implies that $|\widehat{k}(x,x')|\geq 3|k(x,x')|$, and therefore if $|\widehat{k}(x,x')-k(x,x')|\geq 4|k(x,x')|$ and $|\widehat{k}(x,x')|\geq\widehat{\rho}_N$, then it holds that

$$
|\widehat{k}(x,x')-k(x,x')|\geq|\widehat{k}(x,x')|-|k(x,x')|\geq\frac{2}{3}|\widehat{k}(x,x')|\geq\frac{2}{3}\widehat{\rho}_N.
$$

To prove (ii), note that $p\geq 1$ and $\int_D \mathbf{1}\left\{\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq\frac{2}{3}\widehat{\rho}_N\right\}dx'\leq 1$. Next, notice that

$$
\begin{aligned}
\mathbb{P}\left[\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq\frac{2}{3}\widehat{\rho}_N\right] &= \mathbb{P}\left[\frac{2}{3}(\rho_N-\widehat{\rho}_N)+\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq\frac{2}{3}\rho_N\right]\\
&\leq \mathbb{P}\left[\sup_{x\in D}|\widehat{k}(x,x')-k(x,x')|\geq\frac{1}{3}\rho_N\right]+\mathbb{P}\Big[\rho_N-\widehat{\rho}_N\geq\frac{1}{2}\rho_N\Big].
\end{aligned}
$$

Lemma 4.3.4 then implies that

$$\mathbb{P}\Big[\rho_N - \widehat{\rho}_N \geq \frac{1}{2}\rho_N\Big] = \mathbb{P}\Big[\widehat{\rho}_N \leq \frac{1}{2}\rho_N\Big] \lesssim e^{-c_1 N(\rho_N \wedge \rho_N^2)},$$

and Proposition 4.3.3 gives that

$$\mathbb{P}\left[\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \geq \frac{1}{3}\rho_N\right] \lesssim e^{-c_2 N(\rho_N \wedge \rho_N^2)}.$$

Moreover, Corollary 4.3.2 yields that $\Big(\mathbb{E}\big[\|\widehat{k}-k\|_{\max}^{2p}\big]\Big)^{1/2} \lesssim_p \rho_N^p$. Therefore,

$$I_4 \leq \Big(\mathbb{E}\big[\|\widehat{k}-k\|_{\max}^{2p}\big]\Big)^{1/2} \left(\int_D \mathbb{P}\left[\sup_{x\in D} |\widehat{k}(x,x') - k(x,x')| \geq \frac{2}{3}\widehat{\rho}_N\right] dx'\right)^{1/2} \lesssim_p \rho_N^p e^{-cN(\rho_N \wedge \rho_N^2)}.$$

Combining (4.20) with the estimates of I_1, I_2, I_3, and I_4 gives that

$$\mathbb{E}\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\|^p \lesssim_p I_1 + I_2 + I_3 + I_4 \lesssim_p R_q^{pq} \rho_N^{p(1-q)} + \rho_N^p e^{-cN(\rho_N \wedge \rho_N^2)},$$

and hence

$$\big[\mathbb{E}\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\|^p\big]^{\frac{1}{p}} \lesssim_p R_q^q \rho_N^{1-q} + \rho_N e^{-\frac{c}{p}N(\rho_N \wedge \rho_N^2)}.$$ □

4.4 Small Lengthscale Regime

This section studies thresholded estimation of covariance operators under the small lengthscale regime formalized in Assumption 4.2.7. We first present three lemmas which establish the sharp scaling of the L^q-sparsity level, the operator norm of the covariance operator, and the suprema of Gaussian fields in the small lengthscale regime. Combining these lemmas and Theorem 4.2.2, we then prove Theorem 4.2.8. Throughout this section, we use the notation "($\mathcal{B}$), $\lambda \to 0$" to indicate that there is a universal constant $\lambda_0 > 0$ such that if $\lambda < \lambda_0$, the conclusion ($\mathcal{B}$) holds.

The following result establishes the scaling of the L^q-sparsity level in the small lengthscale regime.

Lemma 4.4.1. *Under Assumption 4.2.7, it holds that*

$$\sup_{x \in D} \int_D |k(x,x')|^q dx' \asymp \lambda^d A(d) \int_0^\infty \mathsf{K}(r)^q r^{d-1} dr, \quad \lambda \to 0,$$

where $A(d)$ denotes the surface area of the unit sphere in $\mathbb{R}^d$.

Proof. We have that

$$\begin{aligned}
\sup_{x \in D} \int_D |k(x,x')|^q dx' &\geq \int_{D \times D} k(x,x')^q dx dx' = \int_{[0,1]^d \times [0,1]^d} \mathsf{K}(|x-x'|/\lambda)^q \, dx dx' \\
&= \lambda^{2d} \int_{[0,\lambda^{-1}]^d \times [0,\lambda^{-1}]^d} \mathsf{K}(|x-x'|)^q \, dx dx' \overset{\text{(i)}}{=} \lambda^{2d} \int_{[-\lambda^{-1},\lambda^{-1}]^d} \mathsf{K}(|w|)^q \prod_{j=1}^d (\lambda^{-1} - |w_j|) \, dw \\
&= \lambda^d \int_{[-\lambda^{-1},\lambda^{-1}]^d} \mathsf{K}(|w|)^q \prod_{j=1}^d (1 - \lambda |w_j|) \, dw \overset{\text{(ii)}}{\asymp} \lambda^d \int_{\mathbb{R}^d} \mathsf{K}(|w|)^q \, dw \\
&\overset{\text{(iii)}}{=} \lambda^d A(d) \int_0^\infty \mathsf{K}(r)^q r^{d-1} \, dr, \quad \lambda \to 0, \qquad (4.25)
\end{aligned}$$

where (i) follows by a change of variables $w = x - x'$, $z = x + x'$ and integrating z, (ii) follows by dominated convergence as $\lambda \to 0$, and (iii) follows from the polar coordinate transform in $\mathbb{R}^d$. On the other hand,

$$\begin{aligned}
\sup_{x \in D} \int_D |k(x,x')|^q dx' &\leq \sup_{x \in D} \int_{\mathbb{R}^d} \mathsf{K}(|x-x'|/\lambda)^q dx' \\
&= \int_{\mathbb{R}^d} \mathsf{K}(|x'|/\lambda)^q \, dx' = \lambda^d \int_{\mathbb{R}^d} \mathsf{K}(|x'|)^q \, dx' = \lambda^d A(d) \int_0^\infty \mathsf{K}(r)^q r^{d-1} \, dr,
\end{aligned}$$

which concludes the proof. □

Next, we establish the scaling of the operator norm of the covariance operator.

Lemma 4.4.2. *Under Assumption 4.2.7, it holds that*

$$\|\mathcal{C}\| \asymp \lambda^d A(d) \int_0^\infty \mathsf{K}(r) r^{d-1} dr, \quad \lambda \to 0,$$

where $A(d)$ denotes the surface area of the unit sphere in $\mathbb{R}^d$.

Proof. First, the operator norm can be upper bounded by

$$\|\mathcal{C}\| \leq \sup_{x \in D} \int_D |k(x, x')| dx' \asymp \lambda^d A(d) \int_0^\infty \mathsf{K}(r) r^{d-1} dr, \quad \lambda \to 0,$$

where the last step follows by Lemma 4.4.1.

For the lower bound, taking the test function $\psi(x) \equiv 1$ yields that, as $\lambda \to 0$,

$$\begin{aligned}\|\mathcal{C}\| = \sup_{\|\psi\|_{L^2}=1} \Big(\int_D \Big(\int_D k(x, x')\psi(x') dx' \Big)^2 dx \Big)^{1/2} &\geq \Big(\int_D \Big(\int_D k(x, x') dx' \Big)^2 dx \Big)^{1/2} \\ \overset{\text{(i)}}{\geq} \frac{1}{\sqrt{\mathrm{Vol}(D)}} \int_{D \times D} k(x, x') dx dx' \overset{\text{(ii)}}{=} \int_{D \times D} k(x, x') dx dx' &\overset{\text{(iii)}}{\asymp} \lambda^d A(d) \int_0^\infty \mathsf{K}(r) r^{d-1}\, dr,\end{aligned}$$

where (i) follows by Cauchy-Schwarz inequality, (ii) follows since $\mathrm{Vol}(D) = 1$ for $D = [0, 1]^d$, and (iii) follows from (4.25) with $q = 1$. This completes the proof. □

Finally, we establish the scaling of the suprema of Gaussian fields in the small lengthscale regime.

Lemma 4.4.3. *Under Assumption 4.2.7, it holds that*

$$\mathbb{E}\left[\sup_{x \in D} u(x)\right] \asymp \sqrt{\mathsf{K}(0) d \log\left(\frac{\sqrt{d}}{s\lambda}\right)}, \quad \lambda \to 0,$$

where $s > 0$ is the unique solution of $\mathsf{K}(s) = \frac{1}{2}\mathsf{K}(0)$, which is independent of λ.

Proof. By Fernique's theorem Fernique et al. [1975] and the discussion in [Van Handel, 2014, Theorem 6.19], for the stationary Gaussian random field u, it holds that

$$\mathbb{E}\left[\sup_{x\in D} u(x)\right] \asymp \int_0^\infty \sqrt{\log \mathcal{M}(D, \mathsf{d}, \varepsilon)}\, d\varepsilon, \tag{4.26}$$

where $\mathcal{M}(D, \mathsf{d}, \varepsilon)$ denotes the smallest cardinality of an ε-net of D in the canonical metric d given by

$$\mathsf{d}(x, x') := \sqrt{\mathbb{E}[(u(x) - u(x'))^2]} = \sqrt{2\mathsf{K}(0) - 2\mathsf{K}(\lambda^{-1}|x - x'|)} < \sqrt{2\mathsf{K}(0)}, \quad x, x' \in D.$$

Since under Assumption 4.2.7 the field is isotropic, it is necessarily stationary. Consequently, Fernique's bound implies that $\mathcal{M}(D, \mathsf{d}, \varepsilon) = 1$ for $\varepsilon \geq \sqrt{2\mathsf{K}(0)}$, and hence we can assume without loss of generality that $\varepsilon < \sqrt{2\mathsf{K}(0)}$ in the rest of the proof. Next, notice that

$$\mathsf{d}(x, x') = \sqrt{2\mathsf{K}(0) - 2\mathsf{K}(\lambda^{-1}|x - x'|)} \leq \varepsilon \quad \Longleftrightarrow \quad |x - x'| \leq \lambda \mathsf{K}^{-1}(\mathsf{K}(0) - \varepsilon^2/2),$$

where K^{-1} is the inverse function of K. By the standard volume argument [Vershynin, 2018, Proposition 4.2.12],

$$\begin{aligned}
\mathcal{M}(D, \mathsf{d}, \varepsilon) &= \mathcal{M}\big(D, |\cdot|, \lambda\mathsf{K}^{-1}(\mathsf{K}(0) - \varepsilon^2/2)\big) \\
&\geq \left(\frac{1}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0) - \varepsilon^2/2)}\right)^d \frac{\mathrm{Vol}(D)}{\mathrm{Vol}(B_2^d)} \geq \frac{1}{c_1}\left(\frac{1}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0) - \varepsilon^2/2)}\right)^d \left(\frac{d}{2\pi e}\right)^{d/2},
\end{aligned}$$

where we used that $\mathrm{Vol}(D) = 1$ and that, for the Euclidean unit ball B_2^d, it holds that $\mathrm{Vol}(B_2^d) \leq c_1(2\pi e/d)^{d/2}$ for some absolute constant $c_1 > 1$. On the other hand, using the fact that $D = [0,1]^d \subset \sqrt{d}B_2^d$, as well as $\mathcal{M}(B_2^d, |\cdot|, \varepsilon) \leq (3/\varepsilon)^d$ for $\varepsilon \leq 1$ [Vershynin, 2018, Corollary 4.2.13],

$$\begin{aligned}
&\mathcal{M}(D,\mathsf{d},\varepsilon) = \mathcal{M}\big(D,|\cdot|,\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)\big) \\
&\le \mathcal{M}\big(B_2^d,|\cdot|,d^{-1/2}\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)\big) \le \left[\left(\frac{3}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)}\right)^d d^{d/2}\right] \vee 1, \quad \varepsilon < \sqrt{2\mathsf{K}(0)}.
\end{aligned}$$

Therefore, (4.26) and the bounds we just established on the covering number $\mathcal{M}(D,\mathsf{d},\varepsilon)$ imply that

$$\begin{aligned}
\mathbb{E}\left[\sup_{x\in D} u(x)\right] &\asymp \int_0^{\sqrt{2\mathsf{K}(0)}} \sqrt{\log\left(\frac{1}{c_1}\left(\frac{1}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)}\right)^d \left(\frac{d}{2\pi e}\right)^{d/2}\right)\vee 0}\ d\varepsilon \\
&\asymp \sqrt{d}\int_0^{\sqrt{2\mathsf{K}(0)}} \sqrt{\log\left(\frac{c\sqrt{d}}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)}\right)\vee 0}\ d\varepsilon.
\end{aligned}$$

By a change of variable $t := \sqrt{\log\left(\frac{c\sqrt{d}}{\lambda\mathsf{K}^{-1}(\mathsf{K}(0)-\varepsilon^2/2)}\right)}$, then $\varepsilon = \sqrt{2\left(\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})\right)}$ and

$$\begin{aligned}
\mathbb{E}\left[\sup_{x\in D} u(x)\right] &\asymp \sqrt{d}\int_0^\infty -t\,\frac{d}{dt}\left(\sqrt{\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\right)\ dt \\
&= \sqrt{d}\left(-t\sqrt{\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\,\Big|_0^\infty + \int_0^\infty \sqrt{\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\ dt\right) \\
&= \sqrt{d}\int_0^\infty \sqrt{\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\ dt \\
&= \sqrt{d}\left[\int_{t<\sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}} + \int_{t>\sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}}\right] \sqrt{\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\ dt =: I_1 + I_2,
\end{aligned}$$

where in the second to last equality we used that $\mathsf{K}(0)-\mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2}) \asymp c\lambda^{-1}\sqrt{d}e^{-t^2}$ as $t\to\infty$ since $\mathsf{K}(r)$ is assumed to be differentiable at $r=0$. Further, we let $s>0$ be the unique solution of $\mathsf{K}(s)=\frac{1}{2}\mathsf{K}(0)$, which is independent of λ. For the first term I_1, we have

$$\sqrt{\frac{\mathsf{K}(0)d}{2}}\sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)} \le I_1 \le \sqrt{\mathsf{K}(0)d}\sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}.$$

Therefore, for any $\lambda < c\sqrt{d}/s$, $I_1 \asymp \sqrt{\mathsf{K}(0) d \log\left(\frac{\sqrt{d}}{s\lambda}\right)}$. To bound the second term I_2, we notice that there is some constant $M > 0$ such that $\mathsf{K}(0) - \mathsf{K}(r) \le M\, r$ for $r \in [0, s]$, where M is independent of λ. Therefore,

$$
\begin{aligned}
I_2 &= \sqrt{d} \int_{t > \sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}} \sqrt{\mathsf{K}(0) - \mathsf{K}(c\lambda^{-1}\sqrt{d}e^{-t^2})}\, dt \le \sqrt{d} \int_{t > \sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}} \sqrt{Mc\lambda^{-1}\sqrt{d}e^{-t^2}}\, dt \\
&\lesssim d^{3/4}\lambda^{-1/2} \int_{t > \sqrt{\log\left(\frac{c\sqrt{d}}{s\lambda}\right)}} e^{-\frac{1}{2}t^2}\, dt \lesssim \sqrt{d}\left(\log\left(\frac{c\sqrt{d}}{s\lambda}\right)\right)^{-1/2} \to 0, \quad \lambda \to 0,
\end{aligned}
$$

where we used the tail bound of the Gaussian distribution $\int_x^\infty e^{-\frac{1}{2}t^2} dt \le \frac{1}{x} e^{-\frac{1}{2}x^2}$ for $x > 0$. Since $I_2 \ge 0$, we therefore have that $I_2 \lesssim \sqrt{d}\left(\log\left(\frac{c\sqrt{d}}{s\lambda}\right)\right)^{-1/2} \to 0$ as $\lambda \to 0$. Consequently,

$$
\mathbb{E}\left[\sup_{x \in D} u(x)\right] \asymp I_1 + I_2 \asymp \sqrt{\mathsf{K}(0) d \log\left(\frac{\sqrt{d}}{s\lambda}\right)}, \quad \lambda \to 0. \qquad \square
$$

Remark 4.4.4. *Lemma 4.4.3 admits a clear heuristic interpretation. Consider a uniform mesh $\mathcal{P}$ of the unit cube $D = [0,1]^d$ comprising $(1/\lambda)^d$ points that are distance λ apart. For a random field $u(x)$ with lengthscale λ, the values $u(x_i)$ and $u(x_j)$ at mesh points $x_i \neq x_j \in \mathcal{P}$ are roughly uncorrelated. Thus, $\{u(x_i)\}_{i=1}^{\lambda^{-d}}$ are roughly i.i.d. univariate Gaussian random variables, and, for small λ, we may approximate*

$$
\mathbb{E}\left[\sup_{x \in D} u(x)\right] \approx \mathbb{E}\left[\sup_{x_i \in \mathcal{P}} u(x_i)\right] \approx \sqrt{\log(\lambda^{-d})}.
$$

This heuristic derivation matches the scaling of the expected supremum with λ in Lemma 4.4.3.

Proof of Theorem 4.2.8. In this proof we treat d as a constant. Notice that under Assumption 4.2.1 and Assumption 4.2.7, it holds that $\mathrm{Tr}(\mathcal{C}) = \int_D k(x,x)\, dx = \mathsf{K}(0)\mathrm{Vol}(D) = 1$. For the thresholded estimator, we apply Theorem 4.2.2 with an appropriate choice of the constant $c_0 \in [1, \sqrt{N}]$. By Lemma 4.4.3, $\mathbb{E}[\sup_{x\in D} u(x)] \asymp \sqrt{\log(\lambda^{-d})}$ as $\lambda \to 0$. We assume that $N \geq c_0^2\,(\mathbb{E}[\sup_{x\in D} u(x)])^2 \asymp \log(\lambda^{-d})$, so that the thresholding parameter satisfies

$$\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}\Big[\sup_{x\in D} u(x)\Big] \leq 1.$$

It follows that

$$\begin{aligned}\rho_N e^{-cN(\rho_N\wedge\rho_N^2)} = \rho_N e^{-cN\rho_N^2} &= \rho_N e^{-cc_0^2(\mathbb{E}[\sup_{x\in D} u(x)])^2} \\ &= \rho_N e^{-cc'c_0^2 d\log(1/\lambda)} = \rho_N \lambda^{cc'c_0^2 d} \leq \rho_N^{1-q}\lambda^{cc'c_0^2 d},\end{aligned} \tag{4.27}$$

where c' is an absolute constant. On the other hand, using Lemma 4.4.1 we have that

$$R_q^q \rho_N^{1-q} \asymp \rho_N^{1-q}\lambda^d A(d)\int_0^\infty \mathsf{K}(r)^q r^{d-1}dr. \tag{4.28}$$

Comparing (4.27) with (4.28), we see that if c_0 is chosen so that $cc'c_0^2 > 1$, then the upper bound $R_q^q\rho_N^{1-q} + \rho_N e^{-cN(\rho_N\wedge\rho_N^2)}$ in Theorem 4.2.2 is dominated by $R_q^q\rho_N^{1-q}$ as $\lambda \to 0$. Thus, for sufficiently small λ,

$$\mathbb{E}\|\widehat{\mathcal{C}}_{\widehat{\rho}_N} - \mathcal{C}\| \lesssim R_q^q\rho_N^{1-q} \leq \|\mathcal{C}\|\, c(q)\left(\frac{\log(\lambda^{-d})}{N}\right)^{\frac{1-q}{2}},$$

where $c(q)$ is a constant that only depends on q. □

4.5 Application in Ensemble Kalman Filters

Proof of Theorem 4.2.10. First, we write

$$|\upsilon_n - \upsilon_n^\star| = |(\mathscr{K}(\widehat{\mathcal{C}}) - \mathscr{K}(\mathcal{C}))(y - \mathcal{A}u_n - \eta_n)| \leq \|\mathscr{K}(\widehat{\mathcal{C}}) - \mathscr{K}(\mathcal{C})\||y - \mathcal{A}u_n - \eta_n|. \tag{4.29}$$

For the first term in (4.29), it follows by the continuity of the Kalman gain operator [Kwiatkowski and Mandel, 2015, Lemma 4.1] that

$$\|\mathscr{K}(\widehat{\mathcal{C}}) - \mathscr{K}(\mathcal{C})\| \le \|\widehat{\mathcal{C}} - \mathcal{C}\| \|\mathcal{A}\| \|\Gamma^{-1}\| \left(1 + \|\mathcal{C}\| \|\mathcal{A}\|^2 \|\Gamma^{-1}\|\right). \tag{4.30}$$

Combining the inequalities (4.29) and (4.30) with Theorem 4.2.8 gives that

$$\mathbb{E}\left[|\upsilon_n - \upsilon_n^\star| \mid u_n, \eta_n\right] \lesssim \|\mathcal{A}\| \|\Gamma^{-1}\| |y - \mathcal{A}u_n - \eta_n| \, \mathbb{E}\|\widehat{\mathcal{C}} - \mathcal{C}\| \lesssim c\left(\sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N}\right),$$

where $c = \|\mathcal{A}\| \|\Gamma^{-1}\| \|\mathcal{C}\| |y - \mathcal{A}u_n - \eta_n|$. Applying the same argument to the perturbed observation EnKF update with localization, υ_n^ρ, Theorem 4.2.8 gives that

$$\mathbb{E}\left[|\upsilon_n^\rho - \upsilon_n^\star| \mid u_n, \eta_n\right] \lesssim c\left[c(q)\left(\frac{\log(\lambda^{-d})}{N}\right)^{\frac{1-q}{2}}\right],$$

where $c = \|\mathcal{A}\| \|\Gamma^{-1}\| \|\mathcal{C}\| |y - \mathcal{A}u_n - \eta_n|$ and $c(q)$ is a constant that depends only on q. □

4.6 Conclusions, Discussion, and Future Directions

This chapter has studied thresholded estimation of sparse covariance operators, lifting the theory of sparse covariance matrix estimation from finite to infinite dimension. We have established non-asymptotic bounds on the estimation error in terms of the sparsity level of the covariance and the expected supremum of the field. In the challenging regime where the correlation lengthscale is small, we have shown that estimation via thresholding achieves an exponential improvement in sample complexity over the standard sample covariance estimator. As an application of the theory, we have demonstrated the advantage of using thresholded covariance estimators within ensemble Kalman filters. While our focus has been on studying the statistical benefit of estimation via thresholding, sparsifying the covariance estimator can also lead to significant computational speed-up in downstream tasks Furrer et al. [2006], Chen and Stein [2023], Chen and Anitescu [2023].

As mentioned in the discussion of Theorem 4.2.8, a natural question is whether the convergence rate of our thresholded estimator is minimax optimal. For ℓ_q-sparse covariance matrix estimation, Cai and Zhou [2012b] established the minimax optimality of thresholded estimators. Inspired by the correspondence between our error bound (4.5) and their optimal rate, we conjecture that our thresholded estimator is also minimax optimal in the infinite-dimensional setting.

Another interesting future direction is to relax the assumption of stationarity in our analysis of the small lengthscale regime. In finite dimension, Cai and Liu [2011] proposed adaptive thresholding estimators for sparse covariance matrix estimation that account for variability across individual entries and designed a data-driven choice of the prefactor c_0 through cross-validation. Other interesting extensions include covariance operator estimation for heavy-tailed distributions Abdalla and Zhivotovskiy [2024] and robust covariance operator estimation Goes et al. [2020], Diakonikolas and Kane [2023]. Finally, connections with the thriving topics of infinite-dimensional regression Mollenhauer et al. [2022] and operator learning de Hoop et al. [2023], Jin et al. [2022] will be explored in future work.

4.7 Proof of Lemma 3.4

This section contains the proof of Lemma 3.4. We will use the following auxiliary result, which can be found in [Talagrand, 2022, Lemma 2.10.6].

Lemma 4.7.1. *Under Assumption 2.1 (i), it holds with probability at least* $1-2e^{-t}$ *that*

$$\left| \frac{1}{N} \sum_{n=1}^{N} \sup_{x\in D} u_n(x) - \mathbb{E}\left[\sup_{x\in D} u(x) \right] \right| \le \sqrt{\frac{2t}{N}}.$$

Proof. By Gaussian concentration, $\sup_{x\in D} u(x)$ is $\sup_{x\in D} \operatorname{Var}[u(x)]$-sub-Gaussian. Since under Assumption 2.1 (i), $\sup_{x\in D} \operatorname{Var}[u(x)] = 1$, a Chernoff bound argument gives the result. $\square$

Proof of Lemma 3.4. We first prove (A). Without loss of generality, we assume $c_0 = 1$ in the definition of $\widehat{\rho}_N$ and ρ_N in Theorem 2.2. Let $t > 0$ and define $\mathcal{E}_t$ to be the event on which $\left|\frac{1}{N}\sum_{n=1}^N \sup_{x\in D} u_n(x) - \mathbb{E}\left[\sup_{x\in D} u(x)\right]\right| \le t$. It holds on $\mathcal{E}_t$ that

$$\begin{aligned}
\widehat{\rho}_N &\le \frac{1}{N} \vee \frac{\mathbb{E}[\sup_{x\in D} u(x)] + t}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D} u(x)] + t)^2}{N} \\
&\le \frac{1}{N} \vee \frac{2\mathbb{E}[\sup_{x\in D} u(x)]}{\sqrt{N}} \vee \frac{2t}{\sqrt{N}} \vee \frac{4(\mathbb{E}[\sup_{x\in D} u(x)])^2}{N} \vee \frac{4t^2}{N} \\
&\le 4\rho_N \vee \frac{2t}{\sqrt{N}} \vee \frac{4t^2}{N},
\end{aligned}$$

and $\mathbb{P}\left[\widehat{\rho}_N \le 4\rho_N \vee \frac{2t}{\sqrt{N}} \vee \frac{4t^2}{N}\right] \ge \mathbb{P}\big[\mathcal{E}_t\big] \ge 1 - 2e^{-Nt^2/2}$ by Lemma 4.7.1. It follows then that

$$\begin{aligned}
\mathbb{E}\big[\,\widehat{\rho}_N^{\,p}\big] = p\int_0^\infty t^{p-1}\,\mathbb{P}\big[\,\widehat{\rho}_N \ge t\big]dt &= p\int_0^{4\,\rho_N} t^{p-1}\,\mathbb{P}\big[\,\widehat{\rho}_N \ge t\big]dt + p\int_{4\,\rho_N}^\infty t^{p-1}\,\mathbb{P}\big[\,\widehat{\rho}_N \ge t\big]dt \\
&\le (4\rho_N)^p + 2p\int_{4\,\rho_N}^\infty t^{p-1}\, e^{-\frac{N}{2}\min\{\frac{Nt^2}{4},\frac{Nt}{4}\}}dt \lesssim_p \rho_N^p + \frac{1}{N^p} \lesssim_p \rho_N^p.
\end{aligned}$$

We next show (B). To prove (3.5), we can assume $c_0 = 1$ without loss of generality. Notice that

$$\begin{aligned}
&\mathbb{P}[\widehat{\rho}_N < t\rho_N] \\
&= \mathbb{P}\left[\Big(\frac{1}{N} < t\rho_N\Big)\bigcap\Big(\frac{1}{\sqrt{N}}\Big(\frac{1}{N}\sum_{n=1}^N \sup_{x\in D} u_n(x)\Big) < t\rho_N\Big)\bigcap\Big(\frac{1}{N}\Big(\frac{1}{N}\sum_{n=1}^N \sup_{x\in D} u_n(x)\Big)^2 < t\rho_N\Big)\right] \\
&= 1 - \mathbb{P}\left[\Big(\frac{1}{N} \ge t\rho_N\Big)\bigcup\Big(\frac{1}{\sqrt{N}}\Big(\frac{1}{N}\sum_{n=1}^N \sup_{x\in D} u_n(x)\Big) \ge t\rho_N\Big)\bigcup\Big(\frac{1}{N}\Big(\frac{1}{N}\sum_{n=1}^N \sup_{x\in D} u_n(x)\Big)^2 \ge t\rho_N\Big)\right].
\end{aligned}$$

We consider three cases.

Case 1: If $\mathbb{E}[\sup_{x\in D} u(x)] < \frac{1}{\sqrt{N}}$, then $\rho_N = \frac{1}{N}$ and $\mathbb{P}[\widehat{\rho}_N < t\rho_N] \le 1 - \mathbb{P}\Big[\frac{1}{N} \ge t\rho_N\Big] =$

0.

Case 2: If $\frac{1}{\sqrt{N}} \leq \mathbb{E}[\sup_{x\in D} u(x)] \leq \sqrt{N}$, then $\rho_N = \frac{1}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)]$ and

$$\begin{aligned}
\mathbb{P}[\widehat{\rho}_N < t\rho_N] &\leq 1 - \mathbb{P}\left[\frac{1}{\sqrt{N}}\Big(\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x)\Big) \geq t\rho_N\right] \\
&= 1 - \mathbb{P}\left[\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x) \geq t\,\mathbb{E}[\sup_{x\in D} u(x)]\right] \\
&\leq \mathbb{P}\left[\Big|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x) - \mathbb{E}[\sup_{x\in D} u(x)]\Big| \geq (1-t)\mathbb{E}[\sup_{x\in D} u(x)]\right] \\
&\leq 2\exp\Big(-\frac{1}{2}(1-t)^2 N(\mathbb{E}[\sup_{x\in D} u(x)])^2\Big),
\end{aligned}$$

where the last step follows by Lemma 4.7.1.

Case 3: If $\mathbb{E}[\sup_{x\in D} u(x)] > \sqrt{N}$, then $\rho_N = \frac{1}{N}(\mathbb{E}[\sup_{x\in D} u(x)])^2$ and

$$\begin{aligned}
\mathbb{P}[\widehat{\rho}_N < t\rho_N] &\leq 1 - \mathbb{P}\left[\frac{1}{N}\Big(\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x)\Big)^2 \geq t\rho_N\right] \\
&= 1 - \mathbb{P}\left[\Big|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x)\Big| \geq \sqrt{t}\,\mathbb{E}[\sup_{x\in D} u(x)]\right] \\
&\leq \mathbb{P}\left[\Big|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D} u_n(x) - \mathbb{E}[\sup_{x\in D} u(x)]\Big| \geq (1-\sqrt{t})\mathbb{E}[\sup_{x\in D} u(x)]\right] \\
&\leq 2\exp\Big(-\frac{1}{2}(1-\sqrt{t})^2 N(\mathbb{E}[\sup_{x\in D} u(x)])^2\Big).
\end{aligned}$$

Combining the three cases above and noticing that $(1-\sqrt{t})^2 \leq (1-t)^2$ for $t \in (0,1)$ yields the first inequality in (3.5). To prove (3.6), recall that $1 \leq c_0 \leq \sqrt{N}$ in the definition of ρ_N. If $\mathbb{E}[\sup_{x\in D} u(x)] < 1/\sqrt{N}$, then (3.6) is trivial. If $\frac{1}{\sqrt{N}} \leq \mathbb{E}[\sup_{x\in D} u(x)] \leq \sqrt{N}$, then $\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}[\sup_{x\in D} u(x)]$ and $N(\mathbb{E}[\sup_{x\in D} u(x)])^2 = \frac{N^2\rho_N^2}{c_0^2} \geq N\rho_N^2$, so that

$$2\, e^{-\frac{1}{2}(1-\sqrt{t})^2 N(\mathbb{E}[\sup_{x\in D} u(x)])^2} \mathbf{1}\big\{\mathbb{E}[\sup_{x\in D} u(x)] \geq 1/\sqrt{N}\big\} \leq 2\, e^{-\frac{1}{2}(1-\sqrt{t})^2 N\rho_N^2}.$$

If $\mathbb{E}[\sup_{x\in D} u(x)] > \sqrt{N}$, then $\rho_N = \frac{c_0}{N}(\mathbb{E}[\sup_{x\in D} u(x)])^2$ and $N(\mathbb{E}[\sup_{x\in D} u(x)])^2 = \frac{N^2\rho_N}{c_0} \geq N^{3/2}\rho_N \geq N\rho_N$, so that

$$2\, e^{-\frac{1}{2}(1-\sqrt{t})^2 N(\mathbb{E}[\sup_{x\in D} u(x)])^2} \mathbf{1}\big\{\mathbb{E}[\sup_{x\in D} u(x)] \geq 1/\sqrt{N}\big\} \leq 2\, e^{-\frac{1}{2}(1-\sqrt{t})^2 N\rho_N}. \qquad \square$$

4.8 Additional Results

4.8.1 Bound on Operator Norm

Lemma 4.8.1. *Let $D \subset \mathbb{R}^d$. For an integral operator K on $L^2(D)$,*

$$(K\psi)(x) := \int_D k(x,x')\psi(x')dx', \quad \psi \in L^2(D),$$

it holds that

$$\|K\|^2 \leq \left(\sup_x \int_D |k(x,x')|dx'\right)\left(\sup_{x'} \int_D |k(x,x')|dx\right).$$

Further, if $k(x,x') = k(x',x)$, then

$$\|K\| \leq \sup_x \int_D |k(x,x')|dx'.$$

Proof. For any $\psi \in L^2(D)$ with $\|\psi\|_{L^2(D)} = 1$,

$$\begin{aligned}
\|K\psi\|_{L^2(D)}^2 &= \int_D \Big(\int_D k(x,x')\psi(x')dx' \Big)^2 dx \\
&\leq \int_D \Big(\int_D |k(x,x')| \cdot |\psi(x')|dx' \Big)^2 dx \\
&= \int_D \Big(\int_D \sqrt{|k(x,x')|} \cdot \sqrt{|k(x,x')|}|\psi(x')|dx' \Big)^2 dx \\
&\overset{\text{(i)}}{\leq} \int_D \left(\int_D |k(x,x')|dx' \right) \left(\int_D |k(x,x')|\psi(x')^2 dx' \right) dx \\
&\leq \left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \int_D \int_D |k(x,x')|\psi(x')^2 dx' dx \\
&\leq \left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \int_D \left(\int_D |k(x,x')|dx \right) \psi(x')^2 dx' \\
&\leq \left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \left(\sup_{x'} \int_D |k(x,x')|dx \right) \cdot \left(\int_D \psi(x')^2 dx' \right) \\
&= \left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \left(\sup_{x'} \int_D |k(x,x')|dx \right),
\end{aligned}$$

where (i) follows by Cauchy-Schwarz inequality. Therefore,

$$\|K\|^2 = \sup_{\|\psi\|_{L^2(D)}=1} \|K\psi\|_{L^2(D)}^2 \leq \left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \left(\sup_{x'} \int_D |k(x,x')|dx \right).$$

Further, if $k(x,x') = k(x',x)$, then

$$\|K\| \leq \sqrt{\left(\sup_x \int_D |k(x,x')|dx' \right) \cdot \left(\sup_{x'} \int_D |k(x,x')|dx \right)} = \sup_x \int_D |k(x,x')|dx',$$

which completes the proof. $\square$

4.8.2 Auxiliary Technical Result

Lemma 4.8.2. *For any $\alpha > \beta > 0$ and $q > 0$, it holds that*

$$\int_0^{\alpha-\beta} e^{-\frac{1}{2}t^2}(\alpha - t)^{-q-1}\,dt \leq \frac{2\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{8}} + \frac{1}{q}\Big(\frac{\alpha}{2}\Big)^{-q}.$$

Proof. Integrating by parts gives that

$$\begin{aligned}\int_0^{\alpha-\beta} e^{-\frac{1}{2}t^2}(\alpha - t)^{-q-1}dt &= \frac{\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{2}} - \frac{\alpha^{-q}}{q} + \int_0^{\alpha-\beta} e^{-\frac{1}{2}t^2} t\frac{(\alpha-t)^{-q}}{q}dt \\ &= \frac{\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{2}} - \frac{\alpha^{-q}}{q} + \left(\int_0^{\frac{\alpha-\beta}{2}} + \int_{\frac{\alpha-\beta}{2}}^{\alpha-\beta}\right) e^{-\frac{1}{2}t^2} t\frac{(\alpha-t)^{-q}}{q}dt.\end{aligned}$$

First,

$$\int_0^{\frac{\alpha-\beta}{2}} e^{-\frac{1}{2}t^2} t\frac{(\alpha-t)^{-q}}{q}dt \leq \frac{1}{q}\Big(\alpha - \frac{\alpha-\beta}{2}\Big)^{-q}\int_0^{\frac{\alpha-\beta}{2}} e^{-\frac{1}{2}t^2} t\,dt \leq \frac{1}{q}\Big(\frac{\alpha+\beta}{2}\Big)^{-q} \leq \frac{1}{q}\Big(\frac{\alpha}{2}\Big)^{-q}.$$

Second,

$$\int_{\frac{\alpha-\beta}{2}}^{\alpha-\beta} e^{-\frac{1}{2}t^2} t\frac{(\alpha-t)^{-q}}{q}dt \leq \frac{1}{q}(\alpha - (\alpha-\beta))^{-q}\int_{\frac{\alpha-\beta}{2}}^{\alpha-\beta} e^{-\frac{1}{2}t^2} t\,dt \leq \frac{\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{8}}.$$

Thus,

$$\begin{aligned}\int_0^{\alpha-\beta} e^{-\frac{1}{2}t^2}(\alpha - t)^{-q-1}dt &\leq \frac{\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{2}} - \frac{\alpha^{-q}}{q} + \frac{\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{8}} + \frac{1}{q}\Big(\frac{\alpha}{2}\Big)^{-q} \\ &\leq \frac{2\beta^{-q}}{q} e^{-\frac{(\alpha-\beta)^2}{8}} + \frac{1}{q}\Big(\frac{\alpha}{2}\Big)^{-q}.\end{aligned}$$

□

CHAPTER 5

COVARIANCE OPERATOR ESTIMATION VIA ADAPTIVE THRESHOLDING

This chapter is adapted from the manuscript, which received a minor revision at Stochastic Processes and their Applications, listed below.

O. Al-Ghattas and D. Sanz-Alonso, *Covariance Operator Estimation via Adaptive Thresholding*, *arXiv preprint arXiv:2405.18562*, 2024.

5.1 Introduction

This paper investigates sparse covariance operator estimation in an infinite-dimensional function space setting. Covariance estimation is a fundamental task that arises in numerous scientific applications and data-driven algorithms Anderson [1958], Fan et al. [2008], Hardoon et al. [2004], Tharwat et al. [2017], Al-Ghattas and Sanz-Alonso [2024c], Al-Ghattas et al. [2024a]. The sample covariance is arguably the most natural estimator, and its error in both finite and infinite dimension can be controlled by a notion of *effective dimension* that accounts for spectrum decay Koltchinskii and Lounici [2017], Lounici [2014]. However, a rich literature has identified sparsity assumptions under which other estimators drastically outperform the sample covariance in finite high-dimensional settings Bickel and Levina [2008a,b], El Karoui [2008], Cai and Zhou [2012b], Cai et al. [2016], Wainwright [2019]. This work contributes to the largely unexplored subject of *sparse* covariance operator estimation in infinite dimension. Through rigorous theory and complementary numerical simulations, we demonstrate the benefit of adaptively thresholding the sample covariance. In doing so, this paper contributes to the emerging literature on operator estimation and learning Kovachki et al. [2024], de Hoop et al. [2023], Mollenhauer et al. [2022], emphasizing the importance of exploiting structural assumptions in the design and analysis of estimators.

In this work, we investigate approximate sparsity structure that arises in the nonstationary regime where the marginal variance varies widely in the domain and the correlation lengthscale is small relative to the size of the domain. Covariance estimation for such nonstationary processes is crucial, for instance, in numerical weather forecasting, where local and highly nonstationary phenomena such as cloud formation can significantly impact mid and long-range global forecasts. To study the sparse highly nonstationary regime where the marginal variance varies widely in the domain and the correlation lengthscale is small, we consider a novel class of covariance operators that satisfy a *weighted* L_q-sparsity condition. For covariance operators in this class, we establish a bound on the operator norm error of the *adaptive threshold* estimator in terms of two dimension-free quantities: the sparsity level and the expected supremum of a normalized field. Unlike existing theory that considered unweighted L_q-sparsity (see Section 5.1.1 for a review) our theory allows for covariance models with unbounded marginal variance functions. We then compare our adaptive threshold estimator with other estimators of interest, namely the *universal threshold* and *sample covariance* estimators. For universal thresholding, we prove a lower bound that is larger than our upper bound for adaptive thresholding. In addition, we numerically investigate adaptive thresholding for highly nonstationary covariance models defined through a scalar parameter that controls both the correlation lengthscale and the range of the marginal variance function. In the challenging case where the lengthscale is small and the range of marginal variances is large, we show an exponential improvement in sample complexity of the adaptive threshold estimator compared to the sample covariance. Our numerical simulations clearly demonstrate that universal threshold and sample covariance estimators fail in this regime.

By focusing on the infinite-dimensional setting, our theory reveals the key dimension-free quantities that control the estimation error, and further explains how the correlation lengthscale and the marginal variance function affect the estimation problem. While our infinite-dimensional analysis helps uncover such a connection between interpretable model assumptions and complexity of the estimation task, it poses new challenges that require novel technical tools. In this work, we leverage recent results from empirical process theory Mendelson [2016] to control the error in the estimation of the variance component used to adaptively choose the thresholding radius. Our infinite-dimensional perspective agrees with recent work in operator learning that advocates for the development of theory in infinite dimension, as opposed to the traditional approach in functional data analysis, where it is common to study estimators constructed by first discretizing the data Ramsay and Ramsey [2002], Zhang and Wang [2016]. We will discuss the differences between the two approaches and compare our theory with the existing infinite-dimensional covariance estimation literature Al-Ghattas et al. [2023], Fang et al. [2023]. More broadly, infinite-dimensional analyses that delay introducing discretization have led to numerous theoretical insights and computational advances in mathematical statistics Giné and Nickl [2021], Bayesian inverse problems Stuart [2010], Markov chain Monte Carlo Cotter et al. [2013], importance sampling and particle filters Agapiou et al. [2017], ensemble Kalman algorithms Sanz-Alonso and Waniorek [2024], graph-based learning García Trillos and Sanz-Alonso [2018a], stochastic gradient descent Latz [2021], and numerical analysis and control Zuazua [2005], among many others.

5.1.1 Related Work

For later reference and discussion, here we summarize unweighted and weighted approximate sparsity assumptions in the finite-dimensional thresholded covariance estimation literature, as well as the main sparsity assumptions that have been considered in the infinite-dimensional setting.

Finite Dimension

Thresholding estimators in the finite high-dimensional setting were introduced in the seminal work Bickel and Levina [2008a] and further studied in Rothman et al. [2009], Cai and

Zhou [2012b], Cai et al. [2016], Cai and Liu [2011]. Given d_X-dimensional i.i.d. samples $X_1, \ldots, X_N$ from a centered sub-Gaussian distribution with covariance Σ, the authors demonstrated that thresholding the sample covariance matrix, i.e. $\widehat{\Sigma}^{\mathsf{U}}_{\rho_N} = (\widehat{\Sigma}_{ij}\mathbf{1}\{|\widehat{\Sigma}_{ij}| \geq \rho_N\})$ — where the superscript U denotes *universal* — performed well over the class of covariance matrices with bounded marginal variances and satisfying an ℓ_q-sparsity condition whenever the thresholding parameter ρ_N is chosen appropriately. Specifically, whenever Σ belongs to the class $\mathcal{U}_q(R_q, M)$ for $q \in [0, 1)$, $R_q > 0$ and $M > 0$ where

$$\mathcal{U}_q := \mathcal{U}_q(R_q, M) = \left\{ \Sigma \in \mathbb{R}^{d_X \times d_X} : \Sigma \succ 0, \ \max_{i \leq d_X} \Sigma_{ii} \leq M, \ \max_{i \leq d_X} \sum_{j=1}^{d_X} |\Sigma_{ij}|^q \leq R_q^q \right\}, \quad (5.1)$$

then the operator norm error of universal thresholding estimators is bounded above (up to universal constants) by $R_q^q(M \log d_X/N)^{(1-q)/2}$. The bounded marginal variance assumption is crucial to this theory as it was shown that ρ_N must scale with M in order for the high probability guarantees on $\widehat{\Sigma}^{\mathsf{U}}_{\rho_N}$ to hold. In Cai and Liu [2011], the authors argued that such a bounded variance assumption effectively converted a heteroscedastic problem of covariance estimation into a worst-case homoscedastic one in which $\Sigma_{ii} = M$ for all i for the purposes of choosing a universal thresholding radius. This is problematic whenever (i) no natural upper bound on the marginal variances is known and (ii) the marginal variances vary over a large range. They instead considered the adaptively thresholded covariance estimator $\widehat{\Sigma}^{A}_{\rho_N} = (\widehat{\Sigma}_{ij}\mathbf{1}\{|\widehat{\Sigma}_{ij}| \geq \rho_N \hat{V}_{ij}^{1/2}\})$, where $\hat{V}_{ij}$ is the sample version of the variance component $V_{ij} := \mathrm{var}(X_i X_j)$. This was shown to be optimal over the larger weighted ℓ_q-sparsity covariance matrix class

$$\mathcal{U}_q^* := \mathcal{U}_q^*(R_q) = \left\{ \Sigma \in \mathbb{R}^{d_X \times d_X} : \Sigma \succ 0, \ \max_{i \leq d_X} \sum_{j=1}^{d_X} (\Sigma_{ii}\Sigma_{jj})^{\frac{1-q}{2}} |\Sigma_{ij}|^q \leq R_q^q \right\}, \quad (5.2)$$

with operator norm error bounded above (up to universal constants) by $R_q^q(\log d_X/N)^{(1-q)/2}$. It was further shown that universal thresholding was sub-optimal over the same class. Minimax lower bounds proving the optimality of universal thresholding were studied in Cai and Zhou [2012b]. Distributional assumptions were significantly relaxed by allowing for dependence in Chen et al. [2013], which analyzed thresholding in the high-dimensional time series setting.

Infinite Dimension

In the infinite-dimensional setting, the covariance (operator) estimation problem under sparsity-type constraints has received far less attention. In Al-Ghattas et al. [2023], the authors consider i.i.d. draws of an infinite-dimensional Gaussian process defined over $D = [0,1]^d$ with covariance operator C and corresponding covariance function $k : D \times D \to \mathbb{R}$, denoted $u_1, \ldots, u_N \overset{\text{i.i.d.}}{\sim} \mathrm{GP}(0, C)$. They generalize Bickel and Levina [2008b] to infinite dimensions by considering Gaussian processes that are almost surely continuous with covariance operators $C \in \mathcal{K}_q$, where

$$\mathcal{K}_q := \mathcal{K}_q(R_q, M) = \left\{ C \succ 0 : \sup_{x\in D} k(x,x) \le M, \ \sup_{x \in D} \int_D |k(x,y)|^q dy \le R_q^q \right\}. \qquad (5.3)$$

This class naturally captures approximate sparsity of the covariance, which may arise, for example, from decay of correlations of the Gaussian process at different locations in the domain. It was then shown that for a universally thresholded covariance operator estimator, i.e. with covariance function $\hat{k}(x,y)\mathbf{1}\{|\hat{k}(x,y)| \ge \rho_N\}$, the operator norm error was bounded above by $R_q^q((\mathbb{E}[\sup_{x\in D} u(x)])^2/N)^{(1-q)/2}$, which is a dimension-free quantity. Further, the authors demonstrate that if the covariance function is stationary and depends on a lengthscale parameter λ, then universally thresholded estimators enjoy an exponential improvement over the standard sample covariance estimator in the small λ asymptotic. Notice that here and throughout this paper, d represents the dimension of the physical domain $D = [0,1]^d$ and should not be confused with the dimension of the data points $\{u_n\}_{n=1}^N$, which here represent infinite-dimensional functions.

Motivated by applications in functional data analysis, Fang et al. [2023] considers covariance estimation for a multi-valued process $\mathbf{u} : D \to \mathbb{R}^p$ given independent data

$$\mathbf{u}_n(\cdot) = \big(u_{n1}(\cdot), u_{n2}(\cdot), \ldots, u_{np}(\cdot)\big)^\top, \quad n = 1, \ldots, N.$$

The covariance function now takes the form

$$\mathbf{K} : D \times D \to \mathbb{R}^{p \times p}, \qquad \mathbf{K}(x,y) = \mathrm{Cov}\big(\mathbf{u}_n(x), \mathbf{u}_n(y)\big) = [k_{ij}(x,y)]_{i,j=1}^p,$$

where $k_{ij} : D \times D \to \mathbb{R}$ is a component covariance function. Then, Fang et al. [2023] studies the setting in which the number of component is much larger than the sample size, i.e. $p \gg N$, and under the assumption that the true covariance function belongs to the class

$$\mathcal{G}_q(R_q, \varepsilon) = \left\{ \mathbf{K} \succeq 0 : \max_{i \le p} \sum_{j=1}^p (\|k_{ii}\|_\infty \|k_{jj}\|_\infty)^{\frac{1-q}{2}} \|k_{ij}\|_{\mathrm{HS}}^q \le R_q^q, \max_{i \le p} \|k_{ii}^{-1}\|_\infty \|k_{ii}\|_\infty \le \frac{1}{\varepsilon} \right\}. \tag{5.4}$$

Here, we denote by $\|k\|_{\mathrm{HS}}^2 = \iint k^2(x,y)\, dxdy$ the Hilbert-Schmidt norm, and we denote $\|k\|_\infty = \sup_{x,y \in D} |k(x,y)|$. This class generalizes the class $\mathcal{U}_q^*$ and the authors obtain analogous upper bounds to those of Cai and Liu [2011] for the error of estimation under a functional version of the matrix ℓ_1-norm. We provide further comparisons to this line of work in Remark 5.2.2. Another popular approach in the functional data analysis literature is the partial observations framework, see Yao et al. [2005a], Zhang and Wang [2016] and [Fang et al., 2023, Section 4]. In this setting, observations are comprised of noisy evaluations of the infinite dimensional response function at a set of grid-points located randomly in the domain D. At a high level, much of this literature involves the study of nonparametric estimators (e.g. local polynomials) to recover estimates of the functions underlying the partial observations. These estimates are then used as inputs to the sample covariance estimator. Under general smoothness assumptions, which are necessary to control the bias of the nonparametric estimator, it can be shown that covariance estimators that use these estimated functions are asymptotically equivalent to covariance estimators that use fully observed functional data. We discuss this approach further in Remark 5.2.5.

5.1.2 Outline

Section 5.2 contains the main results of this paper: Theorem 5.2.3 shows an operator norm bound for adaptive threshold estimators, Theorem 5.2.6 states a lower bound for universal thresholding, and Theorem 5.2.10 compares the sample covariance and adaptive threshold estimators. In addition, Section 5.2 also includes numerical simulations in physical dimension $d = 1$; similar results in dimension $d = 2$ are deferred to an appendix. The proof of Theorem 5.2.3 can be found in Section 5.3, and uses a recent result on empirical process theory discussed in Section 5.4. Sections 5.5 and 5.6 contain the proofs of Theorems 5.2.6 and 5.2.10, respectively. The paper closes in Section 5.7 with concluding remarks and suggestions for future work.

Notation Given two positive sequences $\{a_n\}$ and $\{b_n\}$, the relation $a_n \lesssim b_n$ denotes that $a_n \leq cb_n$ for some constant $c > 0$. If both $a_n \lesssim b_n$ and $b_n \lesssim a_n$ hold simultaneously, we write $a_n \asymp b_n$. For an operator $\mathcal{L}$, we denote its operator norm by $\|\mathcal{L}\|$ and its trace by $\text{Tr}(\mathcal{L})$. For a matrix $\Sigma \in \mathbb{R}^{p\times p}$ (resp. operator $C : L_2(D) \to L_2(D)$) we write $\Sigma \succ 0$ (resp. $C \succ 0$) to denote that Σ (resp. C) is positive definite.

5.2 Main Results

This section introduces our framework, assumptions, and main results. In Section 5.2.1, we discuss the data generating mechanism under consideration, and we define the various estimators that will be studied. In Section 5.2.2, we establish our main result, a high probability operator norm error bound for the adaptive threshold estimator. In Section 5.2.3, we provide both theoretical and empirical comparisons of adaptive threshold, universal threshold, and sample covariance estimators.

5.2.1 Setting and Estimators

Let $D \subset \mathbb{R}^d$ and let $u_1, \ldots, u_N$ be i.i.d. copies of a centered square-integrable random field $u : D \to \mathbb{R}$. We are interested in estimating the covariance operator C from the data $\{u_n\}_{n=1}^N$. Recall that the covariance function (kernel) $k : D \times D \to \mathbb{R}$ and covariance operator $C : L_2(D) \to L_2(D)$ are defined by the requirement that, for any $x, y \in D$ and $\psi \in L_2(D)$,

$$k(x, y) := \mathbb{E}[u(x)u(y)], \qquad (C\psi)(\cdot) := \int_D k(\cdot, y)\psi(y)\, dy.$$

That is, C is the integral operator with kernel k. We will focus on (sub-)Gaussian data. Recall that a square-integrable process u is called (sub-)Gaussian if, for any fixed $w \in L_2(D)$, the random variable $\langle u, w\rangle_{L_2(D)}$ is (sub-)Gaussian. We further recall that the process u is called pre-Gaussian if there exists a centered Gaussian process, v, with the same covariance operator as that of u. Following [Ledoux and Talagrand, 2013, page 261], we refer to v as the Gaussian process *associated to* u.

For simplicity, we take $D := [0, 1]^d$ to be the d-dimensional unit hypercube. In the applications that motivate this work, the ambient dimension d is typically $1, 2$, or 3, and so we treat d as a constant throughout. We are interested in applications where the covariance function k exhibits approximate sparsity (which may arise, for instance, due to spatial decay of correlations), and where the marginal variance function $\sigma^2(x) := k(x,x)$ has multiple scales in the domain D. In this setting, the sample covariance function $\hat{k}$ and sample covariance operator $\widehat{C}$ defined by

$$\hat{k}(x,y) := \frac{1}{N}\sum_{n=1}^{N} u_n(x)u_n(y), \qquad (\widehat{C}\psi)(\cdot) := \int_D \hat{k}(\cdot,y)\psi(y)\,dy$$

perform poorly. To improve performance in regard to exploiting approximate sparsity, one can instead consider the universal threshold estimator defined by

$$\hat{k}^{\mathsf{U}}_{\rho_N}(x,y) := \hat{k}(x,y)\mathbf{1}\left\{|\hat{k}(x,y)| \geq \rho_N\right\}, \qquad (\widehat{C}^{\mathsf{U}}_{\rho_N}\psi)(\cdot) := \int_D \hat{k}^{\mathsf{U}}_{\rho_N}(\cdot,y)\psi(y)\,dy,$$

where ρ_N is a tunable thresholding parameter. However, this approach is not well suited if the marginal variance function takes a wide range of values on D, where it becomes essential to consider a spatially varying thresholding parameter. To that end, we define the variance component $\theta : D \times D \to \mathbb{R}_{\geq 0}$ by

$$\theta(x,y) := \operatorname{var}\big(u(x)u(y)\big),$$

To estimate the variance component, we consider the standard *sample-based* estimator given by

$$\hat{\theta}_{\mathsf{S}}(x,y) := \frac{1}{N}\sum_{n=1}^{N} u_n^2(x)u_n^2(y) - \hat{k}^2(x,y).$$

In the Gaussian setting, we additionally consider the *Wick's-based* estimator given by

$$\hat{\theta}_{\mathsf{W}}(x,y) := \hat{k}(x,x)\hat{k}(y,y) + \hat{k}^2(x,y),$$

which is motivated by the following derivation invoking Wick's theorem (also commonly referred to as Isserlis' theorem)

$$\begin{aligned}\theta(x,y) &= \mathbb{E}\big[u^2(x)u^2(y)\big] - \big(\mathbb{E}\big[u(x)u(y)\big]\big)^2 \\ &= \mathbb{E}\big[u^2(x)\big]\mathbb{E}\big[u^2(y)\big] + 2\big(\mathbb{E}\big[u(x)u(y)\big]\big)^2 - \big(\mathbb{E}\big[u(x)u(y)\big]\big)^2 \\ &= k(x,x)k(y,y) + k^2(x,y).\end{aligned}$$

Given an estimator $\hat{\theta}$ of θ, we then define the adaptive threshold estimator

$$\hat{k}^{\mathsf{A}}_{\rho_N} := \hat{k}(x,y)\mathbf{1}\left\{\left|\frac{\hat{k}(x,y)}{\sqrt{\hat{\theta}(x,y)}}\right| \ge \rho_N\right\}, \qquad (\widehat{C}^{\mathsf{A}}_{\rho_N}\psi)(\cdot) := \int_D \hat{k}^{\mathsf{A}}_{\rho_N}(\cdot,y)\psi(y)\,dy,$$

where we set $\mathsf{A} = \mathsf{S}$ when $\hat{\theta} = \hat{\theta}_{\mathsf{S}}$ and $\mathsf{A} = \mathsf{W}$ when $\hat{\theta} = \hat{\theta}_{\mathsf{W}}$. We refer to this as adaptive thresholding (of the sample covariance) since the event in the indicator can be equivalently written as $\{|\hat{k}(x,y)| \ge \rho_N \hat{\theta}^{1/2}(x,y)\}$, and so the level of thresholding varies with the location $(x,y) \in D \times D$. The goal of this paper is to demonstrate through rigorous theory and numerical examples the improved performance of the adaptive threshold estimator relative to the universal threshold and sample covariance estimators.

5.2.2 Error Bound for Adaptive-threshold Estimator

Our theory is developed under the following assumption:

Assumption 5.2.1. *Let $u, u_1, \dots, u_N$ be i.i.d. centered sub-Gaussian and pre-Gaussian random functions on $D = [0,1]^d$ that are Lebesgue almost-everywhere continuous with probability one. It holds that:*

(i) $C \in \mathcal{K}^*_q$ where

$$\mathcal{K}^*_q := \mathcal{K}^*_q(R_q) = \left\{C \succ 0, \sup_{x\in D}\int_D \big(k(x,x)k(y,y)\big)^{\frac{1-q}{2}} |k(x,y)|^q\,dy \le R_q^q\right\}.$$

(ii) There exists a universal constant $\nu > 0$ such that, for any $x, y \in D$,

$$\theta(x, y) \geq \nu k(x, x) k(y, y).$$

(iii) For a sufficiently small constant c, the sample size satisfies

$$\sqrt{N} \geq \frac{1}{c} \mathbb{E} \left[\sup_{x \in D} \frac{u(x)}{\sqrt{k(x, x)}} \right].$$

In contrast to the setting considered in Al-Ghattas et al. [2023], Assumption 5.2.1 allows for sub-Gaussian data and admits covariance functions for which $\sup_{x \in D} k(x, x) \to \infty$. Furthermore, here we only require *Lebesgue almost-everywhere* continuity of the data, whereas Al-Ghattas et al. [2023] requires continuous data. Assumption 5.2.1 (i) specifies that the covariance function k satisfies a weighted L_q-sparsity condition that generalizes the class of row-sparse matrices $\mathcal{U}_q^*(R_q)$ studied in Cai and Liu [2011] to our infinite-dimensional setting.

Assumption 5.2.1 (ii) ensures that consistent estimation of the variance component is possible, and is analogous to requirements in finite dimension [Cai and Liu, 2011, Condition C1]. Assumption 5.2.1 (iii) is imposed for purely cosmetic reasons and can be removed at the expense of a more cumbersome statement of the results and proofs that would need to account for the case in which the sample size is chosen to be insufficiently large. The sample size requirement can also be compared to [Fang et al., 2023, Condition 4], which requires that the pair (N, p) satisfies $\log p / N^{1/4} \to 0$ as $N, p \to \infty$, where p is the number of component random functions as described in Section 5.1.1. In contrast, our assumption is nonasymptotic and stated only in terms of the dimension-free quantity $\mathbb{E}[\sup_{x \in D} u(x)/\sqrt{k(x, x)}]$ as our proof techniques differ from theirs.

Remark 5.2.2 (Comparison to Global-type Sparse Class of Fang et al. [2023])**.** *As noted in Section 5.1.1, Fang et al. [2023] recently studied a notion of sparse covariance functions*

$\mathcal{G}_q(R_q, \varepsilon)$ different to the one considered here. In particular, the class $\mathcal{G}_q(R_q, \varepsilon)$ specifies a global notion of sparsity, in the sense that it imposes conditions on the relationship between the various coordinate covariance functions that make-up $\boldsymbol{K}$. In contrast, the class $\mathcal{K}_q^*$ is local, in that it imposes a sparse structure on the covariance function of a real-valued process. For example, in the single component case $p = 1$, covariance functions k_{11} belonging to $\mathcal{G}_q(R_q, \varepsilon)$ must satisfy $\|k_{11}\|_\infty^{1-q}\|k_{11}\|_{HS}^q \le R_q^q$ and $\|k_{11}^{-1}\|_\infty\|k_{11}\|_\infty \le 1/\varepsilon$. This is equivalent to requiring that the covariance function k_{11} is bounded in Hilbert-Schmidt norm, and that k_{11} and its inverse are bounded in supremum norm. Importantly, in contrast to the class $\mathcal{K}_q^*$ studied here, their assumption does not capture any decay of correlations of the process at two different points in the domain, nor does it permit $\sup_{x\in D} k_{11}(x, x) \to \infty$.

We are now ready to state our main result, which establishes operator norm bounds for adaptive threshold estimators.

Theorem 5.2.3. *Under Assumption 5.2.1, let v be the Gaussian process associated to u, and for a universal constant $c_0 > 0$, let*

$$\rho_N = \frac{c_0}{\sqrt{N}}\mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right].$$

Then, there exists a universal constant $c_1 > 0$ such that, with probability at least $1-c_1 e^{-N\rho_N^2}$,

$$\|\widehat{C}_{\rho_N}^{\mathsf{S}} - C\| \lesssim R_q^q \rho_N^{1-q}.$$

Moreover, if u is a Gaussian process, then for

$$\hat{\rho}_N = \frac{c_0}{\sqrt{N}}\left(\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{u_n(x)}{\hat{k}^{1/2}(x,x)}\right),$$

there exists a universal constant $c_1 > 0$ *such that, with probability at least* $1 - c_1 e^{-N\rho_N^2}$,

$$\|\widehat{C}^{\mathsf{A}}_{\hat{\rho}_N} - C\| \lesssim R_q^q \rho_N^{1-q},$$

where $\mathsf{A} \in \{\mathsf{S}, \mathsf{W}\}$.

Remark 5.2.4. *For adaptive covariance matrix estimation under weighted* ℓ_q*-sparsity described in Section 5.1.1, [Cai and Liu, 2011, Theorem 1] showed that if the (normalized) sample covariance matrix satisfies* $\max_{1\le i,j\le d_X} |\widehat{\Sigma}_{ij} - \Sigma_{ij}|/\hat{V}_{ij}^{1/2} \lesssim \tilde{\rho}_N$*, then the operator norm error of the adaptive thresholding covariance matrix estimator can be bounded by* $\tilde{R}_q^q \tilde{\rho}_N^{1-q}$*, where* $\tilde{R}_q^q$ *controls the row-wise weighted* ℓ_q*-sparsity of* Σ*. The choice of thresholding parameter can be understood by appealing to the analogy that covariance matrix estimation may be interpreted as a heteroscedastic Gaussian sequence model (see [Cai and Liu, 2011, Section 2]), so that roughly speaking, for large* N,

$$\frac{1}{N}\sum_{n=1}^{N} X_{ni}X_{nj} \approx \Sigma_{ij} + \sqrt{\frac{V_{ij}}{N}} Z_{ij}, \qquad 1 \le i, j \le d_X,$$

with $\{Z_{ij}\}$ *i.i.d. standard normal. This explains the choice of the thresholding parameter in the finite-dimensional setting (after normalizing the data by* $\hat{V}_{ij}$*), since*

$$\tilde{\rho}_N \asymp \sqrt{\frac{\log d_X}{N}} \asymp \frac{\mathbb{E}\left[\max_{i,j\le d_X} Z_{ij}\right]}{\sqrt{N}},$$

provides element-wise control on the sample covariance. In the infinite-dimensional setting considered here, we require instead high probability sup-norm concentration bounds for the sample covariance function $\hat{k}(x, y)$ *and the estimated variance component function* $\hat{\theta}(x, y)$*. These bounds are obtained in Section 5.3 utilizing tools adapted from recent advances in the study of multi-product empirical processes Al-Ghattas et al. [2025] via generic chaining. These techniques are discussed in Section 5.4.2. Our results show that the correct thresholding* *radius ρ_N must scale with the expected supremum of the normalized associated Gaussian process, as this is precisely the dimension-free quantity needed to control the sample quantities uniformly over the domain D. Since ρ_N is a population level quantity, we establish in Lemma 5.3.6 that it may be replaced by its empirical counterpart $\hat{\rho}_N$ in the Gaussian setting, yielding a computable estimator $\widehat{C}^{\mathsf{A}}_{\hat{\rho}_N}$. This is possible since the associated Gaussian process agrees with the observed process, i.e. $u = v$ (see also Remark 5.4.6 for a technical discussion of this point). In the sub-Gaussian setting, Theorem 5.2.3 shows that an adaptive threshold estimator with an appropriate choice of thresholding radius ρ_N achieves the same estimation error as in the Gaussian case. In practice, we advocate choosing the thresholding radius ρ_N by cross-validation in non-Gaussian settings.*

Our theory for Gaussian data holds for both the sample-based and the Wick's-based estimators of the variance component. The Wick's-based estimator might be preferred in practice as it only requires estimating the second moment of the process, whereas the sample-based estimator requires estimating both the second and fourth moments, which is more computationally intensive. From a theoretical perspective, the Wick's-based estimator is also easier to analyze using results for quadratic empirical processes (see e.g. Mendelson [2016, 2010]), whereas the sample estimator relies on bounds for higher order multi-product empirical processes as described in Section 5.4.2. We further remark that in contrast to finite-dimensional results in which the high probability guarantee improves as d_X increases, the probability in Theorem 5.2.3 approaches 1 as the expected supremum of the normalized process grows. It is straightforward but tedious to derive high probability bounds that are more general in that they depend additionally on a confidence parameter $t \geq 1$. We provide such bounds for the prerequisite Lemmas 5.3.2 and 5.3.3. In Section 5.2.3, we study an explicit family of processes for which the expected supremum can be expressed in terms of parameters of the covariance kernel. Finally, we note that the pre-factor c_0 is unspecified. In the existing literature (e.g. Cai and Liu [2011], Bickel and Levina [2008b], Al-Ghattas et al. [2023]) it is common to *choose c_0 manually or in a data-driven way, such as via cross-validation. For our purposes, we fix $c_0 = 5$ in our simulated experiments in Section 5.2.3.*

Remark 5.2.5 (Comparison to the Partially Observed Framework)**.** *In this work, we assume access to the (infinite dimensional) Gaussian random functions $u_1, \ldots, u_N$. While in practice we cannot work with such infinite-dimensional functional data, the theory is nonetheless illuminating for finite dimensional discretizations, as demonstrated by our empirical study in Section 5.2.3. An alternative approach, described in Section 5.1.1, is the partial observations framework. While this approach is often considered more practical as real-world data is always discrete, we argue that the partial observations approach inadvertently masks the underlying structure of the problem, as the bounds in that literature necessarily rely on the smoothness exponents (e.g. the exponent of the Hölder condition when the underlying true functions are assumed to be Hölder smooth), and not on the expected supremum of the process, as in our theory. This dependence is an artifact of the smoothness assumption, as opposed to being a quantity that fundamentally characterizes the behavior of the underlying process. In effect, the discretization step is taken far too early in the partial observations framework to uncover the dependence on the expected supremum. We comment that our approach is more in line with the operator learning literature Kovachki et al. [2024], de Hoop et al. [2023], Mueller and Siltanen [2012] and adheres to the philosophy put forward in Dashti and Stuart [2017b], which states "...it is advantageous to design algorithms which, in principle, make sense in infinite dimensions; it is these methods which will perform well under refinement of finite dimensional approximations."*

5.2.3 Comparison to Other Estimators

In this section, we extend our analysis of the adaptive covariance operator estimator by comparing to other candidate estimators, namely the universal thresholding and sample covariance estimators. In Section 5.2.3, we first demonstrate that universal thresholding is inferior to adaptive thresholding over the class $\mathcal{K}_q^*(R_q)$. Next, in Section 5.2.3 we restrict attention to a class of highly nonstationary processes and show that adaptive thresholding significantly improves over the sample covariance estimator. Finally, in Section 5.2.3 we compare all three estimators on simulated experiments.

Inferiority of Universal Thresholding

In this section, we show rigorously that universal thresholding over the class $\mathcal{K}_q^*$ can perform arbitrarily poorly relative to adaptive thresholding. The result extends [Cai and Liu, 2011, Theorem 4], which demonstrates in the finite-dimensional setting that universal thresholding behaves poorly over the class $\mathcal{U}_q^*$. The result relies on a reduction of the infinite-dimensional problem to a finite-dimensional one, to which the existing aforementioned theory can be applied.

Theorem 5.2.6. *Suppose that $R_q^q \geq 8$ and ρ_N is defined as in Theorem 5.2.3. Then, there exists a covariance operator $C_0 \in \mathcal{K}_q^*(R_q)$ such that, for sufficiently large N,*

$$\inf_{\gamma_N \geq 0} \mathbb{E}_{\{u_n\}_{n=1}^N \overset{i.i.d.}{\sim} GP(0,C_0)} \|\widehat{C}_{\gamma_N}^{\mathsf{U}} - C_0\| \gtrsim (R_q^q)^{2-q} \rho_N^{1-q},$$

where $\widehat{C}_{\gamma_N}^{\mathsf{U}}$ is the universal thresholding estimator with threshold γ_N.

Remark 5.2.7. *Since $q \in (0,1)$, the lower bound for universal thresholding in Theorem 5.2.6 is larger than the upper bound for adaptive thresholding in Theorem 5.2.3, and this discrepancy grows as R_q increases. Therefore, Theorem 5.2.6 implies that over the class $\mathcal{K}_q^*$, universal thresholding estimators can perform significantly worse than their adaptive counterparts, regardless of how the universal threshold parameter γ_N is chosen. This agrees with intuition, since if the scale of the marginal variance of the process varies significantly over the domain D, universal thresholding will intuitively need to scale with the largest of these scales, and as a result the thresholding radius will be too large for all but a small portion of the domain.*

See also Remark 5.2.12.

Nonstationary Weighted Covariance Models

In this section, we further demonstrate the utility of adaptive thresholding by applying Theorem 5.2.3 to an explicit class of highly nonstationary covariance models. To that end, we restrict attention to a subset of $\mathcal{K}_q^*$ of operators with covariance functions of the form

$$k(x,y) = \sigma(x)\sigma(y)\tilde{k}(x,y), \tag{5.5}$$

where $\tilde{k}$ is chosen to be an isotropic *base* covariance function and σ will represent a *marginal variance* function. It follows by standard facts on the construction of covariance functions (see e.g. Genton [2001]) that (5.5) defines a valid covariance function. This class is particularly interesting as it can be thought of as a *weighted* version of many standard isotropic covariance functions used in practice, such as the squared exponential $\tilde{k}^{\mathrm{SE}}$ and Matérn $\tilde{k}^{\mathrm{Ma}}$ classes, defined respectively in (5.8). Further, it permits us to express the theoretical quantities of Theorem 5.2.3 in terms of interpretable parameters of the covariance function, as we now describe rigorously.

Throughout this section, we assume that the data $u_1, \dots, u_N$ are Gaussian and that the base function $\tilde{k}$ satisfies the following:

Assumption 5.2.8. *$\tilde{k}$ in (5.5) is a covariance function satisfying:*

(i) $\tilde{k}$ is isotropic and positive, so that for $r = \|x-y\|$, $\tilde{k}(x,y) = \tilde{k}(r) > 0$. Further, $\tilde{k}(r)$ is differentiable, strictly decreasing on $[0,\infty)$, and satisfies $\lim_{r\to\infty}\tilde{k}(r) = 0$.

(ii) $\tilde{k} = \tilde{k}_\lambda$ depends on a correlation lengthscale parameter $\lambda > 0$ such that $\tilde{k}_\lambda(\varphi r) = \tilde{k}_{\lambda\varphi^{-1}}(r)$ for any $\varphi > 0$, and $\tilde{k}_\lambda(0) = \tilde{k}(0)$ is independent of λ.

The class described in Assumption 5.2.8 contains many popular examples of covariance functions, such as the squared exponential (Gaussian) and the Matérn models Williams and

Rasmussen [2006]. Importantly though, functions of the form (5.5) are significantly more general since they are nonstationary (and therefore non-isotropic). This nonstationarity is introduced by the marginal variance functions σ, which are termed as such since $k(x,x) = \sigma^2(x)$. The theoretical study of the small lengthscale regime was initiated in Al-Ghattas et al. [2023] in which the authors considered kernels satisfying Assumption 5.2.8, or equivalently $\sigma(x) \equiv 1$. Our analysis here extends their results to the challenging non-isotropic and nonstationary setting. For concreteness, we focus on a specific class of marginal variance functions detailed in the following assumption, where we let σ depend on the parameter λ in Assumption 5.2.8.

Assumption 5.2.9. *The marginal variance function in* (5.5) *is taken to be either* $\sigma_\lambda(x;\alpha) \equiv 1$ *or* $\sigma_\lambda(x;\alpha) = \exp\left(\lambda^{-\alpha}\|x\|^2\right)$ *for* $\alpha \in (0, 1/2)$.

Functions of the form (5.5) and which additionally satisfy Assumptions 5.2.8 and 5.2.9 are denoted by $k_\lambda(x,y)$. For small λ, the case $\sigma_\lambda(x;\alpha) = \exp\left(\lambda^{-\alpha}\|x\|^2\right)$ results in a particularly challenging estimation problem due to the extreme nonstationarity induced by the wide range of the exponential function across the domain. The case $\sigma_\lambda(x;\alpha) \equiv 1$ allows us to include unweighted covariance functions, thus strictly generalizing the theory in [Al-Ghattas et al., 2023, Section 2.2]. Our main interest in the exponential marginal variance function is further motivated by the following fact regarding the exponential dot-product covariance function, $k_\lambda(x,y) = \exp(x^\top y/\lambda^2)$, which is interesting as it may be viewed as the simplest example of a nonstationary covariance function. Note that we can write

$$\begin{aligned}
\exp\left(\frac{x^\top y}{\lambda^2}\right) &= \exp\left(\frac{\|x\|^2}{2\lambda^2}\right)\exp\left(\frac{\|y\|^2}{2\lambda^2}\right)\exp\left(-\frac{\|x-y\|^2}{2\lambda^2}\right) \\
&= \exp\left(\frac{\|x\|^2}{2\lambda^2}\right)\exp\left(\frac{\|y\|^2}{2\lambda^2}\right)\tilde{k}_\lambda^{\mathrm{SE}}(x,y).
\end{aligned}$$

Figure 5.1 shows random draws when $\tilde{k}$ is chosen to be the squared exponential (SE) covariance function, with varying choices of λ and α. It is immediately clear that for smaller λ,

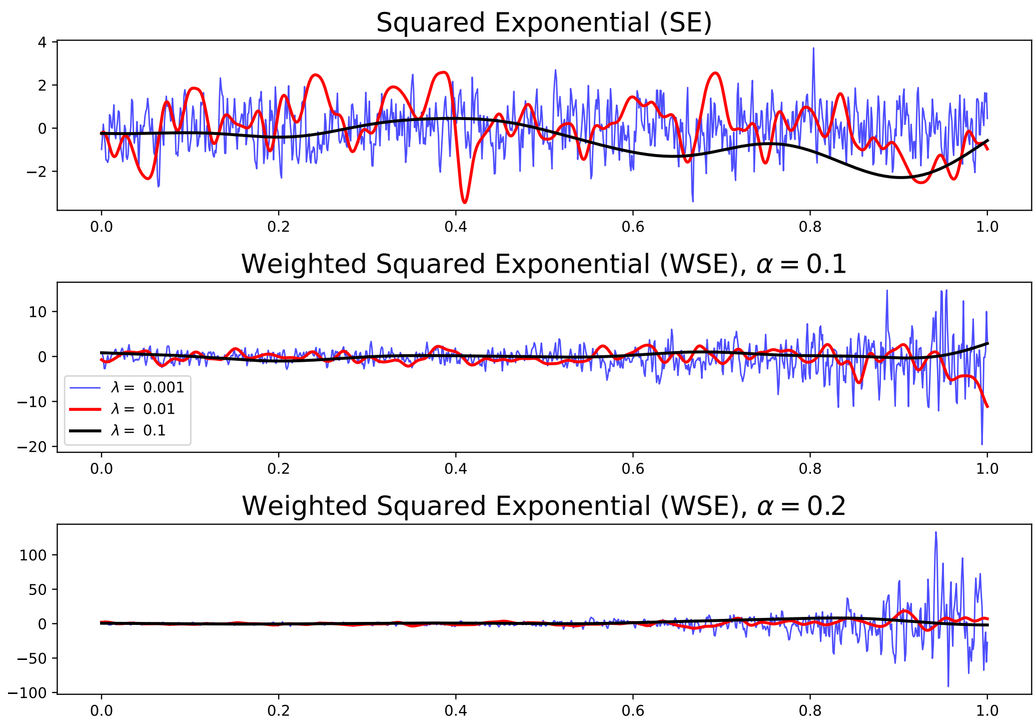

Figure 5.1: Draws from a centered Gaussian process on $D = [0, 1]$ with weighted SE covariance function of the form (5.5) with SE base kernel defined in (5.8). In the first plot, $\sigma = 1$ (unweighted), and in the second and third plots, σ is chosen according to Assumption 5.2.9 and with $\alpha = 0.1, 0.2$ respectively. The scale parameter λ is varied over 0.001 (blue), 0.01 (red) and 0.1 (black)

the processes become more *local*, whereas the role of α is to change the scale of the process across the domain, with this change being more pronounced for larger α. Analogous plots in the $d = 2$ case are presented in Figure 5.6 in the appendix.

Theorem 5.2.10 (Sample Covariance vs. Adaptive Thresholding)**.** *Let* $\widehat{C}$ *and* $\widehat{C}^{\mathsf{A}}_{\hat{\rho}_N}$ *denote the sample covariance and adaptively thresholded estimator respectively. Then, there exists a universal constant* $\lambda_0 > 0$ *such that, for all* $\lambda < \lambda_0$*, it holds with probability at least* $1 - \lambda^d$

that

$$\frac{\|\widehat{C} - C\|}{\|C\|} \asymp \sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N}, \tag{5.6}$$

$$\frac{\|\widehat{C}^{\mathsf{A}}_{\hat{\rho}_N} - C\|}{\|C\|} \lesssim c(q) \left(\frac{\log(\lambda^{-d})}{N}\right)^{(1-q)/2}, \tag{5.7}$$

where $\mathsf{A} \in \{\mathsf{S}, \mathsf{W}\}$ *and* $c(q)$ *is a constant depending only on* q*.*

Remark 5.2.11. *An explicit expression for the constant* $c(q)$ *appearing in Theorem 5.2.10 is provided in the proof of the result. We remark that when* $\tilde{k} := \tilde{k}^{SE}$*, straightforward calculations yield that* $c(q) \asymp q^{-3d/2}$*. We note once more that throughout this work, the dimension* d *of the physical domain* $D = [0, 1]^d$ *is treated as a constant.*

Theorem 5.2.10 — motivated by Al-Ghattas et al. [2023], Koltchinskii and Lounici [2017] — considers the *relative* as opposed to the *absolute* errors commonly used in the sparse estimation literature Bickel and Levina [2008b], Cai and Liu [2011], Fang et al. [2023]. The bound demonstrates that when λ is sufficiently small, the adaptive thresholding estimator exhibits an exponential improvement in sample complexity over the sample covariance estimator. We remark that the bound is identical to [Al-Ghattas et al., 2023, Theorem 2.8] which considers the less general class of unweighted covariance functions, i.e. with $\sigma_\lambda := 1$ in (5.5). The sample covariance bound (5.6) follows by an application of [Koltchinskii and Lounici, 2017, Theorem 9], which shows that with high probability $\|\widehat{C} - C\| \lesssim \|C\|(\sqrt{r(C)/N} \vee r(C)/N)$, where $r(C) = \mathrm{Tr}(C)/\|C\|$ is the effective (intrinsic) dimension of C. To apply this result, it is therefore necessary to derive sharp characterizations for both $\mathrm{Tr}(C)$ and $\|C\|$ in terms of the covariance function parameters α, λ, which we provide in Lemmas 5.6.1 and 5.6.4. The adaptive covariance bound (5.7) follows by an application of our main result, Theorem 5.2.3, and therefore requires a sharp characterization of R_q and ρ_N in terms of the same parameters, provided in Lemma 5.6.3.

Remark 5.2.12 (Universal Thresholding). *Theorem 5.2.10 shows that for sufficiently small λ, adaptive thresholding with an appropriately chosen thresholding parameter will significantly outperform the sample covariance estimator. It is also instructive to consider the universal thresholding estimator in which the same threshold radius ρ_N^{U} is used at all points (x, y) when estimating $k(x, y)$. This estimator was studied in Al-Ghattas et al. [2023] under the assumption that the covariance operator belonged to the class $\mathcal{K}_q$, with $M := \sup_{x\in D} k(x, x) = 1$. Removing the bounded marginal variance assumption, a careful analysis of their theory suggests that the universal threshold radius should be chosen to scale with $\sup_{x\in D} k(x, x)$. This is analogous to the finite d_X-dimensional covariance matrix estimation theory (see e.g. Bickel and Levina [2008b], Cai and Liu [2011]) in which the (universal) thresholding parameter must be chosen to scale with $\max_{i\le d_X} |\Sigma_{ii}|$. For processes with marginal variances that dramatically differ across the domain however, and as noted in Remark 5.2.7, such a scaling causes the estimator to fail as it will necessarily set the estimator to zero for a large proportion of the domain. Specifically in the setting of Assumption 5.2.9, we have that $\sup_{x\in D} k_\lambda(x, x) = e^{2d/\lambda^\alpha}$ and $\inf_{x\in D} k_\lambda(x, x) = 1$. Therefore, as λ decreases, the ratio of largest to smallest marginal variances of the process diverges, and the universal thresholding estimator is zero for larger portions of the domain. This behavior is also borne out in our simulation results (see Figure 5.3) in which a grid of universal threshold parameters is considered and all fail dramatically relative to the adaptive estimator.*

Simulation Results

In this section, we study the behavior of the sample covariance, universal thresholding, and adaptive thresholding estimators. The results provide numerical evidence for our Theorem 5.2.10, and also for the discussion around universal estimators in Remark 5.2.12. Our experiments are carried out in physical dimension $d = 1$ (we also provide results for the case $d = 2$ in 5.9.1). Although our theory works for any base kernel $\tilde{k}$ satisfying Assumption 5.2.8, we focus here on the squared exponential (SE) and Matérn (Ma) classes for simplicity, defined respectively by

$$\begin{aligned}
\tilde{k}_\lambda^{\mathrm{SE}}(x,y) &= \exp\left(-\frac{\|x-y\|^2}{2\lambda^2}\right), \\
\tilde{k}_\lambda^{\mathrm{Ma}}(x,y) &= \frac{2^{1-\nu}}{\Gamma(\nu)}\left(\frac{\sqrt{2\nu}}{\lambda}\|x-y\|\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}}{\lambda}\|x-y\|\right), \qquad \nu > \frac{d-1}{2} \vee \frac{1}{2},
\end{aligned} \tag{5.8}$$

where Γ denotes the Gamma function and K_ν is the modified Bessel function of the second kind. Both covariance functions can be shown to satisfy the assumptions in this work, see [Al-Ghattas et al., 2023, Section 2.2]. Our samples are generated by discretizing the domain $D = [0,1]$ with a uniform mesh of $L = 1000$ points. We consider a total of 30 choices of λ arranged uniformly in log-space and ranging from $10^{-2.5}$ to $10^{-0.1}$. For each λ, with corresponding covariance operator C, the discretized operators are given by the $L \times L$ covariance matrix $C^{ij} = \big(k(x_i, x_j)\big)_{1 \le i,j \le L}$. We sample $N = 5\log(\lambda^{-d})$ realizations of a Gaussian process on the mesh, denoted $u_1, \ldots, u_N \sim N(0, C)$. We then compute the empirical and (adaptively) thresholded sample covariance matrices

$$\widehat{C}^{ij} = \frac{1}{N}\sum_{n=1}^N u_n(x_i)u_n(x_j), \qquad \widehat{C}^{A,ij}_{\hat{\rho}_N} = \widehat{C}^{ij}\mathbf{1}\big\{|\widehat{C}^{ij}| \ge \hat{\rho}_N(\hat{\theta}^{ij})^{1/2}\big\}, \qquad 1 \le i,j \le L,$$

where $\hat{\rho}_N$ is defined as in Theorem 5.2.3, and $\hat{\theta}_{ij}$ are the estimated variance components defined by either

$$(\hat{\theta}_{\mathsf{S}})^{ij} = \frac{1}{N}\sum_{n=1}^N \left(u_n(x_i)u_n(x_j) - \widehat{C}^{ij}\right)^2, \qquad 1 \le i,j \le L,$$

or

$$(\hat{\theta}_{\mathsf{W}})^{ij} = \widehat{C}^{ii}\widehat{C}^{jj} + (\widehat{C}^{ij})^2, \qquad 1 \le i,j \le L.$$

To quantify performance, we consider the relative error of each of the estimators, i.e. $\varepsilon = \|\widehat{C} - C\|/\|C\|$ for the sample covariance, with analogous definitions for the other estimators considered. We repeat the experiment a total of 100 times for each lengthscale, and provide plots of the average relative errors as well as a 95% confidence intervals over the trials. In Figure 5.2, we consider in the first row the (unweighted) squared exponential and Matérn functions (with $\sigma_\lambda := 1$), in the second and third rows we choose σ_λ according to Assumption 5.2.9 with $\alpha = 0.1$ and $\alpha = 0.2$ respectively. In the unweighted case, we also consider the universally thresholded estimator where the threshold is taken to be $\hat{\rho}_N$, which is the correct choice by [Al-Ghattas et al., 2023, Theorem 2.2]. Note that in this case, since $k(x, x) = 1$ for all $x \in D$, all marginal variances are of the same scale, and so the universal and adaptive estimators have the same rate of convergence. While all thresholding estimators exhibit good performance as their relative errors are below 1, it is clear that adaptive thresholding out-performs universal thresholding for the choice of pre-factor 5.

Note that all thresholding estimators significantly outperform the sample covariance estimator for small λ. For the second and third rows, the adaptive estimator continues to significantly outperform the sample covariance and is unaffected by the differences in scale introduced by σ_λ. We observe that the sample-based and Wick's-based adaptive estimators perform similarly, with the Wick's-based estimator exhibiting slightly better performance in all experiments. The results clearly demonstrate our Theorem 5.2.10, since taking only $N = 5\log(\lambda^{-1})$ samples, the relative error of the adaptive estimator remains constant as λ decreases.

Next, in Figure 5.3 we study further the behavior of universal thresholding in the weighted setting with $\alpha = 0.1$. As discussed in Remark 5.2.12, the existing theory suggests to take the threshold radius to scale with $\sup_{x\in D}\sqrt{k_\lambda(x, x)} = e^{d/\lambda^\alpha}$, which becomes extremely large for small λ, and causes the universal threshold estimator to behave effectively like the zero estimator. This choice is reflected by the error $\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$, which has relative error equal to

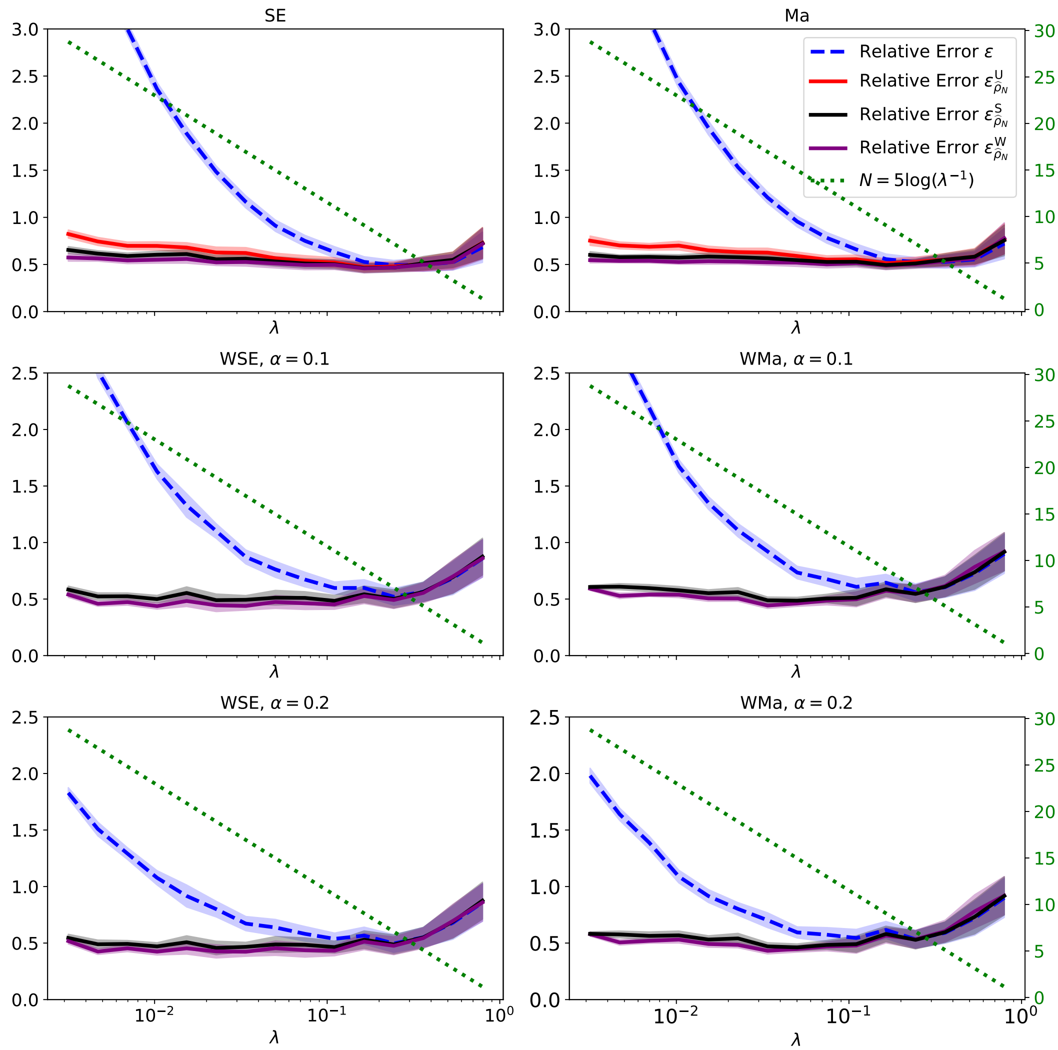

Figure 5.2: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue), universal thresholding ($\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$, red), sample-based adaptive thresholding ($\varepsilon^{\mathsf{S}}_{\hat{\rho}_N}$, black) and Wick's adaptive thresholding ($\varepsilon^{\mathsf{W}}_{\hat{\rho}_N}$, purple) covariance estimators based on a sample size (N, dotted green) for the (weighted) squared exponential (left) and (weighted) Matérn (right) covariance functions in $d = 1$ over 30 Monte-Carlo trials and 30 scale parameters λ ranging from $10^{-2.5}$ to $10^{-0.1}$. The first row corresponds to the unweighted covariance functions and is the only case in which the universal thresholding estimator is considered; the second and third rows correspond to the weighted variants with $\alpha = 0.1, 0.2$ respectively.

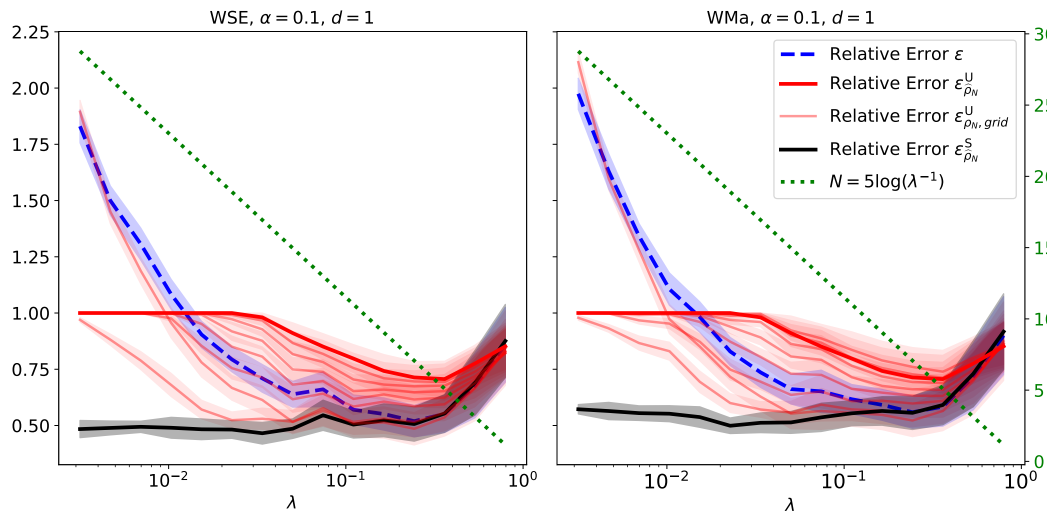

Figure 5.3: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue), universal thresholding ($\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$, red), universal thresholding with data-driven radius ($\varepsilon^{\mathsf{U}}_{\rho_{N,\text{grid}}}$, pink) and sample-based adaptive thresholding ($\varepsilon^{\mathsf{S}}_{\hat{\rho}_N}$, black) covariance estimators based on a sample size (N, dotted green) for the (weighted) squared exponential (left) and (weighted) Matérn (right) covariance functions with $\alpha = 0.1$ in $d = 1$ over 30 Monte-Carlo trials and 30 scale parameters λ ranging from $10^{-2.5}$ to $10^{-0.1}$.

1 for small lengthscales. To further test the universal estimator, for each lengthscale we consider a grid of 10 thresholding radii ranging from 0 (corresponding to just using the sample covariance, with relative error ε) to the one suggested by the theory (corresponding to the theoretically suggested universal estimator, with relative error $\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$). The performance of these 10 estimators is represented in pink. It is clear from these results that regardless of the choice of thresholding radius, the universal estimator performs significantly worse than the adaptive estimator.

The examples considered thus far possess a form of ordered sparsity in that the decay of the covariance function depends monotonically on the physical distance between its two arguments. Although this structure arises in many applications, it is not necessary for the success of thresholding-based estimators. In Figure 5.4, we consider the performance of all estimators when the base kernel exhibits an unordered sparsity pattern. First, we study the

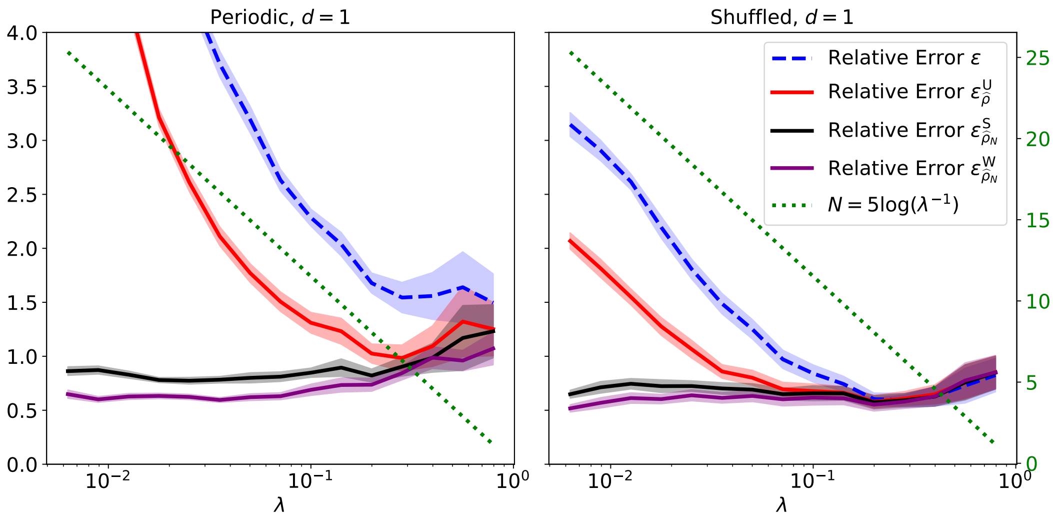

Figure 5.4: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue), universal thresholding ($\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$, red), sample-based adaptive thresholding ($\varepsilon^{\mathsf{S}}_{\hat{\rho}_N}$, black) and Wick's adaptive thresholding ($\varepsilon^{\mathsf{W}}_{\hat{\rho}_N}$, purple) covariance estimators based on a sample size (N, dotted green) for the periodic kernel (left) and shuffled kernel (right) in $d = 1$ over 30 Monte-Carlo trials and 30 scale parameters λ ranging from $10^{-2.2}$ to $10^{-0.1}$.

periodic covariance function $\tilde{k}^{\text{period}}$ given by

$$\tilde{k}^{\text{period}}_{\lambda}(x, y) = \exp\left(-\frac{2\sin^2(\pi\|x-y\|/\eta)}{\lambda^2}\right),$$

where $\eta > 0$ is the periodicity parameter. Intuitively, the periodic covariance function is composed of $\lfloor 1/\eta \rfloor$ bumps spaced uniformly over the domain, each behaving locally like $\tilde{k}^{\text{SE}}_{\lambda}$. Consequently, this kernel is not monotonically decreasing, but it becomes sparser with smaller λ. As another example, we consider the squared-exponential kernel applied to a random permutation of the underlying discretized grid. Shuffling the data breaks the spatial ordering while maintaining the same level of sparsity. For both periodic kernel and shuffled data examples, we choose σ_λ according to Assumption 5.2.9 with $\alpha = 0.1$ and consider 30 scale parameters λ ranging from $10^{-2.2}$ to $10^{-0.1}$. The results demonstrate that adaptive thresholding is superior to both universal thresholding and sample covariance estimators.

Although Theorem 5.2.10 holds for Gaussian data, we investigate numerically here the behavior of all estimators on sub-Gaussian data in the small lengthscale regime. Given two independent centered Gaussian processes $v^{(1)}, v^{(2)}$, with covariance functions $k^1 = k^2$ both satisfying Assumption 5.2.8, we define $u^{(1)} := |v^{(1)}| - \mathbb{E}|v^{(1)}|$ and $u^{(2)} := (|v^{(1)}| - \mathbb{E}|v^{(1)}|)\sin(v^{(2)})$. These transformations ensure the resulting processes are sub-Gaussian (technical details are deferred to 5.9.2). The true covariance matrices are given respectively by

$$C_1^{ij} = \frac{2\sigma_\lambda(x_i)\sigma_\lambda(x_j)}{\pi}\left(\sqrt{1-\tilde{k}_\lambda^2(x_i,x_j)} + \tilde{k}_\lambda(x_i,x_j)\sin^{-1}(\tilde{k}_\lambda(x_i,x_j))\right),$$
$$C_2^{ij} = C_1^{ij} \times e^{-\frac{1}{2}(\sigma_\lambda^2(x_i)+\sigma_\lambda^2(x_j))}\sinh(\sigma_\lambda(x_i)\sigma_\lambda(x_j)\tilde{k}_\lambda(x_i,x_j)), \quad 1 \le i,j \le L.$$

As in Figure 5.3, we consider a grid of 10 threshold radii for each estimator, and for each lengthscale we choose the threshold that gives the smallest average relative error to generate the series in the figure. Throughout we choose σ_λ according to Assumption 5.2.9 with $\alpha = 0.1$, and $\tilde{k} = \tilde{k}^{\mathrm{Ma}}$. Similar results hold in the case $\tilde{k} = \tilde{k}^{\mathrm{SE}}$. The results are presented in Figure 5.5, with both adaptive estimators significantly beating out the sample covariance and universal threshold estimators. The empirical results suggest that the theoretical bound in Theorem 5.2.10 potentially continues to hold beyond the Gaussian setting. We leave a theoretical investigation of this extension to future work.

Remark 5.2.13. *In all of our numerical experiments, the Wick's-based estimator exhibits strong performance at the level of and even superior to that of the sample-based estimator. While our theory suggests that the two estimators have the same rate of convergence, it does not preclude differences owing to the choice of pre-factor c_0 in the choice of sample size. The results therefore indicate that the Wick's-based estimator is more robust to smaller choices of this pre-factor. In the Gaussian setting, this is expected. Recall that Wick's theorem states that for a centered multivariate Gaussian vector $X := (X_1, \dots, X_M)$, then $\mathbb{E}[X_1 X_2 \cdots X_M] = \sum_{\pi\in\Pi_M^2}\prod_{\{i,j\}\in\pi} Cov(X_i, X_j)$, where Π_M^2 is the set of all partitions of*

$\{1, \ldots, M\}$ *of length 2. Therefore, in contrast to the sample-based estimator, the Wick's-based estimator avoids having to compute empirical higher order moments which plausibly leads to a more stable estimator for any given sample size. The situation in the sub-Gaussian setting depicted in Figure 5.5 is somewhat more surprising, given that the Wick's-based estimator is only theoretically justified for Gaussian data. While a rigorous explanation of this phenomena is well beyond the scope of this work, we offer here some intuition as to why the Wick's-based estimator might be competitive even for sub-Gaussian data. A generalization of Wick's theorem to non-Gaussian data given in Leonov and Shiryaev [1959] states that, whenever the joint moment exists,* $\mathbb{E}[X_1 X_2 \cdots X_M] = \sum_{\pi \in \Pi_M} \prod_{a \in \pi} \kappa((X_m)_{m \in a})$, *where* Π_M *is the set of all partitions of* $\{1, \ldots, M\}$, *and* $\kappa((X_m)_{m \in a})$ *is the joint cumulant of the subset* $(X_m)_{m \in a}$. *Therefore, one must estimate all higher-order cumulants as opposed to the Gaussian case in which second-order cumulants suffice. Recall that the cumulants are the coefficients in the Taylor series expansion of the cumulant (log-moment) generating function of* X, $\psi(\gamma) := \log \mathbb{E} e^{\langle \gamma, X \rangle}$, *which for sub-Gaussian* X *is bounded above by* $c\|\gamma\|_2^2$ *for a positive universal constant* c. *In order for this to be true, terms of cubic and higher-order cannot be too large. Consequently, higher-order cumulants (third-order and above) cannot be too large. With this in mind, the Wick-based estimator can be interpreted as a type of penalized estimator that effectively treats these small higher-order cumulants as negligible by approximating them with zero.*

Notice that in Figures 5.2, 5.3, 5.4 and 5.5 thresholding seems to increase the relative error for large λ. We note, however, that our theory holds only in the small λ regime, and consequently the behavior for large λ is not captured. We also note that the increase in error may be solely due to the very small sample size used for large values of λ.

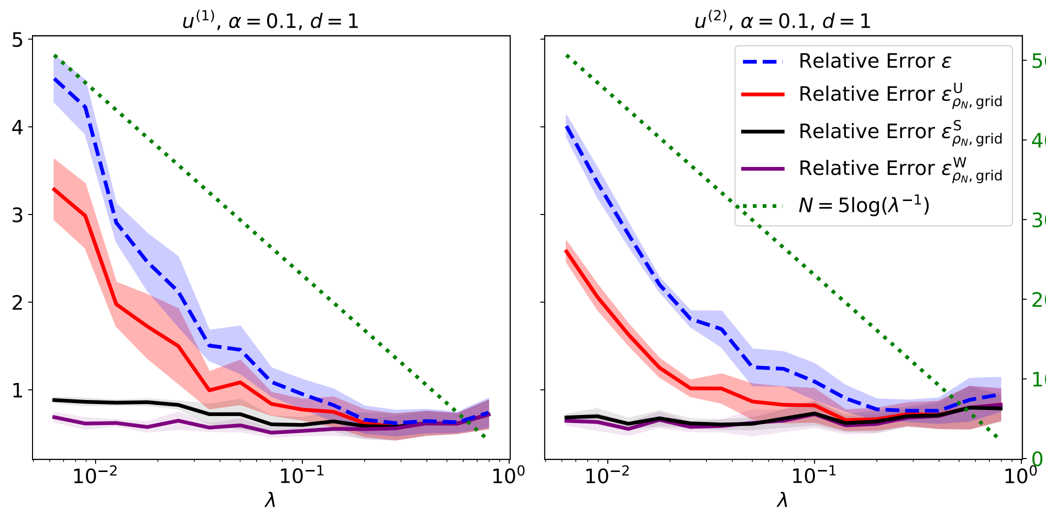

Figure 5.5: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue), universal thresholding with data-driven radius ($\varepsilon^{\mathsf{U}}_{\rho_{N,\text{grid}}}$, pink), sample-based adaptive thresholding with data-driven radius ($\varepsilon^{\mathsf{S}}_{\rho_{N,\text{grid}}}$, black) and Wick's-based adaptive thresholding with data-driven radius ($\varepsilon^{\mathsf{W}}_{\rho_{N,\text{grid}}}$, purple) covariance estimators based on a sample size (N, dotted green) for the sub-Gaussian processes $u^{(1)}$ (left) and $u^{(2)}$ (right). For each data-driven estimator and for each λ, ρ_N is chosen as the error minimizing radius from a set of radii ranging from zero to the choice suggested by the theory in the Gaussian setting. The results are carried out in $d = 1$ over 30 Monte-Carlo trials and 30 scale parameters λ ranging from $10^{-2.2}$ to $10^{-0.1}$.

5.3 Error Analysis for Adaptive-threshold Estimator

In this section, we prove our first main result, Theorem 5.2.3. The proof structure is similar to that for the study of adaptive covariance matrix estimation in Cai and Yuan [2012], but our proof techniques differ in a number of important ways. Chiefly, our results are nonasymptotic and dimension free, owing to our use of recent theory on suprema of product empirical processes put forward in Mendelson [2016] and described in detail in Section 5.4. This new approach allows us to prove Lemmas 5.3.2 and 5.3.3, which provide dimension-free control of the sample covariance and sample variance component. Building on these dimension-free bounds, we show five technical results, Lemmas 5.3.1, 5.3.4, 5.3.5, 5.3.6, and 5.3.7 that are the key building blocks of the proof of the main result. Throughout, we denote the normalized versions of $u, u_1, \ldots, u_N$ by

$$\tilde{u}(\cdot) := \frac{u(\cdot)}{\sqrt{k(\cdot,\cdot)}}, \qquad \tilde{u}_n(\cdot) := \frac{u_n(\cdot)}{\sqrt{k(\cdot,\cdot)}}, \qquad 1 \le n \le N. \tag{5.9}$$

We further denote the Gaussian processes associated to $u, \tilde{u}$ by $v, \tilde{v}$ respectively.

Lemma 5.3.1. *Under Assumption 5.2.1, it holds with probability at least* $1-2e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ *that*

$$\sup_{x,y\in D}\left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\right| \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}},$$

where $\hat{\theta} \in \{\hat{\theta}_{\mathsf{S}}, \hat{\theta}_{\mathsf{W}}\}$ *in the Gaussian setting, and* $\hat{\theta} = \hat{\theta}_{\mathsf{S}}$ *otherwise.*

Proof. We consider first $\hat{\theta}_{\mathsf{S}}$. Assumption 5.2.1 (ii) implies that, for any $x, y \in D$,

$$\begin{aligned}\left|\frac{\hat{\theta}_{\mathsf{S}}(x,y)-\theta(x,y)}{\theta(x,y)}\right| &= \left|\frac{k^2(x,y)-\hat{k}^2(x,y)+\frac{1}{N}\sum_{n=1}^N u_n^2(x)u_n^2(y)-\mathbb{E}[u^2(x)u^2(y)]}{\theta(x,y)}\right| \\ &\le \frac{\left|\hat{k}^2(x,y)-k^2(x,y)\right|}{\nu k(x,x)k(y,y)}+\frac{\left|\frac{1}{N}\sum_{n=1}^N u_n^2(x)u_n^2(y)-\mathbb{E}[u^2(x)u^2(y)]\right|}{\nu k(x,x)k(y,y)} \\ &=: I_1^{\mathsf{S}}+I_2^{\mathsf{S}}.\end{aligned}$$

Controlling I_1^{S}: Note that for constants a, b, we have

$$a^2-b^2=(a-b)(a+b)=(a-b)(a-b+2b)=(a-b)^2+2b(a-b).$$

Therefore,

$$\begin{aligned}I_1^{\mathsf{S}}=\frac{\left|\hat{k}^2(x,y)-k^2(x,y)\right|}{\nu k(x,x)k(y,y)} &\le \frac{1}{\nu}\left|\frac{\hat{k}(x,y)-k(x,y)}{\sqrt{k(x,x)k(y,y)}}\right|^2+2|k(x,y)|\left|\frac{\hat{k}(x,y)-k(x,y)}{\nu k(x,x)k(y,y)}\right| \\ &\le \frac{1}{\nu}\left|\frac{\hat{k}(x,y)-k(x,y)}{\sqrt{k(x,x)k(y,y)}}\right|^2+\frac{2}{\nu}\left|\frac{\hat{k}(x,y)-k(x,y)}{\sqrt{k(x,x)k(y,y)}}\right|,\end{aligned}$$

where we have used that $|k(x,y)| \le \sqrt{k(x,x)k(y,y)}$ by Cauchy-Schwarz. On the event $\Omega_t^{(1)}$ defined in Lemma 5.3.2 it holds that, for all $x, y \in D$,

$$I_1^{\mathsf{S}} \lesssim \frac{1}{\nu}\left(\sqrt{\frac{t}{N}} \vee \frac{t^2}{N} \vee \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^4}{N}\right).$$

Controlling I_2^{S}: On the event $\Omega_t^{(2)}$ defined in Lemma 5.3.3 it holds that, for all $x, y \in D$,

$$I_2^{\mathsf{S}} \lesssim \frac{1}{\nu}\left(\sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}{N}\right).$$

Therefore, by Assumption 5.2.1 (iii) and choosing $t = (\mathbb{E}[\sup_{x\in D} \tilde{v}(x)])^2$, we have

$$I_1^{\mathsf{S}} + I_2^{\mathsf{S}} \lesssim \frac{1}{\nu}\left(\sqrt{\frac{t}{N}} \vee \frac{t^2}{N} \vee \frac{\mathbb{E}[\sup_{x\in D} \tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D} \tilde{v}(x)])^4}{N}\right) \lesssim \frac{1}{\nu}\frac{\mathbb{E}[\sup_{x\in D} \tilde{v}(x)]}{\sqrt{N}}.$$

The proof is completed by noting that the event $A_t := \Omega_t^{(1)} \cap \Omega_t^{(2)}$ has probability at least $1 - 2e^{-t}$ by Lemmas 5.3.2 and 5.3.3.

Next, for $\hat{\theta}_{\mathsf{W}}$, Assumption 5.2.1 (ii) implies that, for any $x, y \in D$,

$$\left|\frac{\hat{\theta}_{\mathsf{W}}(x,y) - \theta(x,y)}{\theta(x,y)}\right| \le \frac{|\hat{k}^2(x,y) - k^2(x,y)|}{\nu k(x,x)k(y,y)} + \left|\frac{\hat{k}(x,x)\hat{k}(y,y) - k(x,x)k(y,y)}{\theta(x,y)}\right| =: I_1^{\mathsf{W}} + I_2^{\mathsf{W}}.$$

Controlling I_1^{W}*:* Since $I_1^{\mathsf{W}} = I_1^{\mathsf{S}}$, it follows that $I_1^{\mathsf{W}} \lesssim \frac{\mathbb{E}[\sup_{x\in D} \tilde{v}(x)]}{\nu\sqrt{N}}$ with probability at least $1 - e^{-(\mathbb{E}[\sup_{x\in D} \tilde{v}(x)])^2}$.

Controlling I_2^{W}*:* Writing

$$\begin{aligned}\hat{k}(x,x)\hat{k}(y,y) - k(x,x)k(y,y) &= (\hat{k}(x,x) - k(x,x))(\hat{k}(y,y) - k(y,y)) \\ &\quad + (\hat{k}(x,x) - k(x,x))k(y,y) + (\hat{k}(y,y) - k(y,y))k(x,x),\end{aligned}$$

and by Assumption 5.2.1, we have that, for any $x, y \in D$,

$$\begin{aligned}I_2^{\mathsf{W}} &\le \left|\frac{(\hat{k}(x,x) - k(x,x))(\hat{k}(y,y) - k(y,y))}{\nu k(x,x)k(y,y)}\right| \\ &\quad + \left|\frac{k(y,y)(\hat{k}(x,x) - k(x,x))}{\nu k(x,x)k(y,y)}\right| + \left|\frac{k(x,x)(\hat{k}(y,y) - k(y,y))}{\nu k(x,x)k(y,y)}\right| \\ &\le \frac{1}{\nu}\sup_{x\in D}\left|\frac{\hat{k}(x,x) - k(x,x)}{k(x,x)}\right|^2 + \frac{2}{\nu}\sup_{x\in D}\left|\frac{\hat{k}(x,x) - k(x,x)}{k(x,x)}\right| =: \frac{1}{\nu}I_{21}^2 + \frac{2}{\nu}I_{21}^{\mathsf{W}}.\end{aligned}$$

Define the event $A := \left\{ I_{21}^{\mathsf{W}} \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} \right\}$, and note that on A,

$$I_2^{\mathsf{W}} \lesssim \frac{1}{\nu}\left(\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}}\right)^2 + \frac{2}{\nu}\left(\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}}\right) \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}}.$$

By Lemma 5.3.2 and Assumption 5.2.1 (iii), we have that $\mathbb{P}(A) \geq 1 - e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$. □

Lemma 5.3.2. *For any $t \geq 1$, define $\Omega_t^{(1)}$ to be the event on which*

$$\sup_{x,y\in D}\left|\frac{\hat{k}(x,y) - k(x,y)}{\sqrt{k(x,x)k(y,y)}}\right| \lesssim \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}{N}.$$

Then, it holds that $\mathbb{P}(\Omega_t^{(1)}) \geq 1 - e^{-t}$.

Proof. The result follows by invoking Lemma 5.4.4 after noting that, for any $x, y \in D$,

$$\begin{aligned}\left|\frac{\hat{k}(x,y) - k(x,y)}{\sqrt{k(x,x)k(y,y)}}\right| &= \left|\frac{1}{N}\sum_{n=1}^{n}\frac{u_n(x)}{\sqrt{\mathbb{E}[u_n^2(x)]}}\frac{u_n(y)}{\sqrt{\mathbb{E}[u_n^2(y)]}} - \mathbb{E}\left[\frac{u(x)}{\sqrt{\mathbb{E}[u^2(x)]}}\frac{u(y)}{\sqrt{\mathbb{E}[u^2(y)]}}\right]\right| \\ &= \left|\frac{1}{N}\sum_{n=1}^{N}\tilde{u}_n(x)\tilde{u}_n(y) - \mathbb{E}[\tilde{u}(x)\tilde{u}(y)]\right|.\end{aligned}$$

□

Lemma 5.3.3. *For any $t \geq 1$, define $\Omega_t^{(2)}$ to be the event on which*

$$\begin{aligned}&\sup_{x,y\in D}\left|\frac{\frac{1}{N}\sum_{n=1}^{N}u_n^2(x)u_n^2(y) - \mathbb{E}[u^2(x)u^2(y)]}{k(x,x)k(y,y)}\right| \\ &\qquad\qquad \lesssim \sqrt{\frac{t}{N}} \vee \frac{t^2}{N} \vee \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^4}{N}.\end{aligned}$$

Then, it holds that $\mathbb{P}(\Omega_t^{(2)}) \geq 1 - e^{-t}$.

Proof. The result follows by invoking Lemma 5.4.5 after noting that, for any $x, y \in D$,

$$\left| \frac{\frac{1}{N}\sum_{n=1}^N u_n^2(x)u_n^2(y) - \mathbb{E}[u^2(x)u^2(y)]}{k(x,x)k(y,y)} \right| = \left| \frac{1}{N}\sum_{n=1}^N \tilde{u}_n^2(x)\tilde{u}_n^2(y) - \mathbb{E}[\tilde{u}^2(x)\tilde{u}^2(y)] \right|.$$

□

Lemma 5.3.4. *Under Assumption 5.2.1, it holds with probability at least* $1-2e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ *that*

$$\sup_{x,y\in D} \left| \frac{\hat{\theta}^{1/2}(x,y) - \theta^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)} \right| \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}},$$

where $\hat{\theta} \in \{\hat{\theta}_{\mathsf{S}}, \hat{\theta}_{\mathsf{W}}\}$ *in the Gaussian setting, and* $\hat{\theta} = \hat{\theta}_{\mathsf{S}}$ *otherwise.*

Proof. Define the event

$$A := \left\{ \sup_{x,y\in D} \left| \frac{\hat{\theta}(x,y) - \theta(x,y)}{\theta(x,y)} \right| \le \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}} \right\}.$$

It holds that $\mathbb{P}(A) \geq 1 - 2e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ by Lemma 5.3.1. Further note that the universal constant in Assumption 5.2.1 (iii) can be taken sufficiently small to ensure that

$\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}} \le \frac{1}{2}$. Then on A, for any $x, y \in D$,

$$
\begin{aligned}
\left|\frac{\hat{\theta}^{1/2}(x,y)-\theta^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| &= \left|\frac{\hat{\theta}^{1/2}(x,y)-\theta^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)}\frac{\theta^{1/2}(x,y)+\hat{\theta}^{1/2}(x,y)}{\theta^{1/2}(x,y)+\hat{\theta}^{1/2}(x,y)}\right| \\
&= \left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\hat{\theta}(x,y)+\theta^{1/2}(x,y)\hat{\theta}^{1/2}(x,y)}\right| \\
&= \left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\frac{\theta(x,y)}{\hat{\theta}(x,y)+\theta^{1/2}(x,y)\hat{\theta}^{1/2}(x,y)}\right| \\
&\le \left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\frac{\theta(x,y)}{\hat{\theta}(x,y)}\right| \le \left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\right|\left|\frac{\theta(x,y)}{\hat{\theta}(x,y)}\right| \\
&\le 2\left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\right| \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}},
\end{aligned}
$$

where the second to last inequality follows since on A we have

$$
|\theta(x,y)| \le |\theta(x,y)-\hat{\theta}(x,y)| + |\hat{\theta}(x,y)| \le \frac{1}{2}|\theta(x,y)| + |\hat{\theta}(x,y)| \implies |\theta(x,y)| \le 2|\hat{\theta}(x,y)|.
$$

□

Lemma 5.3.5. *Under Assumption 5.2.1, it holds with probability at least* $1-3e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ *that*

$$
\sup_{x,y\in D}\left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}},
$$

where $\hat{\theta} \in \{\hat{\theta}_{\mathsf{S}}, \hat{\theta}_{\mathsf{W}}\}$ *in the Gaussian setting, and* $\hat{\theta} = \hat{\theta}_{\mathsf{S}}$ *otherwise.*

Proof. Note that

$$\begin{aligned}
\left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| &\le \left|\frac{\hat{k}(x,y)-k(x,y)}{\theta^{1/2}(x,y)}\right|\left|\frac{\theta^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \\
&\le \left|\frac{\hat{k}(x,y)-k(x,y)}{\theta^{1/2}(x,y)}\right|\left(\left|\frac{\theta^{1/2}(x,y)-\hat{\theta}^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right|+1\right) \\
&\le \left|\frac{\hat{k}(x,y)-k(x,y)}{\sqrt{\nu k(x,x)k(y,y)}}\right|\left(\left|\frac{\theta^{1/2}(x,y)-\hat{\theta}^{1/2}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right|+1\right) \\
&= I_1 \times I_2.
\end{aligned}$$

Controlling I_1 : It holds on the event $\Omega^{(1)}_{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ defined in Lemma 5.3.2 that, for all $x, y \in D$,

$$I_1 \lesssim \frac{1}{\sqrt{\nu}}\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}},$$

and $\mathbb{P}(\Omega^{(1)}_{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}) \ge 1 - e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$.

Controlling I_2*:* Let B be the event on which the bound in Lemma 5.3.4 holds. Then, $\mathbb{P}(B) \ge 1 - 2e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$, and on B

$$I_2 \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}} + 1.$$

Then, on the event $E = \Omega^{(1)}_{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2} \cap B$, we have

$$I_1 \times I_2 \lesssim \frac{1}{\nu^{3/2}}\frac{(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}{N} \vee \frac{1}{\sqrt{\nu}}\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} = \frac{1}{\sqrt{\nu}}\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}}.$$

□

Lemma 5.3.6. *Let* $v, v_1, \dots, v_N$ *denote the Gaussian processes associated to* $u, u_1, \dots, u_N$*,* *which satisfy Assumption 5.2.1. Define*

$$\rho_N = \frac{1}{\nu\sqrt{N}}\mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right], \qquad \hat{\rho}_N = \frac{1}{\nu\sqrt{N}}\left(\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{v_n(x)}{\hat{k}^{1/2}(x,x)}\right).$$

Then, it holds with probability at least $1-4e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ *that* $|\hat{\rho}_N-\rho_N|\lesssim \rho_N$.

Proof.

$$\begin{aligned}
&\nu\sqrt{N}\,|\hat{\rho}_N-\rho_N| = \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{v_n(x)}{\hat{k}^{1/2}(x,x)} - \mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right]\right| \\
&= \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\left(\frac{v_n(x)}{\hat{k}^{1/2}(x,x)} - \frac{v_n(x)}{k^{1/2}(x,x)} + \frac{v_n(x)}{k^{1/2}(x,x)}\right) - \mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right]\right| \\
&\le \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\left(\frac{v_n(x)}{\hat{k}^{1/2}(x,x)} - \frac{v_n(x)}{k^{1/2}(x,x)}\right)\right| + \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{v_n(x)}{k^{1/2}(x,x)} - \mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right]\right| \\
&= I_1 + I_2.
\end{aligned}$$

Controlling I_1*:* We write

$$\frac{v_n(x)}{\hat{k}^{1/2}(x,x)} - \frac{v_n(x)}{k^{1/2}(x,x)} = \frac{v_n(x)}{k^{1/2}(x,x)}\frac{\hat{k}^{1/2}(x,x)-k^{1/2}(x,x)}{\hat{k}^{1/2}(x,x)}.$$

Define the event

$$A := \left\{\sup_{x\in D}\left|\frac{\hat{k}(x,x)-k(x,x)}{k(x,x)}\right| \le \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}}\right\}.$$

By Lemma 5.3.2 and Assumption 5.2.1 (iii), we have that $\mathbb{P}(A)\ge 1-2e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$. Further note that the universal constant in Assumption 5.2.1 (iii) can be taken sufficiently small to ensure that $\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}}\le\frac{1}{2}$. By a similar argument to the one used in Lemma 5.3.4, conditional on A and for any $x \in D$,

$$\left|\frac{\hat{k}^{1/2}(x,x) - k^{1/2}(x,x)}{\hat{k}^{1/2}(x,x)}\right| \lesssim \frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\nu\sqrt{N}} \leq \frac{1}{2}.$$

Therefore, on A it holds that

$$\begin{aligned}
I_1 &= \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\left(\frac{v_n(x)}{\hat{k}^{1/2}(x,x)} - \frac{v_n(x)}{k^{1/2}(x,x)}\right)\right| \\
&\leq \frac{1}{2}\left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{v_n(x)}{k^{1/2}(x,x)}\right| \\
&\lesssim \left|\frac{1}{N}\sum_{n=1}^{N}\sup_{x\in D}\frac{v_n(x)}{k^{1/2}(x,x)} - \mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right]\right| + \mathbb{E}\left[\sup_{x\in D}\frac{v(x)}{k^{1/2}(x,x)}\right] \\
&= I_2 + \nu\sqrt{N}\rho_N.
\end{aligned}$$

Controlling I_2: By [Talagrand, 2022, Lemma 2.4.7], $\sup_{x\in D}\tilde{v}_n(x)$ is $\sup_{x\in D}\mathrm{var}\big(\tilde{v}_n(x)\big)$-sub-Gaussian. Since $\mathrm{var}\big(\tilde{v}_n(x)\big) = 1$, it follows by sub-Gaussian concentration that with probability at least $1 - 2e^{-t}$, $I_2 \leq \sqrt{2t/N}$. Choosing $t = (\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2/2$, we have $I_2 \leq \mathbb{E}[\sup_{x\in D}\tilde{v}(x)]/\sqrt{N}$. Putting the bounds together, we have shown that

$$|\hat{\rho}_N - \rho_N| \leq \frac{1}{\nu\sqrt{N}}(I_1 + I_2) \lesssim \frac{1}{\nu\sqrt{N}}\frac{\mathbb{E}[\sup_{x\in D}\tilde{v}(x)]}{\sqrt{N}} + \rho_N \lesssim \rho_N.$$

□

Lemma 5.3.7. *Under the setting of Lemma 5.3.6, it holds with probability at least* $1 - 7e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2}$ *that*

$$\sup_{x,y\in D}\left|\frac{\hat{k}(x,y) - k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \leq \frac{\hat{\rho}_N}{2},$$

where $\hat{\theta} \in \{\hat{\theta}_{\mathsf{S}}, \hat{\theta}_{\mathsf{W}}\}$ *in the Gaussian setting, and* $\hat{\theta} = \hat{\theta}_{\mathsf{S}}$ *otherwise.*

Proof. Let E be the event on which the bound in the statement of the theorem holds. Then

$$\begin{aligned}\mathbb{P}(E^c) &= \mathbb{P}\left(\sup_{x,y\in D}\left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| + \frac{1}{2}(\rho_N - \hat{\rho}_n) \geq \frac{1}{2}\rho_N\right) \\ &\leq \mathbb{P}\left(\sup_{x,y\in D}\left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \geq \frac{1}{4}\rho_N\right) + \mathbb{P}\left(\rho_N - \hat{\rho}_n \geq \frac{1}{4}\rho_N\right) \\ &\leq 7e^{-(\mathbb{E}[\sup_{x\in D}\tilde{v}(x)])^2},\end{aligned}$$

where the last line follows by Lemmas 5.3.5 and 5.3.6. □

Proof of Theorem 5.2.3. We consider first the Gaussian case. Let $\hat{\theta} \in \{\hat{\theta}_{\mathsf{S}}, \hat{\theta}_{\mathsf{W}}\}$. Define the three events:

$$\begin{aligned}E_1 &:= \left\{\sup_{x,y\in D}\left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \lesssim \frac{\hat{\rho}_N}{2}\right\}, \\ E_2 &:= \left\{\sup_{x,y\in D}\left|\frac{\hat{\theta}(x,y)-\theta(x,y)}{\theta(x,y)}\right| \lesssim \frac{1}{2}\right\}, \\ E_3 &:= \left\{|\hat{\rho}_N - \rho_N| \lesssim \rho_N\right\},\end{aligned}$$

and $E = E_1 \cap E_2 \cap E_3$. The final result holds on E as will be shown below, and so the proof is completed by noting that from Lemmas 5.3.5, 5.3.6 and 5.3.7, $\mathbb{P}(E) \geq 1 - c_1 e^{-(\mathbb{E}[\sup_{x\in D}\tilde{u}(x)])^2}$. Further note that on the event E_2, for any x, y we have the following relation:

$$\frac{1}{2}|\theta(x,y)| \leq |\hat{\theta}(x,y)| \leq 2|\theta(x,y)|. \tag{5.10}$$

Now, defining the set

$$\Omega_x := \left\{y \in D : \left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \geq \frac{\hat{\rho}_N}{2}\right\},$$

we have

$$
\begin{aligned}
&\|\widehat{C}_{\hat\rho_N} - C\| \le \sup_{x\in D}\int_D |\hat k_{\hat\rho_N}(x,y) - k(x,y)|dy \\
&= \sup_{x\in D}\int_{\Omega_x}\left|\frac{\hat k_{\hat\rho_N}(x,y) - k(x,y)}{\hat\theta^{1/2}(x,y)}\right||\hat\theta^{1/2}(x,y)|dy + \sup_{x\in D}\int_{\Omega_x^c}\left|\frac{\hat k_{\hat\rho_N}(x,y) - k(x,y)}{\hat\theta^{1/2}(x,y)}\right||\hat\theta^{1/2}(x,y)|dy \\
&= \sup_{x\in D}\int_{\Omega_x}\left|\frac{\hat k_{\hat\rho_N}(x,y) - \hat k(x,y)}{\hat\theta^{1/2}(x,y)}\right||\hat\theta^{1/2}(x,y)|dy + \sup_{x\in D}\int_{\Omega_x}\left|\frac{\hat k(x,y) - k(x,y)}{\hat\theta^{1/2}(x,y)}\right||\hat\theta^{1/2}(x,y)|dy \\
&+ \sup_{x\in D}\int_{\Omega_x^c}\left|\frac{\hat k_{\hat\rho_N}(x,y) - k(x,y)}{\hat\theta^{1/2}(x,y)}\right||\hat\theta^{1/2}(x,y)|dy =: I_1 + I_2 + I_3.
\end{aligned}
$$

Controlling I_1: For any $x, y \in D$,

$$
\begin{aligned}
\left|\frac{\hat k_{\hat\rho_N}(x,y) - \hat k(x,y)}{\hat\theta^{1/2}(x,y)}\right| &= 0 \times \mathbf{1}\left\{\left|\frac{\hat k(x,y)}{\hat\theta^{1/2}(x,y)}\right| \ge \hat\rho_N\right\} + \left|\frac{\hat k(x,y)}{\hat\theta^{1/2}(x,y)}\right| \times \mathbf{1}\left\{\left|\frac{\hat k(x,y)}{\hat\theta^{1/2}(x,y)}\right| < \hat\rho_N\right\} \\
&\le \hat\rho_N.
\end{aligned}
$$

Therefore,

$$
I_1 \le \hat\rho_N \sup_{x\in D}\int_{\Omega_x}|\hat\theta^{1/2}(x,y)|dy.
$$

By Assumption 5.2.1, we have that

$$
\begin{aligned}
R_q^q &\ge \sup_{x\in D}\int_D (k(x,x)k(y,y))^{(1-q)/2}\,|k(x,y)|^q dy \\
&\ge \sup_{x\in D}\int_{\Omega_x} (k(x,x)k(y,y))^{(1-q)/2}\,|k(x,y)|^q dy \\
&\gtrsim \sup_{x\in D}\int_{\Omega_x} (k(x,x)k(y,y))^{(1-q)/2}\,\hat\rho_N^q|\hat\theta^{q/2}(x,y)|dy \\
&\gtrsim \sup_{x\in D}\int_{\Omega_x} (k(x,x)k(y,y))^{(1-q)/2}\,\hat\rho_N^q|\theta^{q/2}(x,y)|dy,
\end{aligned}
$$

where the third inequality follows by definition of Ω_x, and the final inequality holds by (5.10).

Further, we have

$$\theta(x,y) = \mathrm{var}\big(u(x)u(y)\big) \le \sqrt{\mathbb{E}[u^4(x)]\mathbb{E}[u^4(y)]} \lesssim \mathbb{E}[u^2(x)]\mathbb{E}[u^2(y)] = k(x,x)k(y,y),$$

where the first inequality follows by Cauchy-Schwarz, and the second inequality follows by the L_4-L_2 equivalence property of sub-Gaussian random variables. Therefore, it follows that

$$R_q^q \gtrsim \hat{\rho}_N^q \sup_{x\in D} \int_{\Omega_x} |\theta^{1/2}(x,y)|dy \ge \hat{\rho}_N^q \frac{I_1}{\hat{\rho}_N}.$$

We have therefore shown that $I_1 \lesssim R_q^q \hat{\rho}_N^{1-q}$, and by definition of E_3, it follows immediately that $I_1 \lesssim R_q^q \rho_N^{1-q}$.

Controlling I_2: On E, we have

$$I_2 \lesssim \rho_N \sup_{x\in D} \int_{\Omega_x} |\hat{\theta}^{1/2}(x,y)|dy \lesssim R_q^q \rho_N^{1-q},$$

which can be bounded with an identical argument to the one used to bound I_1.

Controlling I_3: On $E \cap \Omega_x^c$, we have

$$\left|\frac{\hat{k}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \le \left|\frac{\hat{k}(x,y)-k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| + \left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \le \frac{\hat{\rho}_N}{2} + \frac{\hat{\rho}_N}{2} = \hat{\rho}_N.$$

Therefore, $\hat{k}_{\hat{\rho}_N}(x,y) = \hat{k}(x,y)\mathbf{1}\left\{\left|\frac{\hat{k}(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \ge \hat{\rho}_N\right\} = 0$. Now, for any $q \in [0,1)$,

$$
\begin{aligned}
I_3 &\le \sup_{x\in D}\int_D \left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| |\hat{\theta}^{1/2}(x,y)| \mathbf{1}\left\{\left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \le \frac{\hat{\rho}_N}{2}\right\} dy \\
&\le \sup_{x\in D}\int_D \left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| |\hat{\theta}^{1/2}(x,y)| \mathbf{1}\left\{\left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \le \rho_N\right\} dy \\
&\le \rho_N \sup_{x\in D}\int_D \left(\left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| / \rho_N\right)^q |\hat{\theta}^{1/2}(x,y)| \mathbf{1}\left\{\left|\frac{k(x,y)}{\hat{\theta}^{1/2}(x,y)}\right| \le \rho_N\right\} dy \\
&\lesssim \rho_N^{1-q} \sup_{x\in D}\int_D |k(x,y)|^q \, \hat{\theta}(x,y)^{(1-q)/2} dy.
\end{aligned}
$$

The second inequality holds since on E_3, $\hat{\rho}_N \lesssim 2\rho_N$. The third inequality holds since the quantity being taken to the q-th power is smaller than 1 and $q \in [0,1)$. Combining (5.10) with Assumption 5.2.1 (ii) gives that $\hat{\theta}(x,y)^{(1-q)/2} \le \theta(x,y)^{(1-q)/2} \le \big(k(x,x)k(y,y)\big)^{(1-q)/2}$ and so $I_3 \lesssim \rho_N^{1-q} R_q^q$. This completes the proof of the result in the Gaussian case. The proof in the sub-Gaussian setting follows identically except that the events E_1, E_2 are defined with respect to $\hat{\theta}_{\mathsf{S}}$ only, and ρ_N is used in place of $\hat{\rho}_N$. □

5.4 Product Empirical Processes

This section contains the proofs of Lemmas 5.4.4 and 5.4.5, which were used to establish Lemmas 5.3.2 and 5.3.3. The proofs rely on the recent work Al-Ghattas et al. [2025], which provides sharp bounds for suprema of multi-product empirical processes. We begin in Section 5.4.1 by introducing technical definitions as well as the main result regarding multi-product empirical processes from Al-Ghattas et al. [2023]. We then prove in Section 5.4.2 our main results of this section, Lemmas 5.4.4 and 5.4.5. Our proofs have been inspired by the techniques introduced in Koltchinskii and Lounici [2017] as well as Al-Ghattas and Sanz-Alonso [2024c] and Al-Ghattas et al. [2023]. These works deal with product empirical processes in which the product is taken over two sub-Gaussian classes. In contrast, the results here pertain to product empirical processes over a special category of sub-Exponential classes that arise in the nonasymptotic analysis of the variance component $\theta(x, y)$.

5.4.1 Background

Let $X, X_1, \ldots, X_N \overset{\text{i.i.d.}}{\sim} \mathbb{P}$ be a sequence of random variables on a probability space $(\Omega, \mathbb{P})$. The empirical process indexed by a class $\mathcal{F}$ of functions on $(\Omega, \mathbb{P})$ is given by

$$f \mapsto \frac{1}{N}\sum_{n=1}^{N} f(X_n) - \mathbb{E}f(X), \qquad f \in \mathcal{F}.$$

For $s \geq 2$, the order-s multi-product empirical process indexed by $\mathcal{F}$ is given by

$$f \mapsto \frac{1}{N}\sum_{n=1}^{N} f^s(X_n) - \mathbb{E}f^s(X), \qquad f \in \mathcal{F}.$$

For any function f on $(\Omega, \mathbb{P})$ and $\alpha \geq 1$, the Orlicz ψ_α-norm of f is defined as

$$\|f\|_{\psi_\alpha(\mathbb{P})} = \inf\Big\{c > 0 : \mathbb{E}_{X\sim\mathbb{P}}\big[\exp(|f(X)/c|^\alpha)\big] \leq 2\Big\} = \sup_{q\geq 1}\frac{\|f\|_{L_q(\mathbb{P})}}{q^{1/\alpha}}.$$

The base measure will be clear from the context, and so we write $\|f\|_{\psi_\alpha(\mathbb{P})} = \|f\|_{\psi_\alpha}$ and similarly for the L_q-norms. The corresponding Orlicz space L_{ψ_α} contains functions with finite Orlicz ψ_α-norm. A class of functions $\mathcal{G}$ is L-sub-Gaussian if, for every $f, h \in \mathcal{G} \cup \{0\}$,

$$\|f - h\|_{\psi_2} \leq L\|f - h\|_{L_2}.$$

For a sub-Gaussian class $\mathbb{G}$ it holds that, for every $f, h \in \mathcal{G} \cup \{0\}$ and $q \geq 1$,

$$\|f - h\|_{L_q} \leq c\sqrt{q}\|f - h\|_{\psi_2} \leq cL\sqrt{q}\|f - h\|_{L_2}.$$

A class of functions $\mathcal{E}$ is L-sub-Exponential if, for every $f, h \in \mathcal{E} \cup \{0\}$,

$$\|f - h\|_{\psi_1} \leq L\|f - h\|_{L_2}.$$

For a sub-Exponential class $\mathcal{E}$ it holds that, for every $f, h \in \mathcal{E} \cup \{0\}$ and $q \geq 1$,

$$\|f - h\|_{L_q} \leq cq\|f - h\|_{\psi_1} \leq cLq\|f - h\|_{L_2}.$$

Our results depend on Talagrand's γ-functional, whose definition we now recall.

Definition 5.4.1 (Talagrand's γ functional, Talagrand [2022])**.** *Let $(\mathcal{F}, \mathsf{d})$ be a metric space. An admissible sequence of $\mathcal{F}$ is a collection of subsets $\mathcal{F}_s \subset \mathcal{F}$ whose cardinality satisfies $|\mathcal{F}_s| \leq 2^{2^s}$ for $s \geq 1$, and $|\mathcal{F}_0| = 1$. Set*

$$\gamma_2(\mathcal{F}, \mathsf{d}) = \inf \sup_{f \in \mathcal{F}} \sum_{s \geq 0} 2^{s/2} \mathsf{d}(f, \mathcal{F}_s),$$

where the infimum is taken over all admissible sequences, and $\mathsf{d}(f, \mathcal{F}_s) = \inf_{g \in \mathcal{F}_s} \mathsf{d}(f, g)$. We write $\gamma_2(\mathcal{F}, \psi_2)$ when the distance on $\mathcal{F}$ is induced by the ψ_2-norm.

We now introduce a technical result that will be used in the subsequent proofs.

Lemma 5.4.2. *Let $\mathcal{G}, \mathcal{H}$ be arbitrary subsets of a normed space endowed with the norm $\|\cdot\|$. Define $\mathcal{F} = \mathcal{G} + \mathcal{H}$, which inherits this norm. Then*

$$\gamma_2(\mathcal{F}, \mathsf{d}) \leq 2(\sup_{g \in \mathcal{G}} \|g\| + \sup_{h \in \mathcal{H}} \|h\|) + \sqrt{2}(\gamma_2(\mathcal{G}, \mathsf{d}) + \gamma_2(\mathcal{H}, \mathsf{d})),$$

where $\mathsf{d}(a, b) = \|a - b\|$. Moreover, if $\mathcal{G}$ and $\mathcal{H}$ both either contain 0 or are symmetric,

$$\gamma_2(\mathcal{F}, \mathsf{d}) \lesssim \gamma_2(\mathcal{G}, \mathsf{d}) + \gamma_2(\mathcal{H}, \mathsf{d}).$$

Proof. Let $(\mathcal{G}_s)_s, (\mathcal{H}_s)_s$ be admissible sequences for $\mathcal{G}$ and $\mathcal{H}$ respectively. We can construct an admissible sequence for $\mathcal{F}$ as follows. Let $\mathcal{F}_0$ be an arbitrary element of $\mathcal{F}$, and for $s \geq 1$, set $\mathcal{F}_s = \mathcal{G}_{s-1} + \mathcal{H}_{s-1} = \{g + h : g \in \mathcal{G}_{s-1}, h \in \mathcal{H}_{s-1}\}$. This ensures admissibility since $|\mathcal{F}_s| \leq |\mathcal{G}_{s-1}||\mathcal{H}_{s-1}| \leq 2^{2^{s-1}} 2^{2^{s-1}} = 2^{2^s}$. Note then that

$$
\begin{aligned}
\gamma_2(\mathcal{F}, \mathsf{d}) &\leq \sup_{f \in \mathcal{F}} \mathsf{d}(f, \mathcal{F}_0) + \sup_{f \in \mathcal{F}} \sum_{s \geq 1} 2^{s/2} \mathsf{d}(f, \mathcal{F}_s) \\
&= \sup_{f \in \mathcal{F}} \mathsf{d}(f, \mathcal{F}_0) + \sup_{g \in \mathcal{G}, h \in \mathcal{H}} \sum_{s \geq 1} 2^{s/2} \mathsf{d}(g + h, \mathcal{G}_{s-1} + \mathcal{H}_{s-1}) \\
&\leq \sup_{f \in \mathcal{F}} \mathsf{d}(f, \mathcal{F}_0) + \sup_{g \in \mathcal{G}} \sum_{s \geq 1} 2^{s/2} \mathsf{d}(g, \mathcal{G}_{s-1}) + \sup_{h \in \mathcal{H}} \sum_{s \geq 1} 2^{s/2} \mathsf{d}(h, \mathcal{H}_{s-1}) \\
&= \sup_{f \in \mathcal{F}} \mathsf{d}(f, \mathcal{F}_0) + \sqrt{2} \sup_{g \in \mathcal{G}} \sum_{s \geq 0} 2^{s/2} \mathsf{d}(g, \mathcal{G}_s) + \sqrt{2} \sup_{h \in \mathcal{H}} \sum_{s \geq 0} 2^{s/2} \mathsf{d}(h, \mathcal{H}_s).
\end{aligned}
$$

Noting that

$$
\sup_{f \in \mathcal{F}} \mathsf{d}(f, \mathcal{F}_0) \leq \mathrm{diam}(\mathcal{F}) \leq \mathrm{diam}(\mathcal{G}) + \mathrm{diam}(\mathcal{H}) \leq 2(\sup_{g \in \mathcal{G}} \|g\| + \sup_{h \in \mathcal{H}} \|h\|),
$$

and taking the infimum with respect to $(\mathcal{G}_s)_s$ and $(\mathcal{H}_s)_s$ on both sides yields the first result. In the case that $\mathcal{G}$ and $\mathcal{H}$ both either contain 0 or are symmetric, we have that $\sup_{g \in \mathcal{G}} \|g\| \lesssim \gamma_2(\mathcal{G}, \mathsf{d})$ by [Al-Ghattas et al., 2025, Lemma 4.6], and similarly for $\mathcal{H}$. □

The next result provides optimal high probability bounds on order-s multi-product empirical processes.

Theorem 5.4.3 ([Al-Ghattas et al., 2025, Theorem 2.2])**.** *Assume that $0 \in \mathcal{F}$ or that $\mathcal{F}$ is symmetric (i.e., $f \in \mathcal{F} \implies -f \in \mathcal{F}$). For any $s \geq 2$ and $t \geq 1$, it holds with probability at least $1 - e^{-t}$ that, for any $f \in \mathcal{F}$,*

$$
\left| \frac{1}{N} \sum_{n=1}^{n} f^s(X_n) - \mathbb{E} f^s(X) \right| \lesssim_s \frac{\gamma_2(\mathcal{F}, \psi_2) d_{\psi_2}^{s-1}(\mathcal{F})}{\sqrt{N}} \vee \frac{\gamma_2^s(\mathcal{F}, \psi_2)}{N} \vee d_{\psi_2}^s(\mathcal{F}) \left(\sqrt{\frac{t}{N}} \vee \frac{t^{s/2}}{N} \right),
$$

where $\lesssim_s$ indicates that the inequality holds up to a universal positive constant depending only on s, and $d_{\psi_2}(\mathcal{F}) = \sup_{f \in \mathcal{F}} \|f\|_{\psi_2}$.

5.4.2 Product Sub-Gaussian and Sub-Exponential Classes

The goal of this section is to apply Theorem 5.4.3 to the problem of bounding product empirical processes indexed by a function class $\mathcal{F}$, given by

$$f, g \mapsto \frac{1}{N} \sum_{n=1}^{N} f(X_n) g(X_n) - \mathbb{E}[f(X) g(X)], \qquad f, g \in \mathcal{F}.$$

Bounding the suprema of such processes arises in two important ways in this work. First, in establishing uniform bounds on the deviation of the sample covariance function $\hat{k}$ from its expectation, in which case the indexing class $\mathcal{F}$ is sub-Gaussian and we refer to it as a product sub-Gaussian process. Second, in establishing uniform bounds on the deviation of the sample variance component $\hat{\theta}$ from its expectation, in which case $\mathcal{F}$ is sub-Exponential and we refer to it as a sub-Exponential product process.

We now present our two main results of this section. The first bounds the suprema of the product process indexed by two sub-Gaussian classes, and the second bounds the suprema of the product process indexed by two sub-Exponential classes.

We recall here that $u, u_1, \ldots, u_N$ are i.i.d. centered sub-Gaussian and pre-Gaussian random functions on $D = [0,1]^d$ taking values on the real line and with covariance function k. We assume that these functions are Lebesgue almost-everywhere continuous with probability one. Denote by $\tilde{u}, \tilde{u}_1, \ldots, \tilde{u}_N$ their normalized versions as defined in (5.9). Further, recall that a pre-Gaussian process u is one for which there exists a centered Gaussian process, v, that has the same covariance structure as u. Following [Ledoux and Talagrand, 2013, page 261], we refer to v as the Gaussian process *associated to u*.

Lemma 5.4.4. *It holds with probability at least $1 - e^{-t}$ that, for any $x, y \in D$,*

$$\left| \frac{1}{N} \sum_{n=1}^{N} \tilde{u}_n(x)\tilde{u}_n(y) - \mathbb{E}[\tilde{u}(x)\tilde{u}(y)] \right| \lesssim \sqrt{\frac{t}{N}} \vee \frac{t}{N} \vee \frac{\mathbb{E}[\sup_{x\in D} \tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x\in D} \tilde{v}(x)])^2}{N},$$

where $\tilde{v}$ is the Gaussian process associated with $\tilde{u}$.

Proof. For $x \in D$, let $\ell_x : v \mapsto \ell_x(v) = v(x)$ be the evaluation functional at $x \in D$. We then have

$$\begin{aligned}
\sup_{x,y\in D} \left| \frac{1}{N} \sum_{n=1}^{N} \tilde{u}_n(x)\tilde{u}_n(y) - \mathbb{E}[\tilde{u}(x)\tilde{u}(y)] \right| &= \sup_{x,y\in D} \left| \frac{1}{N} \sum_{n=1}^{N} \ell_x(\tilde{u}_n)\ell_y(\tilde{u}_n) - \mathbb{E}[\ell_x(\tilde{u})\ell_y(\tilde{u})] \right| \\
&\le \sup_{x\in D} \left| \frac{1}{N} \sum_{n=1}^{N} \ell_x^2(\tilde{u}_n) - \mathbb{E}[\ell_x^2(\tilde{u})] \right| \\
&\quad + \frac{1}{2} \sup_{x,y\in D} \left| \frac{1}{N} \sum_{n=1}^{N} (\ell_x - \ell_y)^2(\tilde{u}_n) - \mathbb{E}[(\ell_x - \ell_y)^2(\tilde{u})] \right| \\
&\lesssim \sup_{f\in\mathcal{F}} \left| \frac{1}{N} \sum_{n=1}^{N} f^2(\tilde{u}_n) - \mathbb{E} f^2(\tilde{u}) \right|,
\end{aligned}$$

where the first inequality follows by the fact that for two constants a, b, $ab = \frac{1}{2}(a^2 + b^2 - (a-b)^2)$, and the second inequality follows for $\mathcal{F} := \{c_1\ell_x - c_2\ell_y : x, y \in D, c_1, c_2 \in \{0,1\}\}$. Note that $0 \in \mathcal{F}$ since we can take $c_1 = c_2 = 0$. We then have

$$d_{\psi_2}(\mathcal{F}) = \sup_{f\in\mathcal{F}} \|f\|_{\psi_2} \le \sup_{x\in D} \|\ell_x(\tilde{u}_n)\|_{\psi_2} \vee \sup_{x,y\in D} \|(\ell_x - \ell_y)(\tilde{u}_n)\|_{\psi_2} \lesssim 1.$$

Define $\mathcal{G} := \{c\ell_x : x \in D, c \in \{0,1\}\}$, and note that $\mathcal{F} \subset \mathcal{G} - \mathcal{G}$, from which we have

$$\gamma_2(\mathcal{F}, \psi_2) \le \gamma_2(\mathcal{G} - \mathcal{G}, \psi_2) \lesssim \gamma_2(\mathcal{G}, \psi_2) \lesssim \gamma_2(\mathcal{G}, L_2),$$

where the second inequality holds by Lemma 5.4.2 and the third inequality holds by the equivalence of L_2 and ψ_2 norms for linear functionals. Next, let $\tilde{v}$ be the Gaussian process associated to $\tilde{u}$ and define

$$\mathsf{d}_{\tilde{v}}(x,y) := \sqrt{\mathbb{E}\Big[\big(\tilde{v}(x) - \tilde{v}(y)\big)^2\Big]} = \|\ell_x(\cdot) - \ell_y(\cdot)\|_{L_2}, \qquad x, y \in D.$$

Then,

$$\gamma_2(\mathcal{G}, L_2) \lesssim \gamma_2(\{\ell_x : x \in D\}, L_2) = \gamma_2(D, \mathsf{d}_{\tilde{v}}) \asymp \mathbb{E}\left[\sup_{x \in D} \tilde{v}(x)\right],$$

where the first inequality holds by the fact that for any constant $c \in \mathbb{R}$, function class $\mathcal{F}$ and metric d, $\gamma_2(c\mathcal{F}, \mathsf{d}) \leq |c|\gamma_2(\mathcal{F}, \mathsf{d})$, and the second inequality holds by the definition of $\mathsf{d}_{\tilde{v}}$, (see also [Koltchinskii and Lounici, 2017, Theorem 4], [Al-Ghattas et al., 2023, Proposition 3.1]). The final result therefore follows by invoking Theorem 5.4.3 with $s = 2$.

□

Lemma 5.4.5. *It holds with probability at least $1 - e^{-t}$ that, for any $x, y \in D$,*

$$\left| \frac{1}{N} \sum_{n=1}^{N} \tilde{u}_n^2(x) \tilde{u}_n^2(y) - \mathbb{E}[\tilde{u}^2(x)\tilde{u}^2(y)] \right| \lesssim \sqrt{\frac{t}{N}} \vee \frac{t^2}{N} \vee \frac{\mathbb{E}[\sup_{x \in D} \tilde{v}(x)]}{\sqrt{N}} \vee \frac{(\mathbb{E}[\sup_{x \in D} \tilde{v}(x)])^4}{N},$$

where $\tilde{v}$ is the Gaussian process associated with $\tilde{u}$.

Proof. For $x \in D$, let $\ell_x : v \mapsto \ell_x(v) = v(x)$ be the evaluation functional at $x \in D$. We then have

$$\begin{aligned}\sup_{x,y\in D}\left|\frac{1}{N}\sum_{n=1}^{N}\tilde{u}_n^2(x)\tilde{u}_n^2(y)-\mathbb{E}[\tilde{u}^2(x)\tilde{u}^2(y)]\right| &= \sup_{x,y\in D}\left|\frac{1}{N}\sum_{n=1}^{N}\ell_x^2(\tilde{u}_n)\ell_y^2(\tilde{u}_n)-\mathbb{E}[\ell_x^2(\tilde{u})\ell_y^2(\tilde{u})]\right| \\ &\le \frac{1}{3}\sup_{x\in D}\left|\frac{1}{N}\sum_{n=1}^{N}\ell_x^4(\tilde{u}_n)-\mathbb{E}[\ell_x^4(\tilde{u})]\right| \\ &+ \frac{1}{12}\sup_{x,y\in D}\left|\frac{1}{N}\sum_{n=1}^{N}(\ell_x-\ell_y)^4(\tilde{u}_n)-\mathbb{E}[(\ell_x-\ell_y)^4(\tilde{u})]\right| \\ &+ \frac{1}{12}\sup_{x,y\in D}\left|\frac{1}{N}\sum_{n=1}^{N}(\ell_x+\ell_y)^4(\tilde{u}_n)-\mathbb{E}[(\ell_x+\ell_y)^4(\tilde{u})]\right| \\ &\lesssim \sup_{f\in\mathcal{F}}\left|\frac{1}{N}\sum_{n=1}^{N}f^4(\tilde{u}_n)-\mathbb{E}f^4(\tilde{u})\right|,\end{aligned}$$

where the first inequality follows by the fact that for two constants a,b, $a^2b^2=\frac{1}{12}((a+b)^4+(a-b)^4-2a^4-2b^4)$, and the second inequality follows for $\mathcal{F}:=\{c_1\ell_x-c_2\ell_y : x,y\in D, c_1,c_2\in\{-1,0,1\}\}$. Note that $0\in\mathcal{F}$ since we can take $c_1=c_2=0$. We then have

$$d_{\psi_2}(\mathcal{F})=\sup_{f\in\mathcal{F}}\|f\|_{\psi_2}\le\sup_{x\in D}\|\ell_x(\tilde{u}_n)\|_{\psi_2}\vee\sup_{x,y\in D}\|(\ell_x-\ell_y)(\tilde{u}_n)\|_{\psi_2}\lesssim 1.$$

Note that for $\mathcal{G}=\{c\ell_x : x\in D, c\in\{-1,0,1\}\}$, we have $\mathcal{F}\subset\mathcal{G}-\mathcal{G}$. Using once more the fact that for any constant $c\in\mathbb{R}$, function class $\mathcal{F}$ and metric d, $\gamma_2(c\mathcal{F},\mathsf{d})\le|c|\gamma_2(\mathcal{F},\mathsf{d})$, we have that $\gamma_2(\mathcal{G})\lesssim\gamma_2(\{\ell_x : x\in D\})$. By an identical argument to the one used in the proof of Lemma 5.4.4, we have $\gamma_2(\mathcal{F},\psi_2)\lesssim\mathbb{E}[\sup_{x\in D}\tilde{u}(x)]$. The final result therefore follows by invoking Theorem 5.4.3 with $s=4$. □

Remark 5.4.6. *Our results are expressed in terms of the supremum of the Gaussian process v associated with the observed process u rather than directly in terms of the supremum of u. A key step in proving Lemmas 5.4.4 and 5.4.5 is bounding Talagrand's γ-functional, which arises from Theorem 5.4.3. In general, the task of controlling the γ-functional efficiently is* *extremely difficult. One remarkable exception is Talagrand's majorizing measures theorem, which relates $\gamma_2(\mathcal{F}, \mathsf{d})$ to the expected supremum of a Gaussian process. We leverage the pre-Gaussianity of u to establish equivalence between γ_2 functionals defined with respect to $L_2(\mathbb{P})$ (where $\mathbb{P}$ is the law of u) and the natural metric $\mathsf{d}_{\tilde{v}}$ of $\tilde{v}$. Applying Talagrand's theorem, we obtain an upper bound in terms of $\mathbb{E}[\sup_{x\in D} \tilde{v}(x)]$. Extending this approach to instead bound $\mathbb{E}[\sup_{x\in D} \tilde{u}(x)]$ from below in terms of γ_2 or finding an alternative approach that allows a bound in terms of u requires further investigation which we leave to future work.*

5.5 Lower Bound for Universal Thresholding

This section contains the proof of Theorem 5.2.6. The idea is to first reduce the covariance operator estimation problem to a finite-dimensional covariance matrix estimation problem, then apply Theorem 4 in Cai and Liu [2011] which proves a lower bound for universal thresholding in the finite-dimensional covariance matrix estimation problem. The reduction is based on the recent technique developed in [Al-Ghattas et al., 2024b, Proposition 2.6].

Proof of Theorem 5.2.6. For $m \in \mathbb{N}$ to be chosen later, let $\{I_i\}_{i=1}^m$ be a uniform partition of D with $\text{vol}(I_i) = m^{-1}$. For any positive definite matrix $H = (h_{ij}) \in \mathbb{R}^{m\times m}$, define the covariance operator C_H with corresponding covariance function

$$k_H(x, y) = \sum_{i,j=1}^{m} h_{ij}\mathbf{1}_i(x)\mathbf{1}_j(y),$$

where $\mathbf{1}_i(x) := \mathbf{1}\{x \in I_i\}$. Then, C_H is a positive definite operator since, for any $\psi \in L_2(D)$,

$$\int_{D\times D} k_H(x, y)\psi(x)\psi(y)dxdy = \langle H\bar{\psi}, \bar{\psi}\rangle > 0,$$

where $\bar{\psi} = (\bar{\psi}_1, \dots, \bar{\psi}_m)^\top$ and $\bar{\psi}_i = \int_{I_i} \psi(x)dx$. Note further that for $u_n \sim \text{GP}(0, C_H)$, u_n is almost surely a piecewise constant function that can be written as $u_n(x) = \sum_{i=1}^m z_n^{(i)}\mathbf{1}_i(x)$, for $(z_n^{(1)}, \ldots, z_n^{(m)}) \sim N(0, H)$. Consider next the sparse covariance matrix class $\mathcal{U}_q^*(m, R_q)$ studied in Cai and Liu [2011] and defined in (5.2). For any $H \in \mathcal{U}_q^*(m, m^{1/q} R_q)$ it holds that $C_H \in \mathcal{K}_q^*(R_q)$.

$$
\begin{aligned}
&\sup_{x\in D} \int_D (k_H(x,x)k_H(y,y))^{(1-q)/2} |k_H(x,y)|^q dy \\
&= \sup_{x\in D} \int_D \sum_{i,j=1}^m (h_{ii}h_{jj})^{(1-q)/2} |h_{ij}|^q \mathbf{1}_i(x)\mathbf{1}_j(y) dy \\
&= \max_{i\le m} \int_D \sum_{j=1}^m (h_{ii}h_{jj})^{(1-q)/2} |h_{ij}|^q \mathbf{1}_j(y) dy \\
&= \max_{i\le m} \sum_{j=1}^m (h_{ii}h_{jj})^{(1-q)/2} |h_{ij}|^q \int_D \mathbf{1}_j(y) dy \\
&= \frac{1}{m} \max_{i\le m} \sum_{j=1}^m (h_{ii}h_{jj})^{(1-q)/2} |h_{ij}|^q \le R_q^q.
\end{aligned}
$$

Further, for $H \in \mathcal{U}_q^*(m, R_q)$, we have that $mH \in \mathcal{U}_q^*(m, m^{1/q} R_q)$. Now, we will choose $\widetilde{H}_0 = (\widetilde{h}_{0,ij}) \in \mathcal{U}_q^*(m, R_q)$ to be the covariance matrix constructed in the proof of [Cai and Liu, 2011, Theorem 4]. Namely, let $s_1 = \lceil (R_q^q - 1)^{1-q} (\log m/N)^{-q/2} \rceil + 1$, and set

$$
\widetilde{h}_{0,ij} = \begin{cases} 1 & \text{if} \quad 1 \le i = j \le s_1, \\ R_q^q & \text{if} \quad s_1 + 1 \le i = j \le m, \\ 4^{-1} R_q^q \sqrt{\log m/N} & \text{if} \quad 1 \le i \ne j \le s_1, \\ 0 & \text{otherwise.} \end{cases}
$$

Then, $H_0 := m\widetilde{H}_0 \in \mathcal{U}_q^*(m, m^{1/q} R_q)$ and $C_0 := C_{H_0} \in \mathcal{K}_q^*(R_q)$. Next, let $\widehat{C}_{\gamma_N}^{\mathsf{U}}$ have covariance function $\hat{t}_{\gamma_N}(x,y) = \hat{k}(x,y)\mathbf{1}\{|\hat{k}(x,y)| \ge \gamma_N\}$, i.e. $\widehat{C}_{\gamma_N}^{\mathsf{U}}$ is the (universal) thresholding covariance estimator with threshold γ_N. By definition of the operator norm,

$$
\begin{aligned}
\|\widehat{C}^{\mathsf{U}}_{\gamma_N} - C_{H_0}\| &= \sup_{\|f\|_{L_2(D)}=\|g\|_{L_2(D)}=1} \int f(x) \left(\int (\hat{t}_{\gamma_N}(x,y) - k_{H_0}(x,y)) g(y) dy \right) dx \\
&\geq \sup_{a,b\in\mathcal{S}_{m-1}} \int f_a(x) \left(\int (\hat{t}_{\gamma_N}(x,y) - k_{H_0}(x,y)) f_b(y) dy \right) dx \\
&= \sup_{a,b\in\mathcal{S}_{m-1}} m \sum_{i,j=1}^{m} a_i b_j \iint_{D\times D} 1_i(x) 1_j(y) (\hat{t}_{\gamma_N}(x,y) - k_{H_0}(x,y)) dy dx \\
&= \sup_{a,b\in\mathcal{S}_{m-1}} m \sum_{i,j=1}^{m} a_i b_j \left(\iint_{I_i\times I_j} \hat{t}_{\gamma_N}(x,y) dx dy - m^{-2} h_{0,ij} \right) \\
&= m \sup_{a,b\in\mathcal{S}_{m-1}} \left\langle a, \left(\hat{T}_{m,\gamma_N} - m^{-2} H_0 \right) b \right\rangle \\
&= \|m\hat{T}_{m,\gamma_N} - m^{-1} H_0\| = \|m\hat{T}_{m,\gamma_N} - \widetilde{H}_0\|,
\end{aligned}
$$

where $f_a(x) := \sqrt{m}\sum_{i=1}^{m} a_i \mathbf{1}_i(x)$ and the lower bound holds since a is a unit vector and therefore $\|f_a\|_{L_2(D)} = 1$. f_b is defined analogously. Note further that we have defined the $m \times m$ matrix $\hat{T}_{m,\gamma_N}$ with (i,j)-th element

$$
\iint_{I_i\times I_j} \hat{t}_{\gamma_N}(x,y) dx dy = \iint_{I_i\times I_j} \frac{1}{N} \sum_{n=1}^{N} u_n(x) u_n(y) \mathbf{1}\{|\hat{k}(x,y) \geq \gamma_N|\} dx dy,
$$

where $u_1, \ldots, u_n \overset{\text{i.i.d.}}{\sim} \text{GP}(0, C_{H_0})$. Define $v_n = m^{-1/2} u_n$ for $n = 1, \ldots, N$, and so $v_1, \ldots, v_n \overset{\text{i.i.d.}}{\sim} \text{GP}(0, C_{\widetilde{H}_0})$. Then, we have

$$
\begin{aligned}
\inf_{\gamma_N\geq 0} \mathbb{E}_{\{u_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \text{GP}(0,C_{H_0})} \|\widehat{C}^{\mathsf{U}}_{\gamma_N} - C_{H_0}\| &\geq \inf_{\gamma_N\geq 0} \mathbb{E}_{\{u_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \text{GP}(0,C_{H_0})} \|m\hat{T}_{m,\gamma_N} - \widetilde{H}_0\| \\
&= \inf_{\gamma_N\geq 0} \mathbb{E}_{\{v_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \text{GP}(0,C_{\widetilde{H}_0})} \|\widetilde{T}_{m,\gamma_N} - \widetilde{H}_0\|,
\end{aligned}
$$

where $\widetilde{T}_{m,\gamma_N} := m T_{m,\gamma_N}$. Since the samples u_n are piecewise constant functions, we have for any block $I_i \times I_j$ for which the indicator is equal to 1 that

$$\begin{aligned}(\widetilde{T}_{m,\gamma_N})_{ij} &= m^2 \iint_{I_i \times I_j} \frac{1}{N} \sum_{n=1}^{N} v_n(x) v_n(y) \mathbf{1}\{|\hat{k}(x,y)| \geq \gamma_N\} dx dy \\ &= m^2 \iint_{I_i \times I_j} \frac{1}{N} \sum_{n=1}^{N} v_n(x) v_n(y) dx dy \\ &= m^2 \iint_{I_i \times I_j} \frac{1}{N} \sum_{n=1}^{N} w_n^{(i)} w_n^{(j)} dx dy = \frac{1}{N} \sum_{n=1}^{N} w_n^{(i)} w_n^{(j)},\end{aligned}$$

for $w_1, \ldots, w_N \overset{\text{i.i.d.}}{\sim} N(0, \widetilde{H}_0)$ with $w_n = \big(w_n^{(1)}, \ldots, w_n^{(m)}\big)$ for $1 \leq n \leq N$. Therefore, T_{m,γ_N} is a universally thresholded sample covariance matrix estimator. (Note that the scaling inside the indicator is not an issue, as the infimum is over all positive γ_N.) It follows immediately by [Cai and Liu, 2011, Theorem 4] that if m is chosen to satisfy $N^{5q} \leq m \leq e^{o(N^{1/3})}$ and $8 \leq R_q^q \leq \min\Big\{m^{1/4}, 4\sqrt{\frac{N}{\log m}}\Big\}$, then, for sufficiently large N,

$$\inf_{\gamma_N \geq 0} \mathbb{E}_{\{v_n\}_{n=1}^N \overset{\text{i.i.d.}}{\sim} \mathrm{GP}(0, C_{\tilde{H}_0})} \|\widetilde{T}_{m,\gamma_N} - \widetilde{H}_0\| \gtrsim (R_q^q)^{2-q} \left(\frac{\log m}{N}\right)^{(1-q)/2}.$$

Note then that for $u \sim \mathrm{GP}(0, C_{H_0})$, we have that $u(x) = \sum_{i=1}^m z^{(i)} \mathbf{1}_i(x)$, for $z = \big(z^{(1)}, \ldots, z^{(m)}\big) \sim N(0, H_0)$. Further, the normalized process $\tilde{u}(x)$ is of the form $\tilde{u}(x) = u(x)/\sqrt{k_{H_0}(x,x)} = \sum_{i=1}^m g^{(i)} \mathbf{1}_i(x)$, for $g = \big(g^{(1)}, \ldots, g^{(m)}\big) \sim N(0, D_0)$ where D_0 has elements $d_{0,ij} = m\widetilde{h}_{0,ij}/\sqrt{\mathrm{var}(z^{(i)} z^{(j)})}$. By [Van Handel, 2017, Lemma 2.3],

$$\mathbb{E}\left[\sup_{x \in D} \tilde{u}(x)\right] = \mathbb{E}\left[\max_{i \leq m} g^{(i)}\right] \lesssim \max_{i \leq m} \sqrt{d_{0,ii} \log(i+1)}.$$

Moreover, for $1 \leq i \leq m$, using that $\mathbb{E}[(z^{(i)})^\zeta] = (\zeta - 1)!!(\mathrm{var}(z^{(i)}))^{\zeta/2}$ for any non-negative even integer ζ, we have

$$\mathrm{var}\big((z^{(i)})^2\big) = \mathbb{E}\left[(z^{(i)})^4\right] - \left(\mathbb{E}\left[(z^{(i)})^2\right]\right)^2 = 3m^2\widetilde{h}_{0,ii}^2 - m^2\widetilde{h}_{0,ii}^2 = m^2\widetilde{h}_{0,ii}^2.$$

Therefore,

$$d_{0,ii} = \frac{m\widetilde{h}_{0,ii}}{\sqrt{\mathrm{var}((z^{(i)})^2)}} = \frac{1}{\sqrt{2}}.$$

Plugging this in yields $\mathbb{E}[\sup_{x\in D}\tilde{u}(x)] \lesssim \sqrt{\log m}$ and so the bound becomes

$$(R_q^q)^{2-q}\left(\frac{\log m}{N}\right)^{(1-q)/2} \gtrsim (R_q^q)^{2-q}\left(\frac{(\mathbb{E}[\sup_{x\in D}\tilde{u}(x)])^2}{N}\right)^{(1-q)/2},$$

as desired. □

5.6 Error Analysis for Nonstationary Weighted Covariance Models

Throughout this section, u denotes a centered Gaussian process on $D = [0,1]^d$ with covariance function $k_\lambda(x,y)$ satisfying Assumptions 5.2.8 and 5.2.9. We make repeated use of the following easily verifiable facts:

1. $\tilde{k}_\lambda(r) = \tilde{k}_1(\lambda^{-1}r)$.

2. $\sigma_\lambda(\lambda x;\alpha) = \sigma_1(\lambda^{1-\alpha/2}x;\alpha)$.

3. For any $\lambda > 0$, $\alpha \in (0,1/2)$, $1 \le \sigma_\lambda(x;\alpha) \le \exp(d/\lambda^\alpha)$.

In this section, we use the notation "(E), $\lambda \to 0^+$" to mean that there is a universal constant $\lambda_0 > 0$ such that if $\lambda < \lambda_0$, then (E) holds. We interchangeably use the term "*for sufficiently small* λ."

Lemma 5.6.1. *It holds that*

$$Tr(C) = 2^{-d/2}\tilde{k}(0)\lambda^{\alpha d/2}\left(\int_0^{\sqrt{2/\lambda^\alpha}} e^{t^2}\,dt\right)^d.$$

Therefore, for sufficiently small λ, it holds that

$$Tr(C) \asymp \tilde{k}(0)\lambda^{\alpha d} e^{2d/\lambda^{\alpha}}.$$

Proof. By definition,

$$\mathrm{Tr}(C) = \int_D k_\lambda(x,x)dx = \tilde{k}(0)\int_D \sigma_\lambda^2(x)dx = \tilde{k}(0)\|\sigma_\lambda\|_{L_2(D)}^2.$$

Then, note that

$$\begin{aligned}\|\sigma_\lambda\|_{L_2(D)}^2 &= \int_D \exp(2\lambda^{-\alpha}\|x\|^2)dx = \int_D \exp(2\lambda^{-\alpha}\sum_{j=1}^d x_j^2)dx \\ &= \prod_{j=1}^d \int_0^1 \exp(2\lambda^{-\alpha}x_j^2)dx = 2^{-d/2}\lambda^{\alpha d/2}\left(\int_0^{\sqrt{2/\lambda^\alpha}} e^{t^2}dt\right)^d.\end{aligned}$$

In the last line of the above working, we have used the fact that $\int_0^1 e^{cz^2}dz = \frac{1}{\sqrt{c}}\int_0^{\sqrt{c}} e^{t^2}dt$. Then, we have

$$\int_0^{\sqrt{2/\lambda^\alpha}} e^{t^2}dt = e^{2/\lambda^\alpha}\mathcal{D}(\sqrt{2/\lambda^\alpha}),$$

where $\mathcal{D}(\cdot)$ is the Dawson function. By [Abramowitz and Stegun, 1968, Section 7.1], it holds for λ sufficiently small

$$\mathcal{D}(\sqrt{2/\lambda^\alpha}) = \frac{1}{2\sqrt{2}}\lambda^{\alpha/2} + \frac{1}{8\sqrt{2}}\lambda^{3\alpha/2} + \frac{3}{32\sqrt{2}}\lambda^{5\alpha/2} + \cdots \asymp \lambda^{\alpha/2}.$$

Therefore, for λ sufficiently small,

$$\|\sigma_\lambda\|_{L_2(D)}^2 \asymp 2^{-d/2}\lambda^{\alpha d/2}e^{2d/\lambda^\alpha}\lambda^{\alpha d/2} = \lambda^{\alpha d}e^{2d/\lambda^\alpha}.$$

□

Lemma 5.6.2. *It holds that*

$$\sup_{x\in D}\int_D |k_\lambda(x,y)|^q dy \asymp q^{-d}\lambda^{d(\alpha+1)}e^{2dq/\lambda^\alpha}\int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr, \qquad \lambda\to 0^+.$$

Proof. Starting with the upper bound, we have

$$\begin{aligned}\sup_{x\in D}\int_D |k_\lambda(x,y)|^q dy &= \sup_{x\in D}\int_D e^{q\|x\|^2/\lambda^\alpha}e^{q\|y\|^2/\lambda^\alpha}|\tilde{k}_\lambda(\|x-y\|)|^q dy\\ &\le \sup_{x\in D} e^{q\|x\|^2/\lambda^\alpha}\sup_{x\in D}\int_D e^{q\|y\|^2/\lambda^\alpha}|\tilde{k}_\lambda(\|x-y\|)|^q dy\\ &= e^{dq/\lambda^\alpha}\int_D e^{q\|y\|^2/\lambda^\alpha}|\tilde{k}_\lambda(\|y\|)|^q dy\\ &= e^{dq/\lambda^\alpha}\int_D e^{q\|y\|^2/\lambda^\alpha}|\tilde{k}_1(\|\lambda^{-1}y\|)|^q dy\\ &= \lambda^d e^{dq/\lambda^\alpha}\int_{[0,\lambda^{-1}]^d} e^{q\lambda^{2-\alpha}\|y\|^2}|\tilde{k}_1(\|y\|)|^q dy,\end{aligned}$$

where the last line follows by substituting $y\mapsto\lambda^{-1}y$, and noting that the transformation has Jacobian λ^d. Therefore, treating $y=(y_1,\dots,y_d)$ as a random vector with $y_1,\dots,y_d \overset{\text{i.i.d.}}{\sim}$ unif$([0,\lambda^{-1}])$

$$\begin{aligned}\sup_{x\in D}\int_D |k_\lambda(x,y)|^q dy &\le e^{dq/\lambda^\alpha}\left(\lambda^d\int_{[0,\lambda^{-1}]^d} e^{q\lambda^{2-\alpha}\|y\|^2}|\tilde{k}_1(\|y\|)|^q dy\right)\\ &= e^{dq/\lambda^\alpha}\mathbb{E}_y[e^{q\lambda^{2-\alpha}\|y\|^2}|\tilde{k}_1(\|y\|)|^q]\\ &\le e^{dq/\lambda^\alpha}\mathbb{E}_y[e^{q\lambda^{2-\alpha}\|y\|^2}]\mathbb{E}_y[|\tilde{k}_1(\|y\|)|^q]\\ &= e^{dq/\lambda^\alpha}\left(\lambda^d\int_{[0,\lambda^{-1}]^d} e^{q\lambda^{2-\alpha}\|y\|^2}dy\right)\left(\lambda^d\int_{[0,\lambda^{-1}]^d}|\tilde{k}_1(\|y\|)|^q dy\right)\\ &= \lambda^{2d}e^{dq/\lambda^\alpha}\left(\int_{[0,\lambda^{-1}]^d} e^{q\lambda^{2-\alpha}\|y\|^2}dy\right)\left(\int_{[0,\lambda^{-1}]^d}|\tilde{k}_1(\|y\|)|^q dy\right)\\ &=: \lambda^{2d}e^{dq/\lambda^\alpha}I_1\times I_2,\end{aligned}$$

where we have used the covariance inequality [Berger and Casella, 2001, Theorem 4.7.9], which states that for non-decreasing g and non-increasing h, $\mathbb{E}_y[g(y)h(y)] \leq (\mathbb{E}_y[g(y)])(\mathbb{E}_y[h(y)])$. For the first integral, using the substitution $v_j^2 = q\lambda^{2-\alpha}y_j^2$ yields

$$
\begin{aligned}
I_1 &= \prod_{j=1}^{d} \int_0^{\lambda^{-1}} e^{q\lambda^{2-\alpha}y_j^2} dy_j = \prod_{j=1}^{d} \frac{1}{\sqrt{q\lambda^{2-\alpha}}} \int_0^{\lambda^{-1}\sqrt{q\lambda^{2-\alpha}}} e^{v_j^2} dv_j \\
&= \left(\frac{\exp(\lambda^{-2}q\lambda^{2-\alpha})}{\sqrt{q\lambda^{2-\alpha}}} \mathcal{D}(\lambda^{-1}\sqrt{q\lambda^{2-\alpha}}) \right)^d \asymp \left(\frac{\exp(\lambda^{-2}q\lambda^{2-\alpha})}{\sqrt{q\lambda^{2-\alpha}}} \frac{1}{\lambda^{-1}\sqrt{q\lambda^{2-\alpha}}} \right)^d \\
&= \left(\frac{\exp(\lambda^{-2}q\lambda^{2-\alpha})}{\lambda^{-1}q\lambda^{2-\alpha}} \right)^d = \left(\frac{\exp(q\lambda^{-\alpha})}{q\lambda^{1-\alpha}} \right)^d \asymp \frac{e^{dq/\lambda^\alpha}}{q^d\lambda^{d(1-\alpha)}},
\end{aligned}
$$

where $\mathcal{D}(x)$ is the Dawson function, and we have used the fact that $\mathcal{D}(x) \asymp x^{-1}$ for $x \to \infty$ (see the proof of Lemma 5.6.1). For the second integral, by switching to polar coordinates, we have

$$
\begin{aligned}
I_2 &\leq \int_{\mathbb{R}^d} |\tilde{k}_1(\|y\|)|^q dy = \int_0^\infty r^{d-1} \int_{\mathcal{S}_{d-1}} |\tilde{k}_1(r\|u\|)|^q d\mathfrak{s}_{d-1}(u) dr \\
&= \int_0^\infty r^{d-1}\tilde{k}_1(r)^q \int_{\mathcal{S}_{d-1}} d\mathfrak{s}_{d-1}(u) dr = A(d) \int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr,
\end{aligned}
$$

where we have used the fact that $\|u\| = 1$ for any $u \in \mathcal{S}_{d-1}$, and $A(d)$ denotes the surface area of the unit sphere in $\mathbb{R}^d$. Therefore, we have that

$$
\begin{aligned}
\sup_{x\in D} \int_D |k_\lambda(x,y)|^q dy &\leq \lambda^{2d} e^{dq/\lambda^\alpha} I_1 \times I_2 \\
&\leq \lambda^{2d} e^{dq/\lambda^\alpha} \frac{e^{dq/\lambda^\alpha}}{q^d\lambda^{d(1-\alpha)}} A(d) \int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr \\
&= q^{-d}\lambda^{d(1+\alpha)} e^{2dq/\lambda^\alpha} A(d) \int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr.
\end{aligned}
$$

For the lower bound, note that

$$\begin{aligned}
\sup_{x\in D}\int_D |k_\lambda(x,y)|^q dy \geq \int_{D\times D}|k_\lambda(x,y)|^q dxdy &= \int_{[0,1]^d\times[0,1]^d}\sigma_\lambda^q(x)\sigma_\lambda^q(y)|\tilde{k}_\lambda(\|x-y\|)|^q dxdy \\
&= \int_{[0,1]^d\times[0,1]^d}\sigma_\lambda^q(x)\sigma_\lambda^q(y)|\tilde{k}_1(\lambda^{-1}\|x-y\|)|^q dxdy \\
&= \lambda^{2d}\int_{[0,\lambda^{-1}]^d\times[0,\lambda^{-1}]^d}\sigma_\lambda^q(\lambda x')\sigma_\lambda^q(\lambda y')|\tilde{k}_1(\|x'-y'\|)|^q dx'dy',
\end{aligned}$$

where the last line is due to the substitution $x' = \lambda^{-1}x$ and $y' = \lambda^{-1}y$. Now, let $w = x' - y'$ and $z = x' + y'$ and note that the Jacobian of this transformation is 2^{-d}, and also that $x' = (z+w)/2$ and $y' = (z-w)/2$. Continuing from the last line of the above display, the substitution and the fact that $\sigma_\lambda(\lambda x') = \sigma_1(\lambda^{1-\alpha/2}x') = \exp(\lambda^{2-\alpha}\|x'\|^2)$ give that

$$\begin{aligned}
&= \frac{\lambda^{2d}}{2^d}\int_{[-\lambda^{-1},0]^d}\left(\int_{R_1}\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z+w\|^2\right)\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z-w\|^2\right)dz\right)|\tilde{k}_1(\|w\|)|^q dw \\
&+ \frac{\lambda^{2d}}{2^d}\int_{[0,\lambda^{-1}]^d}\left(\int_{R_2}\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z+w\|^2\right)\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z-w\|^2\right)dz\right)|\tilde{k}_1(\|w\|)|^q dw,
\end{aligned}$$

where we have defined

$$\begin{aligned}
R_1 &= \{z : -w_j \leq z_j \leq 2\lambda^{-1} + w_j,\ 1 \leq j \leq d\}, \\
R_2 &= \{z : w_j \leq z_j \leq 2\lambda^{-1} - w_j,\ 1 \leq j \leq d\}.
\end{aligned}$$

For the integral over R_1,

$$\begin{aligned}
&\int_{R_1} \exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z+w\|^2\right) \exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z-w\|^2\right) dz \\
&\geq \int_{R_1} \exp\left(\frac{q\lambda^{2-\alpha}}{2}\|z+w\|^2\right) dz \\
&= \int_{R_3} \exp\left(\frac{q\lambda^{2-\alpha}}{2}\|u\|^2\right) du \\
&= \prod_{j=1}^{d} \int_0^{2(\lambda^{-1}+w_j)} \exp\left(\frac{q\lambda^{2-\alpha}}{2} u_j^2\right) du_j,
\end{aligned}$$

where the inequality holds by the fact that $w \in [-\lambda^{-1}, 0]^d$, and the second line holds by making the substitution $u = z+w$ and defining $R_3 = \{u : 0 \leq u_j \leq 2(\lambda^{-1}+w_j),\ 1 \leq j \leq d\}$. Now, letting $v_j^2 := \frac{q\lambda^{2-\alpha}}{2} u_j^2$ gives

$$\begin{aligned}
\prod_{j=1}^{d} \int_0^{2(\lambda^{-1}+w_j)} \exp\left(\frac{q\lambda^{2-\alpha}}{2} u_j^2\right) du_j &= \prod_{j=1}^{d} \sqrt{\frac{2}{q\lambda^{2-\alpha}}} \int_0^{2(\lambda^{-1}+w_j)\sqrt{\frac{q\lambda^{2-\alpha}}{2}}} e^{v_j^2}\, dv_j \\
&= \prod_{j=1}^{d} \sqrt{\frac{2}{q\lambda^{2-\alpha}}} \exp\left(4(\lambda^{-1}+w_j)^2 \frac{q\lambda^{2-\alpha}}{2}\right) \mathcal{D}\left(2(\lambda^{-1}+w_j)\sqrt{\frac{q\lambda^{2-\alpha}}{2}}\right) \\
&\asymp \prod_{j=1}^{d} \sqrt{\frac{2}{q\lambda^{2-\alpha}}} \exp\left(4(\lambda^{-1}+w_j)^2 \frac{q\lambda^{2-\alpha}}{2}\right) \frac{1}{(\lambda^{-1}+w_j)\sqrt{\frac{q\lambda^{2-\alpha}}{2}}} \\
&\asymp \prod_{j=1}^{d} \exp\left(2q(1+\lambda w_j)^2/\lambda^{\alpha}\right) \frac{1}{q(\lambda^{-1}+w_j)\lambda^{2-\alpha}} \\
&= q^{-d}\lambda^{d(\alpha-1)} \prod_{j=1}^{d} \exp\left(2q(1+\lambda w_j)^2/\lambda^{\alpha}\right) \frac{1}{(1+\lambda w_j)},
\end{aligned}$$

where we have used the same properties of the Dawson function as in the upper bound.

Therefore, we have so far shown that

$$\int_{R_1} \exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z+w\|^2\right)\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z-w\|^2\right)dz$$
$$\gtrsim q^{-d}\lambda^{d(\alpha-1)}\prod_{j=1}^{d}\exp\left(2q(1+\lambda w_j)^2/\lambda^\alpha\right)\frac{1}{(1+\lambda w_j)}.$$

For the integral over R_2 a similar argument shows that

$$\int_{R_2} \exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z+w\|^2\right)\exp\left(\frac{q\lambda^{2-\alpha}}{4}\|z-w\|^2\right)dz$$
$$\gtrsim q^{-d}\lambda^{d(\alpha-1)}\prod_{j=1}^{d}\exp\left(2q(1-\lambda w_j)^2/\lambda^\alpha\right)\frac{1}{(1-\lambda w_j)}.$$

Putting the two bounds together, we have that

$$\begin{aligned}
&\sup_{x\in D}\int_D |k_\lambda(x,y)|^q dy \\
&\geq q^{-d}\lambda^{2d}\int_{[-\lambda^{-1},\lambda^{-1}]^d}\lambda^{d(\alpha-1)}\prod_{j=1}^{d}\exp\left(2q(1-\lambda|w_j|)^2/\lambda^\alpha\right)\frac{1}{(1-\lambda|w_j|)}|\tilde{k}_1(\|w\|)|^q dw \\
&\asymp q^{-d}\lambda^{d(1+\alpha)}\int_{\mathbb{R}^d} e^{2dq/\lambda^\alpha}|\tilde{k}_1(\|w\|)|^q dw \\
&= q^{-d}\lambda^{d(1+\alpha)}e^{2qd/\lambda^\alpha}A(d)\int_0^\infty r^{d-1}\tilde{k}_1(r)dr,
\end{aligned}$$

where the second line holds by the Dominated Convergence Theorem as $\lambda \to 0^+$. The final line follows as in the proof of the upper bound. □

Lemma 5.6.3. *The kernel k_λ satisfies Assumption 5.2.1 with*

$$R_q^q \asymp q^{-d}\lambda^{d(\alpha+1)}e^{2d\lambda^{-\alpha}}\int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr, \qquad \lambda \to 0^+.$$

Proof. The result follows by noting that

$$\sup_{x\in D}\int_D (k_\lambda(x,x)k_\lambda(y,y))^{(1-q)/2}|k_\lambda(x,y)|^q dy = \tilde{k}(0)\sup_{x\in D}\int_D \sigma_\lambda(x)\sigma_\lambda(y)|\tilde{k}_\lambda(x,y)|^q dy,$$

and using an identical approach to the one in Lemma 5.6.2. $\square$

Lemma 5.6.4. *It holds that*

$$\|C\| \asymp \lambda^{d(\alpha+1)} e^{2d\lambda^{-\alpha}} \int_0^\infty r^{d-1}\tilde{k}_1(r)dr, \qquad \lambda \to 0^+.$$

Proof. The proof follows in an identical way to that of [Al-Ghattas et al., 2023, Lemma 4.2], but invokes our novel characterization in Lemma 5.6.2 instead of [Al-Ghattas et al., 2023, Lemma 4.1]. $\square$

Proof of Theorem 5.2.10. For the sample covariance, note that by [Koltchinskii and Lounici, 2017, Theorem 9], for any $t \geq 1$ it holds with probability at least $1 - e^{-t}$ that

$$\frac{\|\widehat{C} - C\|}{\|C\|} \asymp \sqrt{\frac{r_2(C)}{N}} \vee \frac{r_2(C)}{N} \vee \sqrt{\frac{t}{N}} \vee \frac{t}{N}, \qquad r_2(C) := \frac{\mathrm{Tr}(C)}{\|C\|}.$$

By Lemmas 5.6.1 and 5.6.4, we have that for sufficiently small λ,

$$r_2(C) \asymp \frac{\tilde{k}_1(0)\lambda^{\alpha d} e^{2d/\lambda^\alpha}}{\lambda^{d(\alpha+1)} e^{2d/\lambda^\alpha}} \asymp \lambda^{-d}.$$

Therefore, with probability at least $1 - e^{-\log(\lambda^{-d})} = 1 - \lambda^d$

$$\frac{\|\widehat{C} - C\|}{\|C\|} \asymp \sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N} \vee \sqrt{\frac{\log(\lambda^{-d})}{N}} \vee \frac{\log(\lambda^{-d})}{N} = \sqrt{\frac{\lambda^{-d}}{N}} \vee \frac{\lambda^{-d}}{N}.$$

For the adaptive estimator, note first that the normalized process $\tilde{u}(x) = u(x)/k^{1/2}(x,x)$ is isotropic with covariance function $\tilde{k}_\lambda(\|x-y\|)$. Then by [Al-Ghattas et al., 2023, Lemma

4.3], we have that $\mathbb{E}[\sup_{x\in D}\tilde{u}(x)] \asymp \sqrt{\log(\lambda^{-d})}$ for sufficiently small λ. By Theorem 5.2.3, Lemma 5.6.4, and Lemma 5.6.3, we then have with probability at least $1-e^{-\log(\lambda^{-d})} = 1-\lambda^d$

$$
\begin{aligned}
\|\widehat{C}_{\hat{\rho}_N} - C\| &\lesssim R_q^q \left(\frac{\log(\lambda^{-d})}{\sqrt{N}}\right)^{1-q} \\
&\asymp q^{-d}\lambda^{d(\alpha+1)}e^{2d\lambda^{-\alpha}} \int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr \left(\frac{\log(\lambda^{-d})}{\sqrt{N}}\right)^{1-q} \\
&\asymp q^{-d}\|C\| \frac{\int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr}{\int_0^\infty r^{d-1}\tilde{k}_1(r) dr} \left(\frac{\log(\lambda^{-d})}{\sqrt{N}}\right)^{1-q}, \qquad \lambda \to 0^+.
\end{aligned}
$$

Rearranging yields the result with $c(q) := q^{-d}\frac{\int_0^\infty r^{d-1}\tilde{k}_1(r)^q dr}{\int_0^\infty r^{d-1}\tilde{k}_1(r) dr}$. □

5.7 Conclusions and Future Work

In this paper, we have studied covariance operator estimation under a novel sparsity assumption, developing a new nonasymptotic and dimension-free theory. Our model assumptions capture a particularly challenging class of nonstationary covariance models, where the marginal variance may vary significantly over the spatial domain. Adaptive threshold estimators are then shown to perform well over this class from both a theoretical and experimental perspective. The theory developed in this work as well as the connections made to both the (finite) high-dimensional literature and the functional data analysis literature open the door to many interesting avenues for future work, which we outline in the following:

- *Inference for covariance operators*: The techniques developed in this work, in particular the dimension-free bounds for the sample covariance and variance component functions open the door to lifting the finite high-dimensional inference theory to the infinite dimensional setting. The review paper Cai [2017] provides a detailed overview of the covariance inference literature in finite dimensions. Our analysis can likely be used to extend this theory to the covariance operator setting. To the best of our knowledge, the existing literature on inference for covariance operators (see for example Panaretos et al. [2010], Kashlak et al. [2019]) does not account for potential sparse structure in the underlying operators, though such structure can naturally be assumed in many cases of interest.

- *Estimating the covariance operator of a multi-valued Gaussian process*: As described in Section 5.1.1, the paper Fang et al. [2023] considers covariance estimation for multi-valued Gaussian processes under a sparsity assumption on the dependence between the individual component of the process. In contrast to our approach (see also Remark 5.2.2) their assumption does not impose any sparsity on each component covariance function. It would be interesting therefore to extend our analysis to the multi-valued functional data analysis setting where each component satisfies a local-type sparsity constraint, such as belonging to $\mathcal{K}_q^*$.

5.8 Acknowledgments

The authors are grateful for the support of the NSF CAREER award DMS-2237628, DOE DE-SC0022232, and the BBVA Foundation. The authors are also thankful to Jiaheng Chen and Subhodh Kotekal for inspiring discussions.

5.9 Additional Results

5.9.1 Additional Numerical Simulations

This appendix shows random draws and results of operator estimation in $d = 2$. The experimental set-up is identical to the $d = 1$ setting described in Section 5.4.2, however due to computational constraints in the higher dimensional setting, we use different settings for the experimental parameters. Specifically, our samples are generated by discretizing the domain $D = [0, 1]^2$ with a uniform mesh of $L = 10,000$ points. We consider a total of 10 choices of λ arranged uniformly in log-space and ranging from $10^{-2.2}$ to $10^{-0.1}$. As in the $d = 1$ case, the plots indicate that adaptive thresholding is a significant improvement over both the sample covariance and universal thresholding estimators in both the unweighted and weighted settings.

5.9.2 Sub-Gaussian Process Calculations

Let $v^{(1)}, v^{(2)}$ denote independent centered Gaussian processes both with covariance function k^v of the form (5.5), such that $k^v(x, y) = \sigma_\lambda(x)\sigma_\lambda(y)\tilde{k}_\lambda(x, y)$. Consider the process u obtained by transforming $v^{(1)}, v^{(2)}$ and let m^u, k^u denote the mean and covariance functions of u, respectively. In the following sections, we provide explicit expressions for m^u, k^u for various choices of transformation. These results are utilized in the sub-Gaussian portion of Section 5.2.3.

Sine Function

Let $u := \sin(v^{(1)})$. As $v^{(1)}$ is centered, and $\sin(\cdot)$ is an odd function, $\mathbb{E}[\sin(v^{(1)}(x))] = 0$ for any $x \in D$, implying $m^u = 0$. Further, since $\sin(\cdot)$ is bounded, u is a sub-Gaussian process. Note then that

$$\begin{aligned} k^u(x, y) &= \mathbb{E}[\sin(v^{(1)}(x)) \sin(v^{(1)}(y))] \\ &= \frac{1}{2}\mathbb{E}[\cos(v^{(1)}(x) - v^{(1)}(y))] - \frac{1}{2}\mathbb{E}[\cos(v^{(1)}(x) + v^{(1)}(y))], \end{aligned}$$

where we have made use of the identity $\sin(a)\sin(b) = \frac{1}{2}(\cos(a-b) - \cos(a+b))$. Recall that $v^{(1)}(x) - v^{(1)}(y) \sim N(0, k^v(x, x) + k^v(y, y) - 2k^v(x, y))$ and $v^{(1)}(x) + v^{(1)}(y) \sim N(0, k^v(x, x) +$

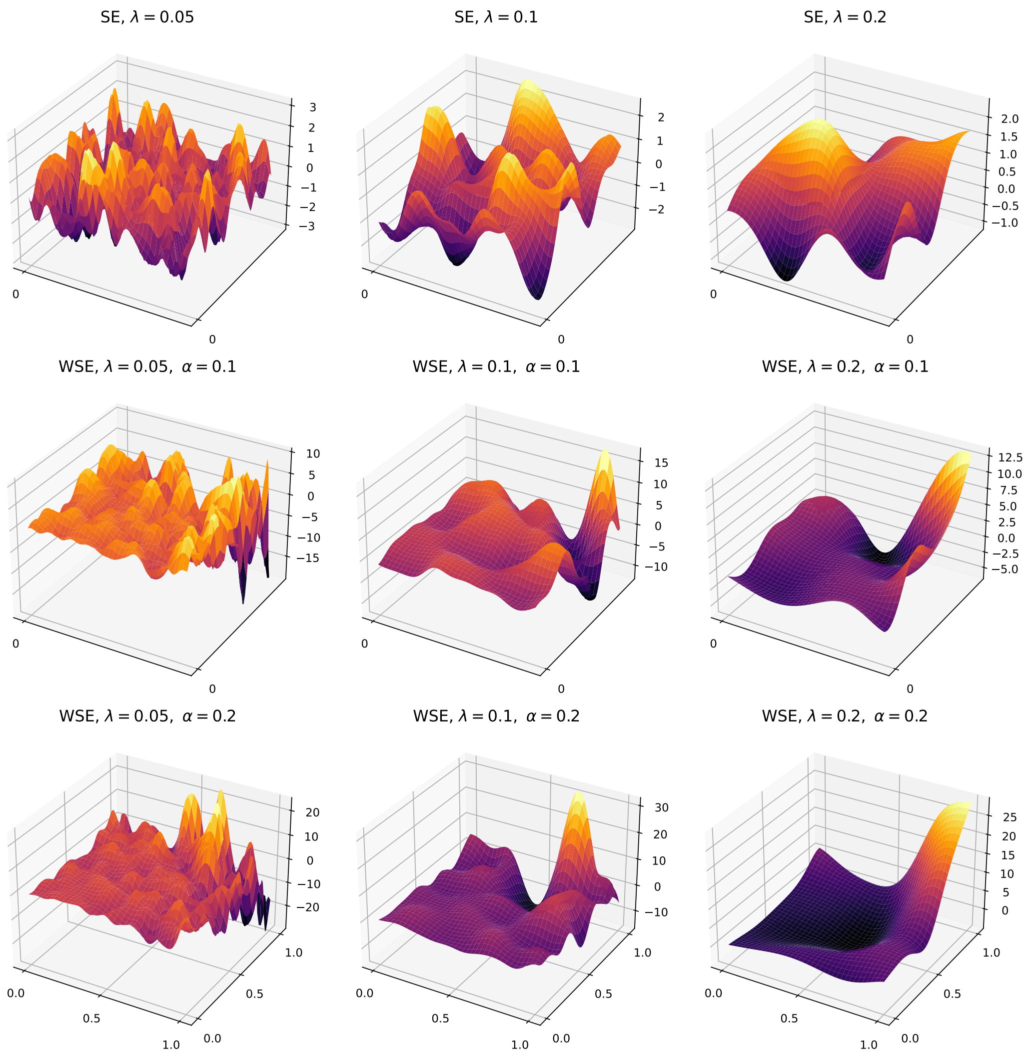

Figure 5.6: Draws from a centered Gaussian process on $D = [0,1]^2$ with covariance function SE in the first row, WSE($\alpha = 0.1$) in the second and WSE($\alpha = 0.2$) in the third, with varying λ parameter.

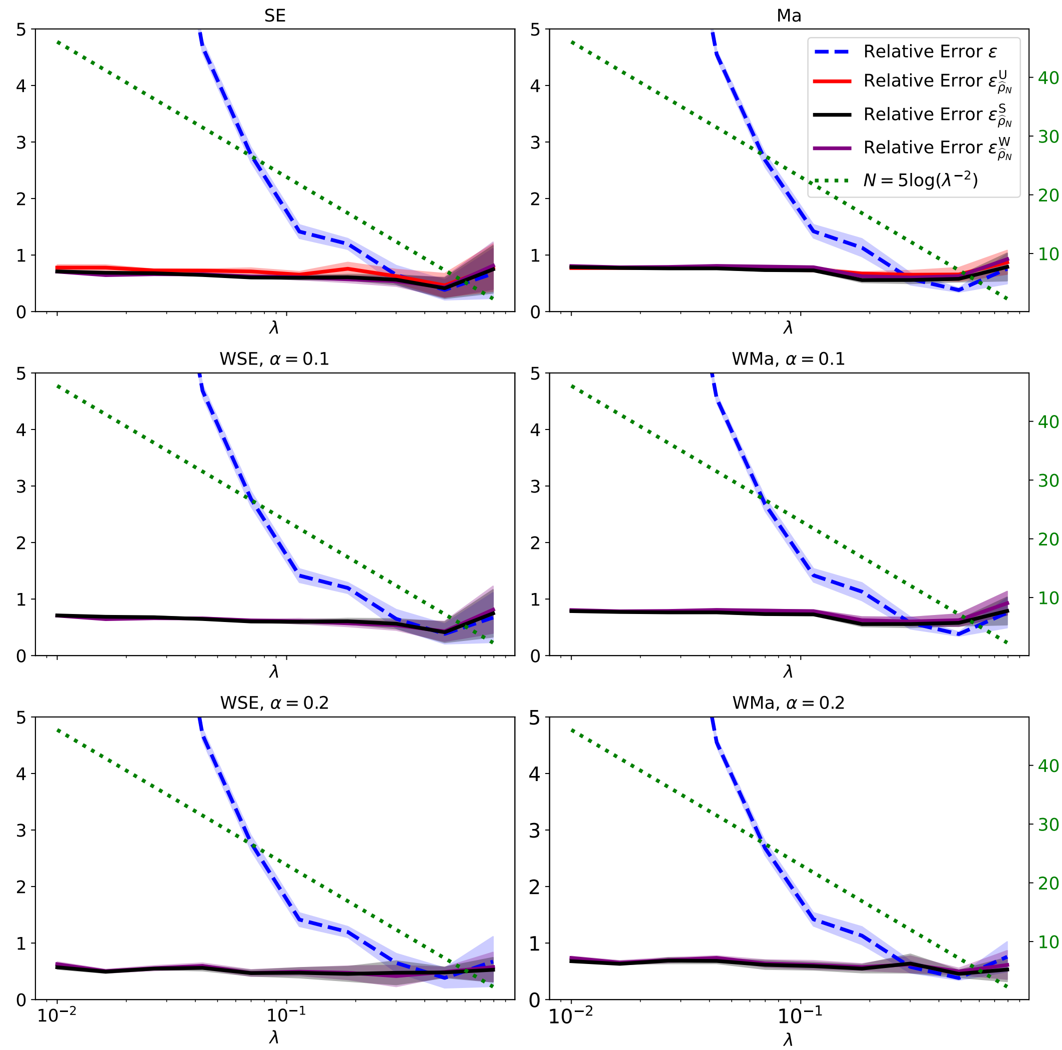

Figure 5.7: Plots of the average relative errors and 95% confidence intervals achieved by the sample (ε, dashed blue), universal thresholding ($\varepsilon^{\mathsf{U}}_{\hat{\rho}_N}$, red), sample-based adaptive thresholding ($\varepsilon^{\mathsf{S}}_{\hat{\rho}_N}$, black) and Wick's adaptive thresholding ($\varepsilon^{\mathsf{W}}_{\hat{\rho}_N}$, purple) covariance estimators based on a sample size (N, dotted green) for the (weighted) squared exponential (left) and (weighted) Matérn (right) covariance functions in $d = 2$ over 10 Monte-Carlo trials and 10 scale parameters λ ranging from 10^{-2} to $10^{-0.1}$. The first row corresponds to the unweighted covariance functions and is the only case in which the universal thresholding estimator is considered; the second and third rows correspond to the weighted variants with $\alpha = 0.1, 0.2$ respectively.

$k^v(y,y) + 2k^v(x,y))$. Further note that for $Z \sim N(0,\tau^2)$, $\mathbb{E}[\cos(Z)] = e^{-\tau^2/2}$. Therefore,

$$\begin{aligned} k^u(x,y) &= \frac{1}{2} e^{-\frac{1}{2}(k^v(x,x)+k^v(y,y)-2k^v(x,y)))} - \frac{1}{2} e^{-\frac{1}{2}(k^v(x,x)+k^v(y,y)+2k^v(x,y)))} \\ &= e^{-\frac{1}{2}(k^v(x,x)+k^v(y,y))} \sinh(k^v(x,y)) \\ &= e^{-\frac{1}{2}(\sigma_\lambda^2(x)+\sigma_\lambda^2(y))} \sinh(\sigma_\lambda(x)\sigma_\lambda(y)\tilde{k}_\lambda^v(x,y)). \end{aligned}$$

Absolute Value Function

Let $u := |v^{(1)}|$. Since $|\cdot|$ is Lipschitz, $u - m^u$ is a sub-Gaussian process. Direct calculation yields $m^u(x) = \mathbb{E}|v^{(1)}(x)| = \sqrt{\frac{2}{\pi} k^v(x,x)} = \sigma_\lambda(x)\sqrt{\frac{2}{\pi}}$ for any $x \in D$. Recall that for any $x, y \in D$, $(v^{(1)}(x), v^{(1)}(y))$ is a centered bi-variate Gaussian vector with $\mathbb{E}[v^{(1)}(x)v^{(1)}(y)] = \sigma_\lambda(x)\sigma_\lambda(y)\tilde{k}_\lambda(x,y)$, by [Li and Wei, 2009, Corollary 3.1]

$$\mathbb{E}[|v^{(1)}(x)||v^{(1)}(y)|] = \frac{2\sigma_\lambda(x)\sigma_\lambda(y)}{\pi} \left(\sqrt{1 - \tilde{k}_\lambda^2(x,y)} + \tilde{k}_\lambda(x,y) \sin^{-1}(\tilde{k}_\lambda(x,y)) \right).$$

Therefore

$$k^u(x,y) = \frac{2\sigma_\lambda(x)\sigma_\lambda(y)}{\pi} \left(\sqrt{1 - \tilde{k}_\lambda^2(x,y)} + \tilde{k}_\lambda(x,y) \sin^{-1}(\tilde{k}_\lambda(x,y)) - 1 \right).$$

Absolute Value × Sine Function

Let $u := (|v^{(1)}| - \mathbb{E}|v^{(1)}|)\sin(v^{(2)})$. By 5.9.2 and the fact that $\sin(v^{(2)})$ is bounded, u is the product of a sub-Gaussian process and a bounded process, and so is itself sub-Gaussian. Next, by independence of $v^{(1)}$ and $v^{(2)}$, $m^u = 0$. Recall that the product of two independent stochastic processes with covariance functions k^1, k^2 has covariance function $k^1 k^2$. Therefore, by the derivations in 5.9.2 and 5.9.2, for any $x, y \in D$,

$$k^u(x,y) = \frac{2\sigma_\lambda(x)\sigma_\lambda(y)}{\pi}\left(\sqrt{1-\tilde{k}_\lambda^2(x,y)} + \tilde{k}_\lambda(x,y)\sin^{-1}(\tilde{k}_\lambda(x,y))\right)$$
$$\times\, e^{-\frac{1}{2}(\sigma_\lambda^2(x)+\sigma_\lambda^2(y))}\sinh(\sigma_\lambda(x)\sigma_\lambda(y)\tilde{k}_\lambda^v(x,y)).$$

REFERENCES

S. I Aanonsen, G. Nævdal, D. S. Oliver, A. C. Reynolds, and B. Vallès. The ensemble Kalman filter in reservoir engineering–a review. *Spe Journal*, 14(03):393–412, 2009.

H. Abarbanel. *Predicting The Future: Completing Models Of Observed Complex Systems*. Springer, 2013.

P. Abdalla and N. Zhivotovskiy. Covariance estimation: Optimal dimension-free guarantees for adversarial corruption and heavy tails. *Journal of the European Mathematical Society*, 2024. doi:10.4171/JEMS/1505.

P. Abrahamsen. A review of Gaussian random fields and correlation functions. *Norsk Regnesentral/Norwegian Computing Center Oslo*, 1997.

M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, volume 55. US Government printing office, 1968.

S. Agapiou, S. Larsson, and A. M. Stuart. Posterior contraction rates for the bayesian approach to linear ill-posed inverse problems. *Stochastic Processes and Their Applications*, 123(10):3828–3860, 2013.

S. Agapiou, J. M. Bardsley, O. Papaspiliopoulos, and A. M. Stuart. Analysis of the Gibbs sampler for hierarchical inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):511–544, 2014.

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.

C. Agrell. Gaussian processes with linear operator inequality constraints. *Journal of Machine Learning Research*, 20(135):1–36, 2019.

B. Ahn, C. Kim, Y. Hong, and Hyunwoo J. Kim. Invertible monotone operators for normalizing flows. *Advances in Neural Information Processing Systems*, 35:16836–16848, 2022.

O. Al-Ghattas and D. Sanz-Alonso. Covariance Operator Estimation via Adaptive Thresholding. *arXiv preprint arXiv:2405.18562*, 2024a.

O. Al-Ghattas and D. Sanz-Alonso. Non-asymptotic analysis of ensemble Kalman updates: effective dimension and localization. *Information and Inference: A Journal of the IMA*, 13(1):iaad043, 2024b.

O. Al-Ghattas and D. Sanz-Alonso. Non-asymptotic analysis of ensemble Kalman updates: effective dimension and localization. *Information and Inference: A Journal of the IMA*, 13(1):iaad043, 2024c.

O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek. Covariance operator estimation: sparsity, lengthscale, and ensemble Kalman filters. *arXiv preprint arXiv:2310.16933*, 2023.

O. Al-Ghattas, J. Bao, and D. Sanz-Alonso. Ensemble Kalman filters with resampling. *SIAM/ASA Journal on Uncertainty Quantification*, 12(2):411–441, 2024a.

O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek. Optimal estimation of structured covariance operators. *arXiv preprint arXiv:2408.02109*, 2024b.

O. Al-Ghattas, J. Chen, D. Sanz-Alonso, and N. Waniorek. Supplement to "Covariance operator estimation: sparsity, lengthscale, and ensemble Kalman filters". *Bernoulli*, 2024c.

O. Al-Ghattas, J. Chen, and D. Sanz-Alonso. Sharp concentration of simple random tensors. *arXiv preprint arXiv:2502.16916*, 2025.

D. Alfke, D. Potts, M. Stoll, and T. Volkmer. NFFT meets Krylov methods: Fast matrix-vector products for the graph Laplacian of fully connected networks. *Frontiers in Applied Mathematics and Statistics*, 4:61, 2018.

J. L. Anderson. An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129(12):2884–2903, 2001.

J. L. Anderson and S. L. Anderson. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127 (12):2741–2758, 1999.

T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*, volume 2. Wiley New York, 1958.

H. Antil and A. K. Saibaba. Efficient algorithms for Bayesian inverse problems with Whittle–Matérn priors. *SIAM Journal on Scientific Computing*, 46(2):S176–S198, 2024.

R. Aoun, M. Banna, and P. Youssef. Matrix Poincaré inequalities and concentration. *Advances in Mathematics*, 371:107251, 2020.

S. Arridge, P. Maass, O. Öktem, and C. Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.

M. Asch, M. Bocquet, and M. Nodet. *Data Assimilation: Methods, Algorithms, and Applications*, volume 11. SIAM, 2016.

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.

Z. D. Bai and Y. Q. Yin. Limit of the smallest eigenvalue of a large dimensional sample covariance matrix. In *Advances In Statistics*, pages 108–127. World Scientific, 2008. doi:https://doi.org/10.1214/aop/1176989118.

R. Baptista, B. Hosseini, N. Kovachki, Y. Marzouk, and A. Sagiv. An approximation theory framework for measure-transport sampling algorithms. *Mathematics of Computation*, 2024a.

R. Baptista, Y. Marzouk, and O. Zahm. On the representation and learning of monotone triangular transport maps. *Foundations of Computational Mathematics*, 24(6):2063–2108, 2024b.

J. F Bard. *Practical Bilevel Optimization: Algorithms and Applications*, volume 30. Springer Science & Business Media, 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. Towards a theoretical foundation for Laplacian-based manifold methods. In *COLT*, volume 3559, pages 486–500. Springer, 2005.

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.

P. C. Bellec, G. Lecué, and A. B. Tsybakov. Slope meets Lasso: improved oracle bounds and optimality. *The Annals of Statistics*, 46(6B):3603–3642, 2018. doi:https://doi.org/10.1214/17-AOS1670.

T. Bengtsson, P. J. Bickel, B. Li, et al. Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. In *Probability and statistics: Essays in honor of David A. Freedman*, pages 316–334. Institute of Mathematical Statistics, 2008.

B. Bercu, E. Gassiat, and E. Rio. Concentration inequalities, large and moderate deviations for self-normalized empirical processes. *The Annals of Probability*, 30(4):1576–1604, 2002.

K. Bergemann and S. Reich. A localization technique for ensemble Kalman filters. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 136(648):701–707, 2010a.

K. Bergemann and S. Reich. A mollified ensemble Kalman filter. *Quarterly Journal of the Royal Meteorological Society*, 136(651):1636–1643, 2010b.

K. Bergemann and S. Reich. An ensemble Kalman-Bucy filter for continuous data assimilation. *Meteorologische Zeitschrift*, 127(5):1417–1440, 2012.

R. L. Berger and G. Casella. *Statistical Inference*. Duxbury, 2001.

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. Bertalmıo, 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. Beskos, A. Jasra, K. Law, R. Tempone, and Y. Zhou. Multilevel Sequential Monte Carlo Samplers. *Stochastic Processes and their Applications*, 127(5):1417–1440, 1994.

J. C. Bezdek, R. J. Hathaway, R. E. Howard, C. A. Wilson, and M. P. Windham. Local convergence analysis of a grouped variable version of coordinate descent. *Journal of Optimization Theory and Applications*, 54(3):471–477, 1987.

P. J. Bickel and E. Levina. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008a.

P. J. Bickel and E. Levina. Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1):199–227, 2008b.

P. J. Bickel, B. Li, and T. Bengtsson. Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh: Sharp failure rates for the bootstrap particle filter in high dimensions. *Institute of Mathematical Statistics*, pages 318–329, 1994.

P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009. doi:https://doi.org/10.1214/08-AOS620.

D. Bigoni, O. Zahm, A. Spantini, and Y. Marzouk. Greedy inference with layers of lazy maps. *arXiv preprint arXiv:1906.00031*, 2019.

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.

A. N. Bishop and P. Del Moral. On the mathematical theory of ensemble (linear-Gaussian) Kalman–Bucy filtering. *Mathematics of Control, Signals, and Systems*, pages 1–69, 2023.

C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

C. H. Bishop, B. J. Etherton, and S. J. Majumdar. Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Monthly Weather Review*, 129(3): 420–436, 2001.

D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

D. Blömker, C. Schillings, and P. Wacker. A strongly convergent numerical scheme from ensemble Kalman inversion. *SIAM Journal on Numerical Analysis*, 56(4):2537–2562, 2018.

D. Blömker, C. Schillings, P. Wacker, and S. Weissmann. Well posedness and convergence analysis of the ensemble Kalman inversion. *Inverse Problems*, 35(8):085007, 2019.

V. I. Bogachev. *Gaussian Measures*. American Mathematical Soc., 1998.

V. I. Bogachev, A. V. Kolesnikov, and K. V. Medvedev. Triangular transformations of measures. *Sbornik: Mathematics*, 196(3):309, 2005.

D. Bolin and K. Kirchner. The rational SPDE approach for Gaussian random fields with general smoothness. *Journal of Computational and Graphical Statistics*, (just-accepted): 1–27, 2019.

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.

S. C. Brenner and L. R. Scott. *The Mathematical Theory of Finite Element Methods*, volume 3. Springer.

J. Bröcker. Existence and uniqueness for four-dimensional variational data assimilation in discrete time. *SIAM Journal on Applied Dynamical Systems*, 16(1):361–374, 2013.

S. Brooks, A. Gelman, G. Jones, and X. Meng. Handbook of Markov chain Monte Carlo. *CRC press*, 2011.

S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.

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.

D. Burago, S. Ivanov, and Y. Kurylev. A graph discretization of the Laplace–Beltrami operator. *Journal of Spectral Theory*, 4(4):675–714, 2015.

G. Burgers, P. Jan van Leeuwen, and G. Evensen. Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126(6):1719–1724, 1998.

Q. Cai, J. Kang, and T. Yu. Bayesian network marker selection via the thresholded graph Laplacian Gaussian prior. *Bayesian Analysis*, 2018.

T. T. Cai. Global testing and large-scale multiple testing for high-dimensional covariance structures. *Annual Review of Statistics and Its Application*, 4:423–446, 2017.

T. T. Cai and W. Liu. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, 106(494):672–684, 2011. doi:https://doi.org/10.1198/jasa.2011.tm10560.

T. T. Cai and M. Yuan. Adaptive covariance matrix estimation through block thresholding. *The Annals of Statistics*, 40(4):2014–2042, 2012.

T. T. Cai and H. H. Zhou. Minimax estimation of large covariance matrices under ℓ_1-norm. *Statistica Sinica*, pages 1319–1349, 2012a.

T. T. Cai and H. H. Zhou. Optimal rates of convergence for sparse covariance matrix estimation. *The Annals of Statistics*, 40(5):2389–2420, 2012b.

T. T. Cai, Z. Ren, and H. H. Zhou. Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electronic Journal of Statistics*, 10:1–59, 2016.

E. Calvello, S. Reich, and A. M. Stuart. Ensemble Kalman methods: a mean field perspective. *arXiv preprint arXiv:2209.11371*, 2022.

D. Calvetti and E. Somersalo. *An Introduction to Bayesian Scientific Computing: Ten Lectures on Subjective Computing*, volume 2. Springer Science & Business Media, 2007.

D. Calvetti and E. Somersalo. Hypermodels in the bayesian imaging framework. *Inverse Problems*, 24(3):034013, 2008.

D. Calvetti, A. Pascarella, F. Pitolli, E. Somersalo, and B. Vantaggi. A hierarchical Krylov–Bayes iterative inverse solver for MEG with physiological preconditioning. *Inverse Problems*, 31(12):125005, 2015.

D. Calvetti, F. Pitolli, E. Somersalo, and B. Vantaggi. Bayes meets Krylov: Statistically inspired preconditioners for CGLS. *SIAM Review*, 60(2):429–461, 2018.

D. Calvetti, A. Pascarella, F. Pitolli, E. Somersalo, and B. Vantaggi. Brain activity mapping from MEG data via a hierarchical Bayesian algorithm with automatic depth weighting. *Brain topography*, 32(3):363–393, 2019a.

D. Calvetti, E. Somersalo, and A Strang. Hierachical Bayesian models and sparsity: L^2-magic. *Inverse problems*, 35(3):035003, 2019b.

D. Calvetti, M. Pragliola, and E. Somersalo. Sparsity promoting hybrid solvers for hierarchical Bayesian inverse problems. *SIAM Journal on Scientific Computing*, 42(6):A3761–A3784, 2020a.

D. Calvetti, M. Pragliola, E. Somersalo, and Alexander Strang. Sparse reconstructions from few noisy data: analysis of hierarchical Bayesian models with generalized gamma hyperpriors. *Inverse Problems*, 36(2):025010, 2020b.

F. Camacho and A. Demlow. L2 and pointwise a posteriori error estimates for fem for elliptic pdes on surfaces. *IMA Journal of Numerical Analysis*, 35(3):1199–1227, 2015.

Y. Canzani. Analysis on manifolds via the laplacian. *Lecture Notes available at: http://www. math. harvard. edu/canzani/docs/Laplacian. pdf*, 2013.

G. Carlier, A. Galichon, and F. Santambrogio. From knothe's transport to brenier's map and a continuation method for optimal transport. *SIAM Journal on Mathematical Analysis*, 41(6):2554–2576, 2010.

A. Carrassi, M. Bocquet, L. Bertino, and G. G. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, 9(5), 2018.

J. A. Carrillo and U. Vaes. Wasserstein stability estimates for covariance-preconditioned Fokker–Planck equations. *Nonlinearity*, 34(4):2275, 2021.

C. M. Carvalho, N. G Polson, and J. G. Scott. Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*, pages 73–80. PMLR, 2009.

N. K. Chada, M. A. Iglesias, L. Roininen, and A. M. Stuart. Parameterizations for ensemble Kalman inversion. *Inverse Problems*, 34(5):055009, 2018.

N. K. Chada, Y. Chen, and D. Sanz-Alonso. Iterative ensemble Kalman methods: A unified perspective with some new variants. *Foundations of Data Science*, 3(3):331–369, 2021.

M. A. J. 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.

S. Chatterjee and P. Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099–1135, 2018.

J. Chen and M. L. Stein. Linear-cost covariance functions for Gaussian random fields. *Journal of the American Statistical Association*, 118(541):147–164, 2023. doi:https://doi.org/10.1080/01621459.2021.1919122.

R. Y. Chen, A. Gittens, and J. A. Tropp. The masked sample covariance estimator: an analysis using matrix concentration inequalities. *Information and Inference: A Journal of the IMA*, 1(1):2–20, 2012.

X. Chen, M. Xu, and W. B. Wu. Covariance and precision matrix estimation for high-dimensional time series. *The Annals of Statistics*, 41(6):2994–3021, 2013.

Y. Chen and M. Anitescu. Scalable physics-based maximum likelihood estimation using hierarchical matrices. *SIAM/ASA Journal on Uncertainty Quantification*, 11(2):682–725, 2023. doi:https://doi.org/10.1137/21M1458880.

Y. Chen, D. Sanz-Alonso, and R. Willett. Autodifferentiable ensemble Kalman filters. *SIAM Journal on Mathematics of Data Science*, 4(2):801–833, 2022.

Y. Chen, D. Sanz-Alonso, and R. Willett. Reduced-order autodifferentiable ensemble Kalman filters. *Inverse Problems*, 39(12):124001, 2023.

N. Chopin and O. Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*, volume 4. Springer, 2020.

A. J. Chorin and M. Morzfeld. Conditions for successful data assimilation. *Journal of Geophysical Research: Atmospheres*, 118(20):11–522, 2013.

E. Cleary, A. Garbuno-Inigo, S. Lan, T. Schneider, and A. M. Stuart. Calibrate, emulate, sample. *Journal of Computational Physics*, 424:109716, 2021.

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.and 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.

S. Cotter, M. Dashti, and A. M. Stuart. Approximation of bayesian inverse problems for pde's. *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.

S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. MCMC methods for functions: Modifying old algorithms to make them faster. *Statistical Science*, 28(3):424–446, 2013.

K. Crane. Keenan's 3d model repository. URL `http://www.cs.cmu.edu/~kmcrane/Projects/ModelRepository`.

D. Crisan and B. Rozovskii. *The Oxford Handbook of Nonlinear Filtering*. Oxford University Press, 2011.

D. Crisan, P. Moral, and T. Lyons. Discrete filtering using branching and interacting particle systems. *Université de Toulouse. Laboratoire de Statistique et Probabilités [LSP]*, 1998.

G. Da Prato and J. Zabczyk. *Stochastic equations in infinite dimensions*. Cambridge university press, 2014.

P. Damlen, J. Wakefield, and S. Walker. Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(2):331–344, 1999.

M. Dashti and A. M. Stuart. Bayesian approach to inverse problems. *Handbook of Uncertainty Quantification*, pages 311–428, 2017a.

M. Dashti and A. M. Stuart. The bayesian approach to inverse problems. In *Handbook of uncertainty quantification*, pages 311–428. Springer, 2017b.

M. Dashti, K. Law, A. M. Stuart, and J. Voss. Map estimators and their consistency in Bayesian nonparametric inverse problems. *Inverse Problems*, 29(9), 2013.

I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 63(1):1–38, 2010.

N. de Freitas, P. Højen-Sørensen, M. I. Jordan, and S. Russell. Variational MCMC. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, UAI'01, page 120–127, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558608001.

M. V. de Hoop, N. B. Kovachki, N. H. Nelsen, and A. M. S. Convergence rates for learning linear operators from noisy data. *SIAM/ASA Journal on Uncertainty Quantification*, 11 (2):480–513, 2023. doi:https://doi.org/10.1137/21M144294.

G. P. Dehaene. Computing the quality of the Laplace approximation. *Neural Information Processing Systems*, 2017.

P. Del Moral. *Feynman-Kac Formulae*. Springer, 2004.

P. Del Moral and J. Tugaut. On the stability and the uniform propagation of chaos properties of ensemble Kalman–Bucy filters. *The Annals of Applied Probability*, 28(2):790–850, 2018.

J. Demmel. The componentwise distance to the nearest singular matrix. *SIAM Journal on Matrix Analysis and Applications*, 13(1):10–19, 1992.

A. Van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.

I. Diakonikolas and D. M Kane. *Algorithmic high-dimensional robust statistics*. Cambridge University Press, 2023. doi:https://doi.org/10.1017/9781108943161.

Z. Ding and Q. Li. Ensemble Kalman inversion: mean-field limit and convergence analysis. *Statistics and Computing*, 31(1):1–21, 2021.

L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

S. Dirksen. Tail Bounds via Generic Chaining. *Electronic Journal of Probability*, 20:1–29, 2015.

A. Doucet, A. M. Johansen, et al. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.

M. M. Dunlop, M. A. Iglesias, and A. M. Stuart. Hierarchical Bayesian level set inversion. *Statistics and Computing*, 27(6):1555–1584, 2017.

M. M. Dunlop, M. A. Girolami, A. M. Stuart, and A. L. Teckentrup. How deep are deep Gaussian processes? *The Journal of Machine Learning Research*, 19(1):2100–2145, 2018.

G. Dziuk and C. M. Elliott. Finite element methods for surface PDEs. *Acta Numerica*, 22: 289–396, 2013.

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.

N. El Karoui. Operator norm consistent estimation of large-dimensional sparse covariance matrices. *The Annals of Statistics*, 36(6):2717–2756, 2008.

T. A. El Moselhy and Y. M. Marzouk. Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23):7815–7850, 2012.

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. England, M. Hanke, and A. Neubauer. *Regularization of inverse problems*. Springer Science and Business Media, 1996.

O. G. Ernst, B. Sprungk, and H.-J. Starkloff. Analysis of the ensemble and polynomial chaos Kalman filters in Bayesian inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):823–851, 2015.

L. C. Evans. *Partial Differential Equations*, volume 19. American Mathematical Soc., 2010.

G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, 99(c5):10143–10162, 1995.

G. Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53:343–367, 2003.

G. Evensen. Sampling strategies and square root analysis schemes for the EnKF. *Ocean Dynamics*, 54(6):539–560, 2004.

G. Evensen. *Data Assimilation: the Ensemble Kalman Filter*. Springer Science and Business Media, 2009.

G. Evensen and P. Van Leeuwen. Assimilation of Geosat altimeter data for the Agulhas current using the ensemble Kalman filter with a quasigeostrophic model. *Monthly Weather Review*, 124(1):85–96, 1996.

G. Evensen, F. C. Vossepoel, and P. J. van Leeuwen. *Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter Estimation Problem.* Springer Nature, 2022.

J. Fan, Y. Fan, and J. Lv. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197, 2008.

Q. Fang, S. Guo, and X. Qiao. Adaptive functional thresholding for sparse covariance function estimation in high dimensions. *Journal of the American Statistical Association*, (546): 1473–1485, 2023. doi:10.1080/01621459.2023.2200522.

A. Farchi and M. Bocquet. On the efficiency of covariance localisation of the ensemble Kalman filter using augmented ensembles. *Frontiers in Applied Mathematics and Statistics*, page 3, 2019.

X. Fernique. Intégrabilité des vecteurs gaussiens. *CR Acad. Sci. Paris Serie A*, 270:1698–1699, 1970.

X. M. Fernique, J. P. Conze, J. Gani, and X Fernique. *Regularité des trajectoires des fonctions aléatoires gaussiennes.* Springer, 1975. doi:http://dx.doi.org/10.1007/BFb0080190.

J. Franklin. Well-posed stochastic extensions of ill-posed linear problems. *Journal of mathematical analysis and applications*, 31(3):682–716, 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.

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.

R. Furrer and T. Bengtsson. Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, 98(2):227–255, 2007.

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. doi:http://dx.doi.org/10.1198/106186006X132178.

D. Gamerman and H. Lopes. Markov chain Monte Carlo: stochastic simulation for Bayesian inference. *CRC Press*, 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.

A. Garbuno-Inigo, F. Hoffmann, W. Li, and A. M. Stuart. Interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM Journal on Applied Dynamical Systems*, 19(1):412–441, 2020.

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, 2018a.

N García Trillos and D. Sanz-Alonso. The Bayesian update: variational formulations and gradient flows. *Bayesian Analysis*, 2018b.

N. García Trillos and D. Slepčev. Continuum limit of total variation on point clouds. *Archive for rational mechanics and analysis*, 220(1):193–241, 2016.

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, Z. Kaplan, and D. Sanz-Alonso. Variational characterizations of local entropy and heat regularization in deep learning. *Entropy*, 21(5):511, 2019b.

N. García Trillos, Daniel Sanz-Alonso, and Ruiyi Yang. Local regularization of noisy point clouds: Improved global geometric estimates and data analysis. *Journal of Machine Learning Research*, 20(136):1–37, 2019c.

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.

M. G. Genton. Classes of kernels for machine learning: a statistics perspective. *Journal of Machine Learning Research*, 2(Dec):299–312, 2001.

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. W. van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*, volume 44. Cambridge University Press, 2017. doi:http://dx.doi.org/10.1017/9781139029834.

A. Gibbs and F. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.

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.

M. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 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.

E. Giné and R. Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press, 2021.

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.

M. S. Gockenbach. *Understanding and implementing the finite element method*, volume 97. Siam, 2006.

J. Goes, G. Lerman, and B. Nadler. Robust sparse covariance estimation by thresholding Tyler's M-estimator. *The Annals of Statistics*, 48(1):86–110, 2020. doi:10.1214/18-AOS1793.

Y. Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50:265–289, 1985. doi:http://dx.doi.org/10.1007/BF02759761.

I. F. Gorodnitsky and B. D. Rao. Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm. *IEEE Transactions on signal processing*, 45(3):600–616, 1997.

G. A. Gottwald and A. J. Majda. A mechanism for catastrophic filter divergence in data assimilation for sparse observation networks. *Nonlinear Processes in Geophysics*, 20(5): 705–712, 2013.

P. J. Green. Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *Journal of the Royal Statistical Society: Series B (Methodological)*, 46(2):149–170, 1984.

D. V. Griffiths, J. Huang, and G. A. Fenton. Influence of spatial variability on slope reliability using 2-D random fields. *Journal of geotechnical and geoenvironmental engineering*, 135 (10):1367–1378, 2009.

Y. Gu and D. S. Oliver. An iterative ensemble Kalman filter for multiphase fluid flow data assimilation. *Spe Journal*, 12(04):438–446, 2007.

S. Guo, D. Li, X. Qiao, and Y. Wang. From sparse to dense functional data in high dimensions: Revisiting phase transitions from a non-asymptotic perspective. *Journal of Machine Learning Research*, 26(15):1–40, 2025.

P. A. Guth, C. Schillings, and S. Weissmann. Ensemble Kalman filter for neural network based one-shot inversion. *arXiv preprint arXiv:2005.02039*, 2020.

M. Hairer. An introduction to stochastic pdes. *arXiv preprint arXiv:0907.4178*, 2009.

M. Hairer, A. M. Stuart, J. Voss, and P. Wiberg. Analysis of SPDEs arising in path sampling. Part I: The Gaussian case. *Communications in Mathematical Sciences*, 3(4):587–603, 2013.

M. Hanke. A regularizing Levenberg-Marquardt scheme, with applications to inverse groundwater filtration problems. *Inverse Problems*, 13(1):79–95, 1997.

A. Harbey. *Forecasting, structural time series models and the Kalman filter*. Cambridge university press, 1964.

D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, 16(12):2639–2664, 2004.

J. Harlim and A. J. Majda. Catastrophic filter divergence in filtering nonlinear dissipative systems. *Communications in Mathematical Sciences*, 8(1):27–43, 2010.

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.

J. Harlim, S. W. Jiang, H. Kim, and D. Sanz-Alonso. Graph-based prior and forward models for inverse problems on manifolds with boundaries. *Inverse Problems*, 38(3):035006, 2022.

K. Hayden, E. Olson, and E. Titi. Discrete data assimilation in the Lorenz and 2D Navier–Stokes equations. *Physica D: Nonlinear Phenomena*, 240(18):1416–1425, 2011.

M. Hein and J. Y. Audibert. Intrinsic dimensionality estimation of submanifolds in r d. In *Proceedings of the 22nd international conference on Machine learning*, pages 289–296. ACM, 2005.

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.

T. Helin and M. Burger. Maximum a posteriori probability estimates in infinite-dimensional bayesian inverse problems. *Inverse Problems*, 31(8), 2015.

M. Herty and G. Visconti. Kinetic methods for inverse problems. *Kinetic & Related Models*, 12(5):1109, 2019.

M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.

R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2012.

P. L. Houtekamer and J. Derome. Methods for ensemble prediction. *Monthly Weather Review*, 123(7):2181–2196, 1995.

P. L. Houtekamer and H. L. Mitchell. Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126(3):796–811, 1998.

P. L. Houtekamer and H. L. Mitchell. A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129(1):123–137, 2001.

P. L. Houtekamer and F. Zhang. Review of the ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144(12):4489–4532, 2016.

J. K. Hunter and B. Nachtergaele. *Applied Analysis*. World Scientific Publishing Company, 2001. doi:http://dx.doi.org/10.1142/4319.

M. A. Iglesias. Iterative regularization for ensemble data assimilation in reservoir models. *Computational Geosciences*, 19:177–212, 2015.

M. A. Iglesias. A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems. *Inverse Problems*, 32(2):025002, 2016.

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.

M. A. Iglesias, K. J. H. Law, and A. M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.

N. J. Irons, M. Scetbon, S. Pal, and Z. Harchaoui. Triangular flows for generative modeling: Statistical consistency, smoothness classes, and fast rates. In *International Conference on Artificial Intelligence and Statistics*, pages 10161–10195. PMLR, 2022.

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.

P. Jaini, K. A. Selby, and Y. Yu. Sum-of-squares polynomial flow. In *International Conference on Machine Learning*, pages 3009–3018. PMLR, 2019.

P. Jaini, I. Kobyzev, Y. Yu, and M. Brubaker. Tails of lipschitz triangular flows. In *International Conference on Machine Learning*, pages 4673–4681. PMLR, 2020.

G. M. James and C. A. Sugar. Clustering for sparsely sampled functional data. *Journal of the American Statistical Association*, 98(462):397–408, 2003. doi:http://dx.doi.org/10.1198/016214503000189.

G. M. James, T. J. Hastie, and C. A. Sugar. Principal component models for sparse functional data. *Biometrika*, 87(3):587–602, 2000. doi:http://dx.doi.org/10.1093/biomet/87.3.587.

S. W. Jiang and J. Harlim. Ghost point diffusion maps for solving elliptic pdes on manifolds with classical boundary conditions. *Communications on Pure and Applied Mathematics*, 76(2):337–405, 2023.

J. Jin, Y. Lu, J. Blanchet, and L. Ying. Minimax optimal kernel operator learning via multilevel training. In *The Eleventh International Conference on Learning Representations*, 2022.

M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.

J. Kaipo and E. Somersalo. Statistical and Computational Inverse Problems. *Springer Science & Business Media*, 160, 2006.

R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2003. doi:https://doi.org/10.1017/CBO9780511802270.

A. B. Kashlak, J. A. D. Aston, and R. Nickl. Inference on covariance operators via concentration inequalities: k-sample tests, classification, and clustering via Rademacher complexities. *Sankhya A*, 81:214–243, 2019.

M. Katzfuss, J. R. Stroud, and C. K. Wikle. Understanding the ensemble Kalman filter. *The American Statistician*, 70(4):350–357, 2016.

D. Kelly and A. M. Stuart. Well-posedness and accuracy of the ensemble Kalman filter in discrete and continuous time. *Nonlinearity*, 27(10), 2014.

D. Kelly and A. M. Stuart. Ergodicity and accuracy of optimal particle filters for bayesian data assimilation. *Chinese Annals of Mathematics, Series B*, 40(5):811–842, 2019.

D. Kelly, A. J. Majda, and X. T. Tong. Concrete ensemble Kalman filters with rigorous catastrophic filter divergence. *Proceedings of the National Academy of Sciences*, 112(34): 10589–10594, 2015.

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. Kim, D. Sanz-Alonso, and A. Strang. Hierarchical ensemble Kalman methods with sparsity-promoting generalized gamma hyperpriors. *Foundations of Data Science*, 5(3): 366–388, 2023.

H. Kim, D. Sanz-Alonso, and R. Yang. Optimization on manifolds via graph gaussian processes. *SIAM Journal on Mathematics of Data Science*, 6(1):1–25, 2024.

D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29, 2016.

B. Klartag and S. Mendelson. Empirical processes and random projections. *Journal of Functional Analysis*, 225(1):229–245, 2005.

B. Knapik, A. van der Vaart, and J. van Zanten. Bayesian inverse problems with Gaussian priors. *The Annals of Statistics*, 39(5):2626–2657, 2011.

H. Knothe. Contributions to the theory of convex bodies. *Michigan Mathematical Journal*, 4(1):39–52, 1957.

V. Koltchinskii and K. Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133, 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. Kontorovich and M. Raginsky. Concentration of measure without independence: a unified approach via the martingale method. In *Convexity and Concentration*, pages 183–210. Springer, 2017.

N. B Kovachki and A. M. Stuart. Ensemble Kalman inversion: a derivative-free technique for machine learning tasks. *Inverse Problems*, 35(9):095005, 2019.

N. B. Kovachki, S. Lanthaler, and A. M. Stuart. Operator Learning: Algorithms and Analysis. *arXiv preprint arXiv:2402.15715*, 2024.

A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. Automatic differentiation variational inference. *The Journal of Machine Learning Research*, 18(1):430–474, 2017.

T. Kühn. Covering numbers of Gaussian reproducing kernel Hilbert spaces. *Journal of Complexity*, 27(5):489–499, 2011.

E. Kwiatkowski and J. Mandel. Convergence of the square root ensemble Kalman filter in the large ensemble limit. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):1–17, 2015.

L. E. Lehmann, and G. Casella. *Theory of Point Estimation.* Springer Science & Business Media, 2006.

A. Lang, J. Potthoff, M. Schlather, and D. Schwab. Continuity of random fields on riemannian manifolds. *arXiv preprint arXiv:1607.05859*, 2016.

T. Lange and W. Stannat. Mean field limit of Ensemble Square Root filters-discrete and continuous time. *Foundations of Data Science*, 3(3):563–588, 2021.

J. Latz. Analysis of stochastic gradient descent in continuous time. *Statistics and Computing*, 31(4):39, 2021.

K. J. H. Law and A. M. Stuart. Evaluating data assimilation algorithms. *Monthly Weather Review*, 140(11):3757–3782, 2012.

K. J. H. Law, A. M. Stuart, and K. Zygalakis. *Data Assimilation.* Springer, 2015.

K. J. H Law, D. Sanz-Alonso, A. Shukla, and A. M. Stuart. Filter accuracy for the Lorenz 96 model: Fixed versus adaptive observation operators. *Physica D: Nonlinear Phenomena*, 325:1–13, 2016a.

K. J. H. Law, H. Tembine, and R. Tempone. Deterministic mean-field ensemble Kalman filtering. *SIAM Journal on Scientific Computing*, 38(3):A1251–A1279, 2016b.

W. G. Lawson and J. A. Hansen. Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. *Monthly Weather Review*, 132(8):1966–1981, 2004.

F. Le Gland, V. Monbet, and V.-D. Tran. *Large sample asymptotics for the ensemble Kalman filter*. PhD thesis, INRIA, 2009.

M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.

Y. Lee. ℓ_p regularization for ensemble Kalman inversion. *SIAM Journal on Scientific Computing*, 43(5):A3417–A3437, 2021.

P. Van Leeuwen. Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Quarterly Journal of the Royal Meteorological Society*, 136(653):1991–1999, 2010.

P. Van Leeuwen, Y. Cheng, and S. Reich. *Nonlinear Data Assimilation*. Springer, 2015.

O. Leeuwenburgh, G. Evensen, and L. Bertino. The impact of ensemble filter definition on the assimilation of temperature profiles in the tropical Pacific. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 131(613):3291–3300, 2005.

M. S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalised random variables. *Inverse Problems*, 5(4):599, 1989a.

M. S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalised random variables. *Inverse Problems*, 5(4):599, 1989b.

A. J. Lemonte and G. M. Cordeiro. The exponentiated generalized inverse Gaussian distribution. *Statistics & Probability Letters*, 81(4):506–517, 2011.

V. P. Leonov and A. N. Shiryaev. On a method of calculation of semi-invariants. *Theory of Probability & its applications*, 4(3):319–329, 1959.

E. Levina and R. Vershynin. Partial estimation of covariance matrices. *Probability Theory and Related Fields*, 153(3-4):405–419, 2012.

C. Li and H. Li. Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics*, 24(9):1175–1182, 2008.

G. Li and A. C. Reynolds. An iterative ensemble Kalman filter for data assimilation. In *SPE annual technical conference and exhibition*. Society of Petroleum Engineers, 2007.

J. Li and D. Xiu. On numerical properties of the ensemble Kalman filter for data assimilation. *Computer Methods in Applied Mechanics and Engineering*, 197(43-44):3574–3583, 2008.

Wenbo V Li and Ang Wei. Gaussian integrals involving absolute value functions. In *High dimensional probability V: the Luminy volume*, volume 5, pages 43–60. Institute of Mathematical Statistics, 2009.

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.

C. Lieberman, C. Willcox, and O. Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2535–2542, 2010.

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.

D. V. Lindley and A. F. M. Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–41, 1972.

T. Lindvall. *Lectures on the Coupling Method*. Springer, 2002.

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.

F. Liu, S. Chakraborty, F. Li, Y. Liu, and A. C. Lozano. Bayesian regularization via graph Laplacian. *Bayesian Analysis*, 9(2):449–474, 2014.

J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Science & Business Media, 2008.

D. M. Livings, S. L. Dance, and N. K. Nichols. Unbiased ensemble square root filters. *Physica D: Nonlinear Phenomena*, 237(8):1021–1028, 2008.

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.

E. N. Lorenz. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, 20(2): 130–141, 1963.

E. N. Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on predictability*, volume 1. Reading, 1996.

K. Lounici. High-dimensional covariance matrix estimation with missing observations. *Bernoulli*, 20(3):1029–1058, 2014.

L. Maestrini, R. G. Aykroyd, and M. P. Wand. A variational inference framework for inverse problems. *Computational Statistics & Data Analysis*, 202:108055, 2025.

A. J. Majda and J. Harlim. *Filtering Complex Turbulent Systems*. Cambridge University Press, 2012.

A. J. Majda and X. T. Tong. Performance of ensemble Kalman filters in large dimensions. *Communications on Pure and Applied Mathematics*, 71(5):892–937, 2018.

A. J. Majda and X. Wang. *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*. Cambridge University Press, 2006.

J. Mandel, L. Cobb, and J. D. Beezley. On the convergence of the ensemble Kalman filter. *Applications of Mathematics*, 56(6):533–541, 2011.

A. Mandelbaum. Linear estimators and measurable linear transformations on a Hilbert space. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(3):385–397, 1984a.

A. Mandelbaum. Linear estimators and measurable linear transformations on a hilbert space. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(3):385–397, 1984b.

Y. Marzouk and D. Xiu. A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics*, 6(4):826–847, 2009.

Y. Marzouk, T. Moselhy, M. Parno, and A. Spantini. Sampling via measure transport: an introduction. *Handbook of Uncertainty Quantification*, pages 1–41, 2016.

A. Matthews. *Scalable Gaussian process inference using variational methods*. PhD thesis, University of Cambridge, 2017.

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.

S. Mendelson. Empirical processes with a bounded ψ_1 diameter. *Geometric and Functional Analysis*, 20(4):988–1027, 2010.

S. Mendelson. Upper bounds on product and multiplier empirical processes. *Stochastic Processes and their Applications*, 126(12):3652–3680, 2016.

S. Mendelson and N. Zhivotovskiy. Robust covariance estimation under L_4-L_2 norm equivalence. 2020.

M. Mercker and A. Marciniak-Czochra. Bud-neck scaffolding as a possible driving force in escrt-induced membrane budding. *Biophysical journal*, 108(4):833–843, 2015.

N. Mollenhauer, N. Mücke, and T. J. Sullivan. Learning linear operators: Infinite-dimensional regression as a well-behaved non-compact inverse problem. *arXiv preprint arXiv:2211.08875*, 2022.

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.

M. Morzfeld, D. Hodyss, and C. Snyder. What the collapse of the ensemble Kalman filter tells us about particle filters. *Tellus A: Dynamic Meteorology and Oceanography*, 69(1): 1283809, 2017.

J. L. Mueller and S. Siltanen. *Linear and nonlinear inverse problems with practical applications*, volume 10. Siam, 2012.

I. Myrseth, J. Sætrom, and H. Omre. Resampling the ensemble Kalman filter. *Computers & Geosciences*, 55:44–53, 2013.

C. Naesseth, S. Linderman, R. Ranganath, and D. Blei. Variational sequential Monte Carlo. In *International conference on artificial intelligence and statistics*, pages 968–977. PMLR, 2018.

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.

F. Nobile and F. Tesei. A Multi Level Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients. *Stoch. Partial Differ. Equ. Anal. Comput.*, 3:398–444, 2015. doi:http://dx.doi.org/10.1007/s40072-015-0055-9.

N. Nüsken and S. Reich. Note on interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler by Garbuno-Inigo, Hoffmann, Li and Stuart. *arXiv preprint arXiv:1908.10890*, 2019.

J. T. Oden and J. N. Reddy. *An Introduction to the Mathematical Theory of Finite Elements*. Courier Corporation, 2012.

E. Ott, B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich, M. Corazza, E. Kalnay, D. J. Patil, and J. A. Yorke. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A: Dynamic Meteorology and Oceanography*, 56(5):415–428, 2004.

A. B. Owen. Statistically efficient thinning of a Markov chain sampler. *Journal of Computational and Graphical Statistics*, 26(3):738–744, 2017.

C. J. Paciorek and M. J. Schervish. Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics: The official journal of the International Environmetrics Society*, 17(5):483–506, 2006.

V. M. Panaretos, D. Kraus, and J. H. Maddocks. Second-order comparison of Gaussian random functions and the geometry of DNA minicircles. *Journal of the American Statistical Association*, 105(490):670–682, 2010.

G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research*, 22(1):2617–2680, 2021.

O. Papaspiliopoulos and M. Ruggiero. Optimal filtering and the dual process. *Bernoulli*, 20 (4):1999–2019, 2014.

O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, pages 59–73, 2007.

A. Papoulis. *Probability and statistics.* Prentice-Hall, Inc., 1990.

R. Petrie. Localization in the ensemble Kalman filter. *MSc Atmosphere, Ocean and Climate University of Reading*, 2008.

J. Ping, D. Zhang, et al. 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.

M. Pourahmadi. *High-dimensional covariance estimation: with high-dimensional data*. John Wiley & Sons, 2013.

X. Qiao, C. Qian, G. M. James, and S. Guo. Doubly functional graphical models in high dimensions. *Biometrika*, 107(2):415–431, 2020. doi:10.1093/biomet/asz072.

J. O. Ramsay and J. B. Ramsey. Functional data analysis of the dynamics of the monthly index of nondurable goods production. *Journal of Econometrics*, 107(1-2):327–344, 2002.

S. Reich. A dynamical systems framework for intermittent data assimilation. *IBIT Numerical Mathematics*, 51(1):235–249, 2017.

S. Reich and C. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, 2015.

S. Remes, M. Heinonen, and S. Kaski. Non-stationary spectral kernels. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4645–4654, 2017.

A. C. Reynolds, M. Zafari, and G. Li. Iterative forms of the ensemble Kalman filter. In *ECMOR X-10th European conference on the mathematics of oil recovery*, pages cp–23. European Association of Geoscientists & Engineers, 2006.

H. E. Robbins. An empirical Bayes approach to statistics. In *Breakthroughs in Statistics*, pages 388–394. Springer, 1992.

G. O. Roberts and S. K. Sahu. Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):291–317, 1997.

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.

M. Rosenblatt. Remarks on a multivariate transformation. *The annals of mathematical statistics*, 23(3):470–472, 1952.

M. Roth, G. Hendeby, C. Fritsche, and F. Gustafsson. The ensemble Kalman filter: a signal processing perspective. *EURASIP Journal on Advances in Signal Processing*, 2017(1): 1–16, 2017.

A. J. Rothman, E. Levina, and J. Zhu. Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association*, 104(485):177–186, 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 and P. Guttorp. Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association*, 87(417):108–119, 1992.

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 A. M. Stuart. Long-time asymptotics of the filtering distribution for partially observed chaotic dynamical systems. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):1200–1220, 2015.

D. Sanz-Alonso and Z. Wang. Bayesian update with importance sampling: Required sample size. *Entropy*, 23(1):22, 2021.

D. Sanz-Alonso and N. Waniorek. Analysis of a computational framework for Bayesian inverse problems: Ensemble Kalman updates and MAP estimators under mesh refinement. *SIAM/ASA Journal on Uncertainty Quantification*, 12(1):30–68, 2024.

D. Sanz-Alonso and R. Yang. The SPDE approach to Matérn fields: Graph representations. *Statistical Science*, 37(4):519–540, 2022a.

D. Sanz-Alonso and R. Yang. Unlabeled data help in graph-based semi-supervised learning: a bayesian nonparametrics perspective. *Journal of Machine Learning Research*, 23(97): 1–28, 2022b.

D. Sanz-Alonso, A. M. Stuart, and A. Taeb. *Inverse Problems and Data Assimilation*, volume 107. Cambridge University Press, 2023a.

D. Sanz-Alonso, Andrew Stuart, and Armeen Taeb. *Inverse Problems and Data Assimilation*. London Mathematical Society Student Texts. Cambridge University Press, 2023b.

S. Särkkä. Linear operators and stochastic partial differential equations in gaussian process regression. In *Artificial Neural Networks and Machine Learning–ICANN 2011: 21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14-17, 2011, Proceedings, Part II 21*, pages 151–158. Springer, 2011.

S. Särkkä. *Bayesian Filtering and Smoothing*, volume 3. Cambridge University Press, 2013.

S. Särkkä and L. Svensson. *Bayesian Filtering and Smoothing*, volume 17. Cambridge University Press, 2023.

C. Schillings and A. M. Stuart. Analysis of the ensemble Kalman filter for inverse problems. *SIAM Journal on Numerical Analysis*, 55(3):1264–1290, 2017.

C. Schillings, B. Sprungk, and P. Wacker. On the convergence of the Laplace approximation and noise-level-robustness of Laplace-based Monte Carlo methods for Bayesian inverse problems. *Numerische Mathematik*, 145(4):915–971, 2020.

C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in UQ. *Analysis and Applications*, 17(01): 19–55, 2019.

Z. Shi and J. Sun. Convergence of the point integral method for poisson equation on point cloud. *arXiv preprint arXiv:1403.2141*, 2014.

A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.

C. Snyder. Particle filters, the “optimal” proposal and high-dimensional systems. In *Proceedings of the ECMWF Seminar on Data Assimilation for Atmosphere and Ocean*, 2011.

C. Snyder, T. Bengtsson, and M. Morzfeld. Performance bounds for particle filters using the optimal proposal. *Monthly Weather Review*, 143(11):4750–4761, 2015.

C. Snyder, T. Bengtsson, P. J. Bickel, and J. L. Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, 2016.

A. Spantini, D. Bigoni, and Y. Marzouk. Inference via low-dimensional couplings. *The Journal of Machine Learning Research*, 19(1):2639–2709, 2018.

A. Spantini, R. Baptista, and Y. Marzouk. Coupling techniques for nonlinear ensemble filtering. *SIAM Review*, 64(4):921–953, 2022.

C. Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the sixth Berkeley symposium on mathematical statistics and probability, volume 2: Probability theory*, volume 6, pages 583–603. University of California Press, 1972.

M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging.* Springer Science & Business Media, 2012.

G. Strang and G. J. Fix. *An Analysis of the Finite Element Method.* 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.

M. Talagrand. *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes.* Springer, 2005.

M. Talagrand. *Upper and Lower Bounds for Stochastic Processes*, volume 60. Springer, 2014.

M. Talagrand. *Upper and Lower Bounds for Stochastic Processes: Decomposition Theorems*, volume 60. Springer Nature, 2022. doi:https://doi.org/10.1007/978-3-030-82595-9.

A. Tarantola. Inverse problem theory and methods for model parameter estimation. *SIAM*, 2015a.

A. Tarantola. Towards adjoint-based inversion for rheological parameters in nonlinear viscous mantle flow. *Physics of the Earth and Planetary Interiors*, 234:23–34, 2015b.

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.

A. L. Teckentrup. Convergence of Gaussian process regression with estimated hyper-parameters and applications in Bayesian inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 8(4):1310–1337, 2020.

A. Tharwat, T. Gaber, A. Ibrahim, and A. E. Hassanien. Linear discriminant analysis: A detailed tutorial. *AI Communications*, 30(2):169–190, 2017.

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.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

M. K. Tippett, J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker. Ensemble square root filters. *Monthly Weather Review*, 131(7):1485–1490, 2003.

S. Tokdar, S. Kass, and R. Kass. Importance sampling: a review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(1):54–60, 2010.

X. T. Tong. Performance analysis of local ensemble Kalman filter. *Journal of Nonlinear Science*, 28(4):1397–1442, 2018.

X. T. Tong and M. Morzfeld. Localized ensemble Kalman inversion. *Inverse Problems*, 39 (6):064002, 2023.

X. T. Tong, A. J. Majda, and D. Kelly. Nonlinear stability of the ensemble Kalman filter with adaptive covariance inflation. *Nonlinearity*, 29(2):54–60, 2015.

X. T. Tong, A. J Majda, and D. Kelly. Nonlinear stability and ergodicity of ensemble based Kalman filters. *Nonlinearity*, 29(2):657, 2016.

F. Tonolini, J. Radford, A. Turpin, D. Faccio, and R. Murray-Smith. Variational inference for computational imaging inverse problems. *Journal of Machine Learning Research*, 21 (179):1–46, 2020.

A. Trevisan and F. Uboldi. Assimilation of standard and targeted observations within the unstable subspace of the observation–analysis–forecast cycle system. *Journal of the Atmospheric Sciences*, 61(1):103–113, 2004.

J. A. Tropp. *An Introduction to Matrix Concentration Inequalities*, volume 8. Now Publishers, Inc., 2015.

P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494, 2001.

A. B. Tsybakov. *Introduction to Nonparametric Estimation.* Springer Series in Statistics. Springer New York, 2008. ISBN 9780387790527. URL `https://books.google.com/books?id=mwB8rUBsbqoC`.

R. Tuo and W. Wang. Kriging prediction with isotropic Matérn correlations: Robustness and experimental designs. *Journal of Machine Learning Research*, 21(1):7604–7641, 2020.

S. Ungarala. On the iterated forms of Kalman filters using statistical linearization. *Journal of Process Control*, 22(5):935–943, 2012.

S. A. van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge University Press, 2000.

R. Van Handel. Probability in high dimension. Technical report, Princeton University, 2014.

R. Van Handel. On the spectral norm of Gaussian random matrices. *Transactions of the American Mathematical Society*, 369(11):8161–8178, 2017.

P. J. Van Leeuwen. A consistent interpretation of the stochastic version of the Ensemble Kalman Filter. *Quarterly Journal of the Royal Meteorological Society*, 146(731):2815–2825, 2020.

R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge University Press, 2018.

U. Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.

M. J. Wainwright and M. I. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Now Publishers Inc, 2008.

S. Wang and Y. Marzouk. On minimax density estimation via measure transport. *arXiv preprint arXiv:2207.10231*, 2022.

X. Wang, C. H. Bishop, and S. J. Julier. Which is better, an ensemble of positive–negative pairs or a centered spherical simplex ensemble? *Monthly Weather Review*, 132(7):1590–1605, 2004.

L. Wasserman. *All of Nonparametric Statistics*. Springer Science & Business Media, 2006.

A. Wehenkel and G. Louppe. Unconstrained monotonic neural networks. *Advances in neural information processing systems*, 32, 2019.

H. Wendland. *Scattered Data Approximation*, volume 17. Cambridge University Press, 2004. doi:http://dx.doi.org/10.1017/CBO9780511617539.

A. Wiens, D. Nychka, and W. Kleibe. Modeling spatial data using local likelihood estimation and a matérn to sar translation. *arXiv preprint arXiv:2002.01124*, 2020.

C. K. I. Williams and C. E. Rasmussen. *Gaussian Processes for Machine Learning*, volume 2. MIT press Cambridge, MA, 2006.

J. Wu, J. X. Wang, and S. C. Shadden. Improving the convergence of the iterative ensemble Kalman filter by resampling. *arXiv preprint arXiv:1910.04247*, 2019.

J. Wu, L. Wen, and J. Li. Resampled ensemble Kalman inversion for Bayesian parameter estimation with sequential data. *Discrete & Continuous Dynamical Systems-Series S*, 15 (4), 2022.

W. B. Wu and M. Pourahmadi. Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika*, 90(4):831–844, 2003. doi:http://dx.doi.org/10.1093/biomet/90.4.831.

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.

F. Yao, H. G. Müller, and J. L. Wang. Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590, 2005a. doi:http://dx.doi.org/10.1198/016214504000001745.

F. Yao, H. G. Müller, and J. L. Wang. Functional linear regression analysis for longitudinal data. *The Annals of Statistics*, 33(6):2873–2903, 2005b. doi:http://dx.doi.org/10.1214/009053605000000660.

L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Advances in neural information processing systems*, pages 1601–1608, 2005.

C. Zhang, J. Bütepage, H. Kjellström, and S. Mandt. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):2008–2026, 2018.

X. Zhang and J. Wang. From sparse to dense functional data and beyond. *The Annals of Statistics*, pages 2281–2321, 2016.

Y. Zhang and D. S. Oliver. Improving the ensemble estimate of the Kalman gain by bootstrap sampling. *Mathematical Geosciences*, 42:327–345, 2010.

S. Zhe, Syed A. Z. Naqvi, Y. Yang, and Y. Qi. Joint network and node selection for pathway-based genomic data analysis. *Bioinformatics*, 29(16):1987–1996, 2013.

D. X. Zhou. The covering number in learning theory. *Journal of Complexity*, 18(3):739–767, 2002.

X. Zhu, J. Lafferty, and Z. Ghahramani. Semi-supervised learning: from Gaussian fields to Gaussian processes. In *School of CS, CMU*. Citeseer, 2003.

H. Zou, T. Hastie, and R. Tibshirani. On the "degrees of freedom" of the lasso. *The Annals of Statistics*, 35(5):2173–2192, 2007.

E. Zuazua. Propagation, observation, and control of waves approximated by finite difference methods. *SIAM Review*, 47(2):197–243, 2005.