THE UNIVERSITY OF CHICAGO

HYPOXIA IMAGING WITH POSITRON EMISSION TOMOGRAPHY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON MEDICAL PHYSICS

BY
LUCAS BERENS

CHICAGO, ILLINOIS
DECEMBER 2026

© 2026 by Lucas Berens

ORCID iD: https://orcid.org/0000-0002-8958-3446

To my parents, my brother, and Quinn

"It's a magical world, Hobbes, ol' buddy... let's go exploring!"

— Bill Watterson

TABLE OF CONTENTS

LIST OF FIGURES viii

LIST OF TABLES xiv

LIST OF ACRONYMS AND INITIALISMS xv

ACKNOWLEDGMENTS xviii

ABSTRACT xxi

1 INTRODUCTION 1
1.1 An overview of PET 3
1.2 FMISO-PET 4
1.3 Positronium lifetime imaging 7
1.4 Roadmap 10

I AN ANALYSIS OF THEORETICAL AND COMPUTATIONAL RESULTS USING A MOMENT-BASED METHOD FOR POSITRON EMISSION TOMOGRAPHY-BASED HYPOXIA IMAGING

2 A MOMENT-BASED METHOD FOR ORTHOPOSITRONIUM LIFETIME ESTIMATION 13
2.1 The lifetime model 13
2.2 Moments and exponential weighting 15
2.3 The moment-ratio estimator for a uniform sample 15
2.4 Extension to positronium lifetime tomography 16
2.5 Estimation of the moment data from list-mode TOF-PET data 19
2.5.1 Per-event lifetime 19
2.5.2 Random density and weighted moments 23

3 POSITRON ANNIHILATION LIFETIME SPECTROSCOPY 25
3.1 Introduction 25
3.2 Materials and Methods 25
3.2.1 Simulated data 27
3.2.2 Implementation and numerical studies 28
3.3 Results 29
3.3.1 Lifetime estimate versus order of moment n 29
3.3.2 Lifetime estimate versus number of events N 31
3.3.3 Lifetime estimate versus cutoff and background 33
3.3.4 Monotonicity 35
3.4 Discussion 36
3.4.1 Effect of the number of events 36

3.4.2 Short o-Ps lifetimes *in vivo* . . . 38
3.4.3 Drawbacks and future work . . . 38
3.5 Conclusion . . . 39

4 POSITRONIUM LIFETIME TOMOGRAPHY . . . 41
4.1 Introduction . . . 41
4.2 Simulation studies . . . 43
4.3 Results . . . 46
4.3.1 Minimum $\overline{\text{NMSE}}$ . . . 46
4.3.2 $\overline{\text{NMSE}}$ and MLEM iteration . . . 47
4.3.3 CRC . . . 48
4.3.4 Mean and single-realization images . . . 51
4.3.5 Non-o-Ps component perturbations . . . 53
4.4 Discussion . . . 54
4.4.1 Influence of moment and CTR . . . 54
4.4.2 Bias and lookup table correction . . . 57
4.5 Conclusions . . . 61

II AN ANALYSIS OF EXPERIMENTAL RESULTS FOR POSITRON EMISSION TOMOGRAPHY-BASED HYPOXIA IMAGING

5 ADDITION OF CONFORMAL RADIATION THERAPY TO A PRECLINICAL MULTI-MODAL IMAGING PROTOCOL . . . 63
5.1 Introduction . . . 63
5.2 Radioisotope production . . . 64
5.3 Animal preparation and minimization of distress . . . 65
5.4 Preclinical magnetic resonance imaging (MRI) . . . 65
5.5 Tumor growth . . . 67
5.6 Injections . . . 67
5.7 Preclinical EPR-PET combined scanner . . . 68
5.8 Preclinical CT and CRT . . . 70
5.9 Conclusion of experimental protocol . . . 70

6 PRELIMINARY SURVIVAL STUDY OF DCE-MRI-BASED CRT CORRECTION IN A PRECLINICAL FIBROSARCOMA MODEL . . . 72
6.1 Introduction . . . 72
6.2 Brief reintroduction of dynamic contrast-enhanced (DCE) . . . 73
6.3 DCE correction of ^{18}F-fluoromisonidazole (FMISO) images . . . 73
6.4 Results . . . 74
6.4.1 Cohort and baseline characteristics . . . 74
6.4.2 Survival by protocol . . . 75
6.4.3 Hypoxia metric structure . . . 76
6.5 Conclusion . . . 77

7 OPTIMAL THRESHOLD ANALYSIS OF FMISO IMAGES IN PRECLINICAL FIBROSARCOMA MODELS . . . 80
7.1 Introduction . . . 80
7.2 Methods . . . 80
7.2.1 Imaging data and region of interest . . . 80
7.2.2 Definition of standardized uptake value (SUV) and tumor-to-muscle ratio (TMR) . . . 81
7.2.3 Thresholding strategies . . . 81
7.2.4 Agreement metrics and the composite score . . . 82
7.2.5 Threshold optimization . . . 83
7.2.6 Voxel-level discrimination . . . 84
7.2.7 Statistical analysis . . . 86
7.3 Results . . . 87
7.3.1 Cohort-level hypoxia statistics . . . 87
7.3.2 Voxel-level discrimination . . . 88
7.3.3 Optimal thresholds and strategy comparison . . . 88
7.3.4 Internal validation . . . 90
7.4 Discussion . . . 92
7.5 Conclusions . . . 94

8 THE FIRST POSITRONIUM LIFETIME TOMOGRAPHY DONE USING ^{44g}Sc WITH GROUND TRUTH EPR $\mathbf{pO_2}$ . . . 96
8.1 Introduction . . . 96
8.2 Methods . . . 97
8.3 Results . . . 98
8.4 Discussion . . . 100
8.5 Conclusion . . . 101

9 CONCLUSIONS AND FUTURE DIRECTIONS . . . 102

APPENDIX A SUPPORTING PROOFS AND DETAILS . . . 104
A.1 Moment properties . . . 104
A.2 The EMG, its tail, and the lower integration limit . . . 105
A.3 Moments of the EMG . . . 107
A.4 The moment-ratio relation . . . 112
A.5 List-mode evaluation of $\tilde{Y}_i^{(n)}(s)$ . . . 118
A.6 Linearity of $\tilde{Y}_i^{(n)}(s)$ . . . 119

REFERENCES . . . 120

LIST OF FIGURES

1.1 Chemical structure of FMISO (1-(2-Nitro-imidazolyl)-3-[18F]fluoro-2-propanol). 5

2.1 Geometry of a triple coincidence event. The annihilation point $\vec{r}_e$ lies on the line of response between the two annihilation photon detection positions $\vec{r}_1$ and $\vec{r}_2$. . 20

3.1 A simulated histogram of t_d having 2×10^5 events and using a 60 ps bin size. It was generated by using the following parameters: $\lambda_1 c_1$, $\lambda_2 c_2$, $\lambda_3 c_3 = 1.548$, 0.8, 0.12, τ_1, τ_2, $\tau_3 = 0.388\,\text{ns}$, 0.125 ns, 2.5 ns, $\sigma = 0.169\,\text{ns}$, $b = 10$, and the Gaussian rise position was offset by 9.97 ns so the peak was near-centered after the convolution. $\lambda_1, \lambda_2, \lambda_3 = 2.58\,\text{ns}^{-1}$, $8\,\text{ns}^{-1}$, $0.4\,\text{ns}^{-1}$. The individual contributions from the direct-annihilation (DA), parapositronium (p-Ps), and orthopositronium (o-Ps) components are shown in green, yellow, and magenta, respectively. 28

3.2 Estimated o-Ps lifetime as a function of the order of moment n. The shaded areas for each curve represent the ±1 standard deviation on each data point. The true τ_3 values shown in the legend are those used for producing simulation data. . . . 30

3.3 Similar to Fig. 3.2, this shows two practical cases for o-Ps lifetimes which have been measured recently *in vivo* 32

3.4 Estimated o-Ps lifetimes as a function of the number of events N in the histogram. The shaded areas again indicate ±1 standard deviation about the mean. 33

3.5 A simulated histogram for $\tau_3 = 1.5\,\text{ns}$ and $N = 1,500$, including the noise-free contributions of DA, p-Ps, and o-Ps. 34

3.6 Estimated o-Ps lifetime as a function of the upper truncation threshold T_2. Note that here $\hat{\tau}_3/\tau_3$ is plotted. The horizontal line indicates the perfect estimate given by $\hat{\tau}_3/\tau_3 = 1$. 35

3.7 The ratio of the standard deviation to the mean of the curves from Fig. 3.6. The ratios are seen to decrease as the upper truncation decreases. 36

3.8 Estimated o-Ps lifetime versus the true value for three background levels b with $n = 11$, $N = 2\times10^5$ and no upper truncation. The region close to $\tau_3 = 0.5\,\text{ns}$ may not be monotonic for some cases due to the curvature. 37

3.9 Several estimates of the $\tau_3 = 1\,\text{ns}$ case from Figure 3.4, however, with δ varied to demonstrate how the standard deviation variability can be handled for small τ_3 values. Note that increasing the value of δ increases the bias, but decreases the standard deviation for any given number of counts. 39

4.1 Ground truth maps used in this study. A) o-Ps rate constant map with a linear overlay (red) indicating the rate-constant profile taken for each simulated phantom. B) Ground truth activity distribution for the 1×10^8 total events case. 44

4.2 Mean minimum $\overline{\text{NMSE}}$ observed for each even moment for coincidence timing resolution (CTR)s of 200 ps, 600 ps, and 3000 ps, for A) 1×10^5 counts, B) 1×10^6 counts, C) 1×10^7 counts, and D) 1×10^8 counts, across all maximum-likelihood expectation-maximization (MLEM) iterations. 47

4.3 Specific cases of the data from Fig. 4.2. The plots are those where the minimum $\overline{\text{NMSE}}$ points for A) 1×10^5 counts, B) 1×10^6 counts, C) 1×10^7 counts, and D) 1×10^8 counts were observed. The shaded areas represent the ± 1 standard deviations. 48
4.4 A) The contrast recovery coefficient (CRC) for 1×10^8 counts versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A) 49
4.5 A) The CRC for 1×10^7 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A) 49
4.6 A) The CRC for 1×10^6 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A) 50
4.7 A) The CRC for 1×10^5 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A) 50
4.8 The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 200 ps, $n = 8$, 1×10^8 total counts, and taken from the 7th MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean. 52
4.9 A) A single realization of an image that was used to produce the data in Fig. 4.8. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 52
4.10 The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^7 total counts, and taken from the 4th MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean. 53
4.11 A) A single realization of an image that was used to produce the data in Fig. 4.10. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 53
4.12 The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^6 total counts, and taken from the 1st MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean. 54
4.13 A) A single realization of an image that was used to produce the data in Fig. 4.12. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 54

4.14 The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 3000 ps, $n = 12$, 1×10^5 total counts, and taken from the 1st MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean. 55
4.15 A) A single realization of an image that was used to produce the data in Fig. 4.14. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 55
4.16 The same reconstruction data as Fig. 4.11, but with the p-Ps and DA lifetimes allowed to be perturbed randomly (same perturbation for all voxels) by $\pm 10\%$. This data produced using CTR = 600 ps, $n = 8$, 1×10^7 counts, and taken from the 4th MLEM iteration. A) A single realization of the mean data. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 56
4.17 The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^6 total counts, and taken from the 4th MLEM iteration, as opposed to the 1st like in Fig. 4.12. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean. 58
4.18 A) A single realization of an image that was used to produce the data in Fig. 4.17. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm. 59
4.19 Ground-truth, uncorrected, and look-up table (LUT)-corrected o-Ps rate-constant maps for the structured phantom at the 1st MLEM iteration. A) Ground-truth map, B) uncorrected reconstructed map, and C) LUT-corrected reconstructed map. 60
4.20 A) The LUT, shown via the true versus estimated o-Ps rate constant values inside the region-of-interest (ROI). The shaded region denotes the ± 1 standard deviation about the mean. Although 20 nodes are plotted (one for each rate constant), 19 are used as two were degenerate. B) The resulting line profile from applying the LUT to the case from Fig. 4.12. 60

5.1 Timeline of a single imaging and treatment protocol for a general experiment. Times are given relative to FMISO injection at 0 h. Radioisotope production begins at −4 h and MRI is acquired at −1 h for tumor sizing. Following a 2 h tracer uptake period, positron emission tomography (PET) and electron paramagnetic resonance (EPR) imaging are performed, followed by a second MRI acquisition for tumor contouring and DCE. Next is treatment-planning computed tomography (CT) at 6 h and irradiation at 7 h. The lower axis expands the 6 h to 7 h planning-to-treatment window. After CT, the tumor is contoured (6.25 h), all modalities are registered to the common EPR-based fiducial space (6.5 h), the boost plan is generated (6.75 h), and treatment is delivered (7 h). The animal is maintained under anesthesia throughout. Times shown are approximate. 64
5.2 The combined EPR-PET small-animal scanner. A preclinical PET insert is mounted coaxially within the bore of the EPR imaging magnet, so that FMISO PET and oxygen EPR images can be acquired simultaneously and without repositioning the animal. Left: the full system on its mobile rack. Right: detail of the imaging bore showing the EPR magnet, the PET seated inside it, the mouse table, and the isoflurane anesthesia line. 69
5.3 Representative co-registered images of a tumor-bearing leg. All panels show the same axial slice, with the tumor contour (magenta) drawn on the T2-weigted (T2W) MRI and propagated to every modality through the fiducial-based registration. Top row: T2W MRI and CT provide anatomy and landmarks, with fiducial markers, bone, bed, and immobilizing cast labeled. The EPR map gives oxygen partial pressure in mmHg, with the hypoxic tumor core appearing as a dark-blue low-pO_2 region. Bottom row: FMISO PET shown as tumor-to-muscle ratio (TMR), and the DCE-derived K^{trans} (min^{-1}) and v_e (fractional) maps. Elevated FMISO uptake spatially corresponds to the EPR-defined hypoxic region in this example. 71

6.1 Kaplan-Meier recurrence free survival for the corrected and uncorrected arms over the 90 day window. Shaded bands are 95% confidence intervals. The curves overlap throughout, and a log-rank test found no significant difference ($P = 0.76$). 76
6.2 locally estimated scatterplot smoothing (LOESS)-smoothed scatter of hypoxic surface area against HF10, with the red curve as the LOESS fit. Hypoxic surface area rises roughly monotonically with HF10, a strong positive relationship. . . . 77
6.3 LOESS-smoothed scatter of the number of hypoxic regions against HF10, with the red curve as the LOESS fit. The number of hypoxic regions rises then falls across the HF10 range, a weak and nonmonotonic relationship. 78
6.4 LOESS-smoothed scatter of tumor growth rate against HF10, with the red curve as the LOESS fit. Tumor growth rate shows no trend with HF10. 78

7.1 Histogram of per-tumor hypoxic fraction, the fraction of ROI voxels at or below 10 mmHg, across all tumors with a valid ROI. The dashed line is the 5% inclusion floor and the shaded region marks the tumors that were excluded. 87

7.2 Distribution of voxel-wise partial pressure of oxygen (pO_2) pooled across all included tumors within the ROI. The dashed line marks the 10 mmHg hypoxia threshold and the shaded region the hypoxic range. 88
7.3 Voxel-level receiver operating characteristic curve (ROC) curves for TMR (left) and SUV (right) as scores to predict for pO_2-defined hypoxia. Grey curves represent ROC curves for individual tumors, the red curve is the mean across tumors, and the shaded band denotes the ± standard error (SE). The TMR and SUV curves are expected to be identical. 89
7.4 Voxel-wise relationship between pO_2 and PET signal, pooled across the included tumors, with TMR on the left and SUV on the right. The vertical dashed line is the 10 mmHg hypoxia threshold and the horizontal dashed line is the optimized θ using the absolute rule. 89
7.5 Mean overall hypoxic similarity (OHS) as a function of the decision threshold for each thresholding method, with the shaded band showing ± SE across tumors. The top row is the absolute and per-image-factor TMR, and the bottom row the absolute and per-image-factor SUV. The dashed line is the optimal threshold and, on the absolute-TMR panel, the dotted line is the TMR ≥ 1.4 baseline. 91
7.6 Mean Dice similarity coefficient (DSC), $\hat{d}_{H,95}$, and $\widehat{BA}$, overlaying TMR and SUV, with bands showing ± SE. The top row shows the per-image factor rules and the bottom row shows the absolute cutoffs. The figures are arranged in columns for the DSC, $\hat{d}_{H,95}$, and $\widehat{BA}$, respectively. The dashed vertical lines are the locations of the respective optimal thresholds, and the dotted line in the bottom-left panel is the baseline threshold scheme. The vertical lines are the OHS maximum locations (not the individual metric maximum locations). 92
7.7 OHS for each thresholding strategy across the included tumors after optimization. Boxes show the median and interquartile range (IQR) with whiskers to the full range if there are no outliers, or to $1.5 \times$ IQR otherwise. 93
7.8 Distributions of the enabled component metrics by thresholding strategy, from left to right the DSC, the $\hat{d}_{H,95}$ in voxels, and $\widehat{BA}$. Higher is better for all metrics. 93
7.9 leave-one-out (LOO) cross-validation. On the left the cross-validated OHS for each strategy, and on the right the in-sample versus held-out mean OHS for the four tunable strategies with the optimism gap annotated. 94

8.1 Experimental setup for the two positronium lifetime imaging (PLI) acquisitions. The gas phantoms were glass ampoules of deionized water through which air, argon, or carbogen was bubbled continuously during imaging, each loaded with approximately 170 µCi of ^{44g}Sc-prostate-specific membrane antigen (PSMA)-617. The mouse, bearing an lymph node carcinoma of the prostate - androgen receptor (LNCaP-AR) prostate tumor, received approximately 335 µCi of the same radiotracer. Both were imaged in the combined EPR-PET scanner following the protocol of Chapter 5. 99

8.2 Reconstructed o-Ps lifetime map of the gas phantoms for a selected axial slice. The two left panels show the PET moment image used to define the phantom regions. The two right panels show the reconstructed o-Ps lifetime $\hat{\tau}_3$ map and the same map restricted to those regions. The three segmented regions had mean lifetimes of 2.31 ± 0.18 ns for the oxygen phantom, 2.31 ± 0.18 ns for the air phantom, and 2.59 ± 0.18 ns for the argon phantom, where the stated uncertainty is the voxel-level standard deviation within each region. 99
8.3 Reconstructed o-Ps lifetime map of the tumor-bearing mouse for a selected axial slice. The two left panels show the T2W-MRI used to segment the tumor and muscle, over the full field of view and then restricted to those regions. The two right panels show the reconstructed o-Ps lifetime $\hat{\tau}_3$ map and the same map restricted to the segmented regions. The tumor and muscle had mean lifetimes of 2.12 ± 0.39 ns and 1.97 ± 0.40 ns, where the stated uncertainty is the voxel-level standard deviation within each region. 100

LIST OF TABLES

3.1 Bias-minimizing estimates of Fig. 3.2, shown with the moment n which minimized the bias. We again emphasize that this moment-based method requires n to be sufficiently large. 31
3.2 Selected estimates from Fig. 3.4. 32

6.1 Metrics used in this study for correlations to hypoxia. 75
6.2 Baseline hypoxia metrics by protocol. Continuous metrics are mean ± SD and are compared with Welch's two-sample t-test. The P column compares the corrected and uncorrected arms. Covariate data were unavailable for one uncorrected animal and therefore the first six rows are computed on 12 corrected and 13 uncorrected tumors while the recurrence row uses all 26 animals. 75
6.3 Pearson correlation coefficients among the six hypoxia metrics. 77

7.1 Performance metrics for each thresholding method predicting EPR-defined hypoxia across the 20 included tumors, reported as median and IQR. All four columns are on a common zero-to-one scale for which higher values indicate better agreement with ground truth. For the SUV$\geq$ 0.19 method the $\hat{d}_{H,95}$ entry is computed from 19 tumors, since in one tumor no voxel exceeded the cutoff and the predicted mask was empty, leaving $\hat{\mathrm{d}}_{\mathrm{H,95}}$ undefined. A fifth method, the optimized absolute TMR cutoff, converged to 1.40 and is not listed separately because it is identical to the fixed TMR$\geq$ 1.4 baseline shown in the first row. . . 90
7.2 Optimized thresholds and held-out mean OHS under LOO, repeated five-fold cross validation, 0.632-bootstrap validation, and the optimism. Optimal θ maximized the in-sample mean OHS. For reference, the fixed TMR$\geq$ 1.4 baseline had a mean OHS of 0.484. 94

LIST OF ACRONYMS AND INITIALISMS

$d_{H,95}$	95th percentile Hausdorff distance
AUC	area under the curve
BA	balanced accuracy
CRC	contrast recovery coefficient
CRT	conformal radiation therapy
CT	computed tomography
CTR	coincidence timing resolution
DA	direct-annihilation
DCE	dynamic contrast-enhanced
DNA	deoxyribonucleic acid
DSC	Dice similarity coefficient
EBRT	external-beam radiation therapy
EMG	exponentially modified Gaussian
EPR	electron paramagnetic resonance
EPRI	electron paramagnetic resonance imaging
EPROI	electron paramagnetic resonance oxygen imaging
FAZA	^{18}F-fluoroazomycin arabinoside
FETNIM	^{18}F-fluoroerythronitromidazole
FMISO	^{18}F-fluoromisonidazole
FPF	false positive fraction
FSa	fibrosarcoma
FWHM	full width at half maximum
HIF	hypoxia-inducible factor

IMRT intensity-modulated radiation therapy

IQR interquartile range

iSAIRR Integrated Small Animal Imaging Research Resource

LET linear energy transfer

LNCaP-AR lymph node carcinoma of the prostate - androgen receptor

LOESS locally estimated scatterplot smoothing

LOO leave-one-out

LOR line of response

LUT look-up table

MCa4 mammary carcinoma

MLEM maximum-likelihood expectation-maximization

MRI magnetic resonance imaging

NMSE normalized mean squared error

o-Ps orthopositronium

OHS overall hypoxic similarity

OOB out-of-bag

p-Ps parapositronium

PALS positron annihilation lifetime spectroscopy

PAVA pool-adjacent-violators algorithm

PDF probability density function

PET positron emission tomography

PLI positronium lifetime imaging

PLT positronium lifetime tomography

pO_2 partial pressure of oxygen

Ps	positronium
PSMA	prostate-specific membrane antigen
RMS	root mean squared
ROC	receiver operating characteristic curve
ROI	region-of-interest
ROS	reactive oxygen species
SCCVII	squamous cell carcinoma VII
SE	standard error
SUV	standardized uptake value
T1W	T1-weigted
T2W	T2-weigted
TCD15	tumor control dose 15%
TCD95	tumor control dose 95%
TMR	tumor-to-muscle ratio
TOF	time-of-flight
TPF	true positive fraction

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to my mentors, friends, colleagues, and family members who have contributed, both directly and indirectly, to my dissertation and scientific training over the past five years. The list of individuals who have supported me throughout this journey is possibly too long for me and my fallible memory to include in full, and I apologize in advance for any inadvertent omissions.

To my advisors, Howard Halpern and Chien-Min Kao, thank you for your guidance, support, and especially your patience. I am immensely lucky to have such mentors who knew when to challenge me to grow and when to offer encouragement. Your confidence in my abilities, thoughtful advice, and unwavering support have shaped not only this work but also my development as a researcher and as a person. I am deeply grateful for everything you have taught me and for the opportunity to learn under your mentorship.

To my Committee Chair, Chin-Tu Chen, thank you as well for your wisdom and support. I am continuously surprised at how much time you were able to dedicate to my success, and I am eternally grateful for your assistance throughout my journey. To Maryellen Giger, thank you for being the only non-nuclear medicine scientist on my committee. I was ecstatic when you agreed to join, and I have never once regretted it. Your suggestions, insights, and feedback have been among the most useful I have ever received.

As will hopefully become apparent during a read-through of this dissertation, many, many people contributed to this work. For example, the experimental protocol described in this dissertation involved more than a dozen people working in preparation for, during, or after the protocol. In short, these individuals are truly the backbone of my experimental data analysis. They include Heejong Kim, Mihai Giurcanu, Richard Freifelder, Lucas Gonzalez, Eugene Barth, John Lukens, Jason Meier, Boris Epel, Subramanian Sundramoorthy, Mohammed Bhuiyan, Kaustab Ghosh, Swapna Chitneni, Xiaobing Fan, Erica Markiewicz, Brian B. Roman, Andrew Miller, and Satish Chitneni. A sincere thank you to you all, and a

special thank you to Gene, John, Erica, Brian, Andrew, Boris, Heejong, Fan, and Mihai for your unrelenting helpfulness and kindness during many early mornings and late evenings. I also would like to acknowledge the mice who unwillingly gave their lives for this work.

To current and past graduate students in medical physics, I sincerely thank you for your friendship and for having my back during graduate school. I especially want to acknowledge my cohort, Gia Jadick, Chisondi Warioba, and Amy Renne. You all are some of the best people on the planet. Thank you as well to Inna Gertsenshteyn, without whom my work would simply not have been possible, and Jason Meier for the unending guidance you continue to give me. Thank you also to Hadley DeBrosse, Geneva Anderberg, Linnea Kremer, Mira Liu, Mena Shenouda, Renu Rajamagesh, Chris Valdes, Ethan Stolen, Tabitha Welch, Sagada Penano, Kyle Lleras, and JP Phillips for keeping the program's culture of friendship and outreach alive.

Aside from my colleagues and dissertation advisors, some of the medical physics faculty and staff have played no small part in my graduate training. These include, but are certainly not limited to, Sam Armato, to whom I owe a great deal for support in my career decisions, Lili González, whose patience is something to strive for, Elena Rizzo and Jasmine Fields, for making this entire task manageable, Patrick La Riviere for mentoring me during my first-year rotation, Emily Marshall, for being the absolute greatest advocate for outreach we could have asked for, and Yulei Jiang, for his extraordinary generosity with his time and effort in supporting my grant submissions and outreach events. Further thanks are due to Ingrid Reiser, Zheng Feng Lu, Zhen Tian, Heather Whitney, Baihui Yu, Kevin Little, Steffen Sammet, Cameron Kofler, and Hania Al-Hallaq for your teaching, your conversations, and your well wishes on more occasions than I can remember. I would be remiss not to mention my undergraduate advisor Aaron McGowan, without whom I would not have reached graduate school in the first place.

I am incredibly fortunate to have been part of a spectacular set of extracurricular groups

during my time in graduate school. The people I have met because of this have been absolutely vital to my success. I want to thank Jim Poueymirou, for being possibly the greatest mentor I have experienced outside of graduate school. You truly live up to the scouting ideals and the world would be a lesser place without your presence in it. To Megan Mekinda, Rosie Huggins, and Michelle Domecki: in my opinion you are doing some of the most important work at the university. Thank you for your time and effort, and for allowing me to be a part of it. Margarite Matossian, Nada Aboelella, and Kevin Chang, I thank you for being such a fantastic group to work with, and I hope your careers blossom. To Erin Adams, Beth Wickware, Candice Lewis, Nick Foster, John Anderson, and Shawn Hannan, thank you for everything you have helped me with.

To my friends, I owe everything. In no particular order I want to acknowledge Alec, Casey, Dom, Heather, Michael, Vicky, Stevie, Morgan, Alan, Marek, Dominika, Ellen, Isaac, Artur, Rustin, Jessalin, Toby, Jamie, Allie, Malena, Noah, Catherine, Tasha, Grayson, Jonathan, Spencer, Helle, Jared, and especially Brandon and Isa, who gave me quite possibly the greatest gift of my life thus far. If you are reading this and wondering why you are not on this list, well, bring it up next time you see me and I promise to make it up to you.

Last, but certainly not least, thank you Mom, Dad, grandparents, Chaney, Quinn, Jacqueline, cousins, aunts, and uncles for your encouragement and love.

ABSTRACT

Tumor hypoxia, a state of low oxygen within a tumor, causes resistance to radiotherapy, chemotherapy, and immunotherapy and is associated with poor patient outcomes. Because hypoxic regions can be targeted with an external-beam radiation therapy (EBRT) dose boost, the ability to map oxygen concentration *in vivo* could improve treatment. This dissertation focuses on two complementary positron emission tomography (PET) approaches to accomplish this: 1) a moment-based reconstruction method for positronium lifetime imaging/tomography (PLI)/(PLT), and 2) a correction-based analysis of the hypoxia radiotracer ^{18}F-fluoromisonidazole (FMISO) using dynamic contrast-enhanced (DCE) magnetic resonance imaging (MRI), leveraging electron paramagnetic resonance (EPR) imaging to measure ground truth oxygenation.

The first part introduces a fast, analytic, moment-based estimator of the orthopositronium (o-Ps) lifetime in positron annihilation lifetime spectroscopy (PALS), a quantity sensitive to and potentially able to quantify the local oxygen concentration. Our estimator is calculated from the ratio of two successive, exponentially weighted, modified moments of the elapsed-time annihilation histogram. The estimator isolates the o-Ps component without fitting and extends to tomography as a pair of ordinary reconstructions followed by a voxel-wise ratio. It is shown to be stable and monotonic on one-dimensional PALS data and to yield robust two-dimensional lifetime images in simulation for PLI/PLT.

The second part focuses on experimental methods in a preclinical fibrosarcoma (FSa) model imaged with a combined PET-EPR scanner that provides ground-truth oxygen maps. A DCE MRI correction to FMISO PET is evaluated in a survival study, and the optimal FMISO threshold for defining hypoxia is re-examined for variable thresholding approaches.

Finally, elements from both parts were combined for an experiment which produced the first PLT images acquired with ^{44g}Sc alongside a ground-truth oxygen reference from EPR imaging. Together these results advance both the established and the emerging routes to

PET-based hypoxia imaging for oxygen-guided EBRT.

CHAPTER 1

INTRODUCTION

The complex system of organs and tissues that form the human body, operating via deeper cellular, biochemical, and biophysical mechanisms, is an evolutionary marvel. Its reliability and robustness are astonishing, yet, the delicate homeostasis of the body can be disrupted. In this dissertation, we focus on tumor hypoxia, a condition in which normal oxygen transport and delivery to tumorigenic tissues is disrupted [1]. Normal tissue vasculature can be characterized by well-organized blood vessels which, among other functions, deliver adequate oxygen and nutrients levels to the surrounding tissue [2,3]. The growth of tumors, however, is often accompanied by tumor angiogenesis, due to the high metabolic demand, which results in a chaotic and disorganized vasculature [2,4]. This tortuous vasculature can lead to regions of low oxygenation, or hypoxia, which is a hallmark of most solid tumors [5,6]. We define tumor hypoxia as the measured local partial pressure of oxygen (pO_2) in a tumor tissue that is below 10 mmHg [7]. Normal oxygen concentrations in the body vary based on tissue type (from 8 mmHg in skin dermis to 108 mmHg in the alveoli [8]), and a review of many values in various tissues and organs can be found in Ortiz-Prado *et al.* [8].

Counterintuitively, tumor cells in hypoxic environments often remain viable [9] due to the overexpression of hypoxia-inducible factors (HIFs) [10]. Specifically, the α subunit of HIF-1 is stabilized in hypoxic conditions, which leads to the transcription of genes that promote angiogenesis, metabolic reprogramming, and cell survival [10]. These adaptations allow tumor cells to survive hypoxic conditions and to develop resistance to many therapies [11]. Hypoxic tumors are associated with resistance to chemotherapy, due to the fact that many chemotherapeutic agents either directly require normoxic conditions to function as designed, or target cellular mechanisms which have increased resistance in hypoxic environments [12]. Radiotherapy resistance is also observed in hypoxic tumor environments [13]. Ionizing radiation therapy relies on the generation of reactive oxygen species (ROS) to in-

duce deoxyribonucleic acid (DNA) damage in tumor cells [13]. For low linear energy transfer (LET) radiation such as that from X- and γ-rays, most ionizations are caused from indirect action, where ROS are generated from the radiolysis of water and then diffuse to the DNA to cause damage [13–15][1]. This is in contrast to direct action, where the radiation directly ionizes atoms in the DNA molecule (or target of interest). In hypoxic conditions, the concentration of ROS from conventional photon radiation therapy is reduced, leading to an increase in radiation resistance [13]. For completeness, the immune function of the body is also suppressed in hypoxic tumor environments [17], which also may contribute to immunotherapeutic resistances [17, 18]. The presence of hypoxia in tumors therefore, has significant implications for the effectiveness of cancer therapies, and it is known that hypoxic tumors are associated with worsened patient outcomes due to the aforementioned factors [19].

It is then natural to consider cancer therapies that target hypoxic tumor environments. Because hypoxic tissue regions are located relatively far from blood vessels, they are often difficult to reach with pharmaceutical therapies delivered via the cardiovascular system [20]. Despite this challenge, there is great interest in developing hypoxia-targeting strategies [21]. A recent review of therapeutic strategies for targeting HIFs can be found in [22].

In parallel, it is important to consider non-pharmaceutical therapies as well. Intensity-modulated radiation therapy (IMRT) is a form of radiotherapy that allows for the delivery of highly targeted radiation doses to a tumor which may be used to target hypoxic regions [23]. Although hypoxic regions are more resistant to radiation therapy, they can be targeted with higher doses of radiation, which can be delivered using conformal radiation therapy (CRT) [24]. We have previously shown that the use of CRT to target mice with hypoxic

1. Strictly speaking, LET is defined for charged particles and does not apply directly to photons such as X- and γ-rays. In this context, "low-LET radiation" is used conventionally to refer to the secondary charged particles, primarily electrons, produced by photon interactions in tissue. For a more in-depth treatment of LET, the definition of it with regard to photons, and other metrics for radiation and biological effects, we note two specific useful publications of the International Commission on Radiological Protection [16], and the Internation Atomic Energy Agency [15]. In our context the low-LET radiation causes secondary electrons to be formed, mainly through Compton scattering, which then ionizes the DNA [16].

tumors can lead to improved tumor control and survival outcomes [24, 25]. We hypothesize that if a reliable and accurate method for mapping tumor hypoxia, or more generally oxygen concentration *in vivo*, can be developed, it will allow for IMRT to be used to boost dose to hypoxic tumor regions, which will improve patient outcomes. This dissertation work focuses on two methods using positron emission tomography (PET) to map oxygen concentrations. The first is a method using the oxygen-associated positronium (Ps) lifetime, and the second is a method using the radiotracer ^{18}F-fluoromisonidazole (FMISO).

Mapping tumor oxygenation *in vivo* has been a long-standing scientific goal [26], and many methods have been developed to address this. The historical reference standard is the polarographic Eppendorf electrode, which samples pO_2 directly but is invasive and restricted to accessible tumors [27–30]. Markers such as HIFs, carbonic anhydrase IX, pimonidazole, and pentafluorinated etanidazole report hypoxia on excised tissue by immunohistochemistry [6]. Non-invasive imaging methods span blood-oxygen-level-dependent and oxygen-enhanced magnetic resonance imaging (MRI), and nuclear methods using hypoxia-seeking PET tracers [6,26,31,32]. Among techniques that image oxygen directly, electron paramagnetic resonance (EPR) oxygen imaging provides quantitative, absolute pO_2 maps and serves as the ground-truth reference throughout this work [7, 24, 33, 34].

1.1 An overview of PET

PET is a functional imaging modality that maps the *in vivo* distribution of a positron-emitting radiotracer [35]. A small quantity of a radiopharmaceutical, labeled with a radioisotope that decays by positron (β^+) emission, is administered to the subject, where it accumulates according to the physiological or molecular process the tracer is designed to probe. Each emitted positron travels a short distance through tissue, on the order of millimeters (for example, ^{18}F has an average positron range of 0.36 mm while for ^{82}Rb it is an average of 3.85 mm [35, 36]), losing energy until it thermalizes and annihilates with an elec-

tron. The annihilation converts the rest mass of the electron-positron pair into two 511 keV gamma photons emitted almost exactly 180° apart. Their near-collinearity is useful and a consequence of conservation of energy and momentum [35]. The spatial resolution of PET is ultimately limited by the finite positron range and the small deviations from collinearity of the annihilation photons [35].

To perform PET, the subject is surrounded by a ring of scintillation detectors, and any two photons registered within a narrow coincidence timing window are attributed to a single annihilation. Each such pair defines a line of response (LOR), the line joining the two detection points, along which the annihilation is taken to have occurred. From a large number of LORs, the three-dimensional tracer distribution is recovered by analytic or, more commonly, iterative reconstruction such as maximum-likelihood expectation-maximization (MLEM) [35]. Time-of-flight (TOF) systems additionally measure the difference in arrival times of the two photons and use it to constrain the annihilation position along the LOR, improving the signal-to-noise ratio of the reconstructed image, with the precision of this localization set by the coincidence timing resolution (CTR) of the scanner. Because the reconstructed image gives intensity based on where the tracer accumulates, the tracer itself determines which biological process is visualized, which motivates the hypoxia tracers considered next.

1.2 FMISO-PET

A family of ^{18}F-based PET tracers have been investigated for targeting hypoxia. Specifically, FMISO, is a well-known radiotracer that has been studied as a probe for hypoxia [37]. Fig. 1.1 shows the structure of FMISO. The accumulation of FMISO in hypoxic tissues is thought to be due to reduction of its nitro group, which allows the resulting molecule to bind covalently to macromolecules [37]. Other mechanisms that trap the tracer inside hypoxic cells are likely also at play, which may include for example, glutathione-dependent conjugation [38, 39].

The reduction step is straightforward and illustrates how FMISO accumulation is oxygen

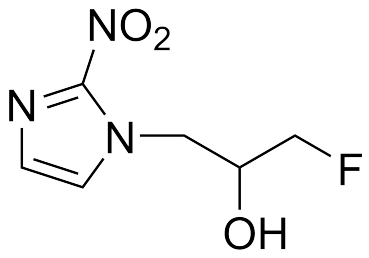

Figure 1.1: Chemical structure of FMISO (1-(2-Nitro-imidazolyl)-3-[18F]fluoro-2-propanol).

dependent. After passive diffusion into the cellular compartment, FMISO is reduced to an unstable radical,

$$\mathrm{FMISO} + \mathrm{e}^- \rightleftharpoons \mathrm{FMISO}^{\bullet -} \tag{1.1}$$

Under normoxic conditions, this radical is rapidly reoxidized by molecular oxygen, regenerating FMISO and producing superoxide, allowing the tracer to diffuse back into the vascular compartment,

$$\mathrm{FMISO}^{\bullet -} + \mathrm{O_2} \longrightarrow \mathrm{FMISO} + \mathrm{O_2}^{\bullet -} \tag{1.2}$$

Under hypoxic conditions, the absence of oxygen prevents reoxidation of the radical, allowing successive enzymatic reductions of the nitro group through nitroso and hydroxylamine intermediates to the amine [38, 40],

$$\mathrm{R{-}NO_2} + 6\,\mathrm{H^+} + 6\,\mathrm{e^-} \longrightarrow \mathrm{R{-}NH_2} + 2\,\mathrm{H_2O}, \tag{1.3}$$

which subsequently forms covalent adducts with intracellular macromolecules, resulting in irreversible trapping of FMISO within hypoxic tissue.

Certain biomarkers derived from FMISO and its second generation counterpart ^{18}F-fluoroazomycin arabinoside (FAZA) have been shown to correlate with local regional control [41]. To date, however, there has been insufficient evidence to support the use of external beam radiation boosts to FMISO-defined hypoxic regions in head and neck cancer [42]. For

example, a phase II clinical trial [43] using FMISO-PET to provide radiation boosts to hypoxic locally advanced head-and-neck cancer subvolumes failed to show a significance in local tumor control between boosted and non-boosted radiation treatment therapies. However, this trial was terminated early because of low enrollment. Additionally, in 2017, a phase II clinical trial by Vera *et al.* [44], showed that there was no significant benefit to dose escalation to hypoxic areas in non-small cell lung carcinoma based on FMISO-PET. While FAZA-PET and other PET-based hypoxia tracers have yet to undergo clinical trials at the time of this writing, they each face unique challenges that must be addressed before use in dose-escalation therapy. For example, FMISO-PET suffers from dependence on blood perfusion kinetics and lipophilicity [45]. FAZA-PET (which is hydrophilic and has better kinetics than FMISO-PET [46]) is possibly the most promising hypoxia tracer but cannot easily cross the blood–brain barrier [47]. Perspectives and outlooks in those and other tracers, such as the [^{18}F]EF family, [^{18}F]flortanidazole ([^{18}F]HX4), ^{18}F-fluoroerythronitromidazole (FETNIM), and non ^{18}F-based compounds can be found in several recent reviews including Gouel *et al.* [48], Perez *et al.* [49], and Huang *et al.* [50].

Despite these limitations, FMISO-PET remains the reference standard for hypoxia imaging and the most extensively validated of the available tracers [51], which makes it the natural basis for the correction developed in this work. Its dependence on perfusion and delivery kinetics is the confounder that the dynamic contrast-enhanced (DCE) MRI-based correction of Chapter 6 is designed to address. FMISO is used in this work because it is well-characterized and because the testing of a correction-based method may be developed more readily. Combined with its availability the absence of an alternative tracer that is both more capable and clinically established, these considerations motivate the continued use of FMISO-PET here.

1.3 Positronium lifetime imaging

Ps has been studied since the mid-20th century [52], and the sensitivity of the orthopositronium (o-Ps) lifetime to its surroundings has long been exploited in materials science [53], where positron annihilation lifetime spectroscopy (PALS) probes the free-volume voids and porosity of polymers and solids [54]. Its translation to medicine is recent and still in-development and investigation. Building on the sensitivity of the o-Ps lifetime to nanoscale structure and dissolved oxygen, Moskal *et al.* proposed Ps imaging as a new source of contrast, obtained simultaneously with conventional tracer-uptake PET [55, 56]. Unlike standard PET, which images only the location of tracer accumulation, positronium lifetime imaging (PLI) provides a physical property of the tissue itself.

In a conventional PET scan, positrons emitted by the radioisotope thermalize in tissue and annihilate with local electrons. Most annihilate directly, a process referred to below as direct-annihilation (DA), producing the pair of coincident 511 keV gamma rays from which the image is formed. In roughly 40 % of cases, however, the positron and an electron first bind into a metastable state called Ps [55]. Positronium exists in two ground states, parapositronium (p-Ps) with spin $S = 0$ and o-Ps with $S = 1$. In vacuum p-Ps self-annihilates into two photons within about 125 ps, whereas o-Ps decays into three photons only after about 142 ns [57]. The thousandfold difference in timing makes o-Ps useful to medical applications. Its relatively long lifetime gives it time to interact with the surrounding molecules and decay through either of two additional routes before self-annihilating. The first route is pick-off annihilation, in which the positron annihilates directly with an electron from a neighboring molecule at a rate set by certain tissue properties including, for example, the electron density. In water, pick-off contributes roughly $513\,\mu s^{-1}$ against an o-Ps self-annihilation rate of $7.04\,\mu s^{-1}$, collapsing the o-Ps mean lifetime from 142 ns to about 1.8 ns [58].

Neither of the shorter-lived components is affected to a comparable degree. Pick-off is spin independent and so adds the same rate to p-Ps, but against a self-annihilation rate of

7990.9 μs^{-1} that addition shortens the p-Ps mean lifetime by only about 7 ps [58]. Direct annihilation involves no bound state and is subject to neither pick-off nor spin exchange, with a mean time of 400 ps to 450 ps reported in water [58].

The second route is spin exchange with a paramagnetic species, which converts o-Ps into p-Ps and so ends the state within picoseconds. This route requires the presence of local unpaired spins. As unpaired electrons are found in abundance from O_2 molecules, the local concentration of O_2 may significantly affect the o-Ps lifetime to the degree that o-Ps lifetime in tissue may be an effective metric for *in vivo* O_2 concentration. In *in vitro* studies, a nearly linear dependence of o-Ps lifetime with the O_2 concentration in water has been reported by Shibuya *et al.* [59]. In this study, the shortest published o-Ps lifetime value using a phantom was observed to be 1.8239 ns [59]. This was measured in fully O_2 saturated water. In a more artificial environment, Stepanov et al. [60] bubbled oxygen gas through water and measured an o-Ps lifetime of 1.746 ns. With argon gas bubbled through water, however, a lifetime of 1.833 ns was measured. This is evidence that even extremely well-oxygenated *in vivo* environments will not have o-Ps lifetimes of less than 1.746 ns. In a recent review, Moskal and Stępień concluded that the mean biological o-Ps lifetime would be approximately 2 ns [61]. This is substantially shorter than its 142 ns lifetime in vacuum and depends on local tissue properties. In contrast, the p-Ps lifetime in tissue is not significantly different from its value in vacuum. Overall, the observed mean o-Ps lifetime and the mean positron lifetime (population average of the DA, p-Ps, and o-Ps lifetimes) can reflect differences in the tissue microenvironment and produce contrast for differentiating normal and tumorous tissues.

This sensitivity has been observed directly in tissue. *Ex vivo* PALS of human samples has found that both the o-Ps lifetime and its intensity differ between tissue pathologies [62, 63], and the o-Ps lifetime separates soft-tissue types such as adipose, muscle, and hepatic tissue [64]. The effect persists *in vivo*, where a shorter mean o-Ps lifetime was measured inside a

brain glioma than in adjacent healthy tissue [65]. Together these support the o-Ps lifetime as a tissue-differentiating and potentially oxygen-sensitive biomarker.

To measure the positron lifetime with TOF PET, Moskal *et al.* [66] proposed using an isotope that emits a prompt gamma (γ) photon at essentially the same time as a positron (β^+) decay, called a $\beta^+\gamma$ decay below, and modifying the scanner to also detect triple-coincidence events that include a prompt gamma and two annihilation photons. The time difference between the release of the prompt gamma and the occurrence of annihilation (the positron lifetime), can be estimated from the detection times of the prompt gamma and annihilation photons. ^{44g}Sc has been identified as the most-promising isotope for this task [56, 67]. Nearly all of its decays are $\beta^+\gamma$ decays.

Feasibility of PLI in clinical settings has also begun to be demonstrated. Moskal *et al.* reported the first *in vivo* imaging of o-Ps lifetime in a patient with a primary brain glioma [65]. They observed a reduced mean lifetime within the tumor compared to adjacent healthy brain tissue, measuring 1.77 ± 0.58 ns in glioma tissue versus 2.72 ± 0.72 ns in normal tissue. Furthermore, Steinberger *et al.* [68] and Mercolli *et al.* [69] demonstrated that the clinically available Biograph Vision Quadra PET scanner can reliably measure o-Ps lifetimes, both in reference radioisotopes embedded in inert materials [68] and in human subjects [69]. We hypothesize that not only may the o-Ps lifetime be used for soft-tissue contrast enhancement, but that it also may be used to create hypoxic maps within tumors and therefore allow for dose enhancement.

In previous studies, Moskal *et al.* and Steinberger *et al.* did not truly reconstruct the mean lifetime images. Instead, they assigned each $\beta^+\gamma$ decay to the point on the LOR connecting the detection positions of the annihilation photons according to the measured TOF. Using this assignment, they obtained a lifetime histogram at an image voxel or in a region-of-interest (ROI) and fit the histogram to derive the mean o-Ps lifetime [56, 65, 68]. The accuracy of this method is limited by the CTR of current PET systems, which is in the

range of 200 ps to 600 ps full width at half maximum (FWHM) and corresponds to a 3.0 cm to 9.0 cm FWHM uncertainty in positioning events. Consequently, the lifetime histogram obtained at a voxel or in an ROI contains events originating in a considerably larger volume than the voxel and ROI. The resulting mean o-Ps lifetime may be incorrect when the lifetime distribution in this volume is heterogeneous.

Accurate reconstruction of the mean lifetime images of o-Ps or the positron is a new research topic. Maximum-likelihood-based methods have been developed [70]. However, these methods are computationally intensive and can be sensitive to mismatches between the assumed model and the data. Huang *et al.* proposed an efficient statistical method that reconstructs several lifetime-weighted images and derives the o-Ps lifetime image from them [71]. This group also reported an even more efficient method for reconstructing the mean positron lifetime instead of the mean o-Ps lifetime [72]. This method was extended by Huang and Qi to estimate the o-Ps lifetime by using a series of equations from moment-based images [73]. In this dissertation, we present an efficient and robust moment-ratio estimator for reconstructing the o-Ps lifetime image.

1.4 Roadmap

The remainder of this dissertation is organized into two parts. Part I develops and characterizes the moment-based method related to PLI computationally. Chapter 2 derives the moment-ratio estimator for the o-Ps lifetime and its extension to tomography, Chapter 3 validates the estimator against one-dimensional PALS measurements, while Chapter 4 evaluates it for tomographic reconstruction. Part II discusses experimental work related to both FMISO and PLI. Chapter 5 details the preclinical multi-modal imaging protocol, and Chapter 6 presents a survival study of a DCE MRI-based FMISO correction using electron paramagnetic resonance imaging (EPRI) as the ground truth standard under that protocol. Chapter 7 determines the optimal FMISO threshold for defining hypoxia against EPR ground

truth. Finally, Chapter 8 reports the first PLI acquired with ^{44g}Sc alongside a ground truth oxygen reference, using many of the experimental protocol tools and expertises developed in Chapter 5.

Part I

An Analysis of Theoretical and Computational Results using a Moment-Based Method for Positron Emission Tomography-Based Hypoxia Imaging

CHAPTER 2

A MOMENT-BASED METHOD FOR ORTHOPOSITRONIUM LIFETIME ESTIMATION

Positronium lifetime imaging (PLI) exploits the fact that the mean lifetime of positronium is sensitive to the local tissue environment. Extracting that lifetime from measured data is the central estimation problem addressed in this work. This chapter develops the analytic framework used in the two studies that follow, with a specific focus on the orthopositronium (o-Ps) decay-rate constant/lifetime. We first construct an estimator of the o-Ps rate constant for a spatially uniform sample from the ratio of two successive exponentially weighted moments of the elapsed-time/lifetime histogram. We then show that the same construction extends, essentially unchanged, to tomographic imaging, where it reduces to two ordinary emission-tomography reconstructions followed by a voxel-wise ratio. Finally, we give the practical estimators needed to form the required quantities from list-mode time-of-flight (TOF) positron emission tomography (PET) data. Selected proofs and auxiliary results are collected in the Appendix. The two subsequent chapters evaluate this framework. The first is for the one-dimensional positron annihilation lifetime spectroscopy (PALS) measurements, and the second is for two-dimensional tomographic reconstruction. In theory this can be readily extended to full three-dimensional reconstruction. The method generalizes the analytic estimator we introduced previously for PALS measurements [74], and the tomographic extension follows the statistical-imaging strategy of Huang *et al.* [72].

2.1 The lifetime model

Let t_d denote the *lifetime*, i.e. the time difference between the emission of the prompt gamma and the subsequent annihilation. The distribution of t_d is comprised of three components: 1) the direct-annihilation (DA) group, in which the positron annihilates without forming

positronium, 2) the parapositronium (p-Ps) group, and 3) the o-Ps group. The lifetime distribution of each component is given by the probability density function (PDF) of an exponential distribution,

$$\mathrm{EXP}(t_d;\lambda_k) = \lambda_k\, e^{-\lambda_k t_d}\, u(t_d), \tag{2.1}$$

where $\lambda_k > 0$ is the rate constant of component k ($k = 1$ for DA, $k = 2$ for p-Ps, $k = 3$ for o-Ps) and $u(t_d)$ is the unit step function, equal to 1 for $t_d \geq 0$ and 0 otherwise. Because o-Ps is the longest-lived component, its rate constant is the smallest, and throughout we assume

$$\lambda_3 < \lambda_1,\, \lambda_2. \tag{2.2}$$

This λ_3 is the quantity to be estimated. In some sections, we estimate the lifetime $\tau_3 = 1/\lambda_3$ instead.

The finite temporal resolution of the detector blurs each exponential by a Gaussian. Modeling the time-measurement uncertainty as a zero-mean Gaussian of standard deviation σ,

$$\mathcal{N}(t_d;\sigma) = \frac{1}{\sqrt{2\pi}\,\sigma}\, \exp\!\left\{-\frac{t_d^2}{2\sigma^2}\right\}, \tag{2.3}$$

the observed distribution of each component is the exponentially modified Gaussian (EMG) [56],

$$\mathrm{EMG}(t_d;\lambda_k,\sigma) = \mathrm{EXP}(t_d;\lambda_k) * \mathcal{N}(t_d;\sigma), \tag{2.4}$$

where $*$ denotes convolution. In the mean, the full lifetime histogram is

$$S(t_d) = A \cdot p(t_d) + b, \qquad p(t_d) = \sum_{k=1}^{3} c_k \cdot \mathrm{EMG}(t_d;\lambda_k,\sigma), \tag{2.5}$$

where $b \geq 0$ accounts for background, $A > 0$ is the expected total number of signal events, and the fractional intensities satisfy $c_k \geq 0$ and $\sum_k c_k = 1$, so that $p(t_d)$ is a probability density. We write $\tilde{S}(t_d) = S(t_d) - b$ for the background-corrected histogram in the mean.

2.2 Moments and exponential weighting

The method is built on moments of the histogram. For a function $f(t)$, we define the nth moment with translation a by

$$\mu_n^{(a)}\{f(t)\} = \int t_d^n \, f(t_d - a)\, dt_d = \int (t_d + a)^n \, f(t_d)\, dt_d, \tag{2.6}$$

and write $\mu_n\{f\} = \mu_n^{(0)}\{f\}$ for the ordinary moment. Viewed as an operator, $\mu_n\{\cdot\}$ weights the data by t_d^n. As n grows, the small-t_d region, where the short-lived DA and p-Ps components dominate, is progressively suppressed relative to the long-lived o-Ps tail. This mechanism selectively isolates λ_3.

Raising data to the nth power also amplifies noise at large t_d, where $\tilde{S}$ is small but its statistical fluctuations are not. To control this, we weight the histogram by e^{-st_d} with $s > 0$, which attenuates the large-t_d contribution, and work with $\mu_n^{(a)}\{e^{-st_d}\tilde{S}(t_d)\}$. The factor e^{-st_d} conversely amplifies noise for $t_d < 0$. However, the left tail of $\tilde{S}$ is Gaussian-like and negligible below a sufficient bound such as $t_{d,L} = -2.5\sigma$ (Appendix, Section A.2), so the lower limit of integration may be set to $t_{d,L}$ without loss. Throughout we take the translation to be $a = \sigma^2 s$, which is the choice that makes the exponential weighting interact cleanly with the EMG (Eq. (A.16)).

2.3 The moment-ratio estimator for a uniform sample

Consider a spatially uniform sample described by Eq. (2.5). Define the ratio of two successive exponentially weighted moments,

$$r^{(n)}(s) = \frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st_d}\,\tilde{S}(t_d)\}}{\mu_n^{(\sigma^2 s)}\{e^{-st_d}\,\tilde{S}(t_d)\}}. \tag{2.7}$$

The central result of the method, proved as Lemma A.6 in Section A.4, is that for a sufficiently large even n,

$$r^{(n)}(s) \approx \frac{n+1}{s+\lambda_3}. \tag{2.8}$$

This can be understood intuitively. Each EMG component contributes to Eq. (2.7), but the contributions of the two short-lived components ($k = 1, 2$) are suppressed relative to the o-Ps component by a factor that decays geometrically in n (Lemma A.5), because $s + \lambda_3$ is the smallest of the shifted rate constants. For large even n only the $k = 3$ term survives, and the ratio of successive moments of a single EMG component is $(n + 1)/(s + \lambda_3)$ (Lemma A.4).

Solving Eq. (2.8) for the rate constant gives $s + \lambda_3 \approx (n + 1)/r^{(n)}(s)$, and hence an estimator for the o-Ps lifetime,

$$\tau_3 \approx \frac{r^{(n)}(s)}{n + 1 - s \cdot r^{(n)}(s) + \delta}, \tag{2.9}$$

where δ is a small positive constant added to the denominator to mitigate the numerical instability that arises when dividing by a quantity that noise can drive close to zero. Odd n may also be permissible (and is explored in Chapter 3) if conditions are met (see Appendix, Lemma A.4). Eqs. (2.7)–(2.9) constitute the estimator for a uniform sample. In some sections of this work, the reciprocal of Eq. (2.9) is used instead, to estimate the o-Ps rate constant λ_3.

2.4 Extension to positronium lifetime tomography

In imaging, the rate constants λ_k and intensities c_k vary with position, and the task is to recover the spatial distribution of λ_3. We now attach a pixel index j to the quantities of Section 2.1, writing $c_{k,j}$, $\lambda_{k,j}$, and

$$p_j(t_d) = \sum_{k=1}^{3} c_{k,j}\,\mathrm{EMG}(t_d; \lambda_{k,j}, \sigma), \qquad q_j(t_d) = \sum_{k=1}^{3} c_{k,j}\,\mathrm{EXP}(t_d; \lambda_{k,j}), \tag{2.10}$$

so that q_j is the lifetime density for decays occurring in pixel j before time blurring, and $p_j = q_j * \mathcal{N}(\cdot;\sigma)$ is its blurred counterpart.

Conventional TOF-PET detects annihilation events, defined as a pair of 511 keV photons in coincidence. Let x_j be the number of decays occurring in image voxel j. Let i be a specific line of response (LOR) for non-TOF-PET, or a specific TOF bin on a specific LOR for TOF-PET. Ignoring scatter, the mean number of annihilation events y_i recorded in channel i is

$$y_i = \sum_j H_{ij}\, x_j + b_i, \tag{2.11}$$

where H_{ij} is the probability that a positron released in pixel j produces an annihilation event in channel i, and b_i is the number of random coincidences in channel i. In PLI with TOF-PET, an event is instead a triple coincidence comprising the annihilation pair and one prompt gamma, and so carries two pieces of information: 1) the channel number i and 2) the lifetime t_d. In place of y_i we therefore have a density $y_i(t_d)$, where $y_i(t_d)\, dt_d$ is the mean number of events in channel i whose lifetime lies in $[t_d, t_d + dt_d)$.

Let $\gamma_j < 1$ be the sensitivity of the system to prompt photons originating in pixel j. Assuming independent detection of annihilation photons and prompt gammas and neglecting positron range, a decay in pixel j produces a channel-i event with probability $\gamma_j H_{ij}$, and, with the common Gaussian timing uncertainty $\mathcal{N}(\cdot;\sigma)$,

$$\begin{aligned} y_i(t_d)\, dt_d &= \left\{ \sum_j \gamma_j H_{ij} \cdot x_j\, q_j(t_d)\, dt_d \right\} * \mathcal{N}(t_d;\sigma) + r_i\, dt_d \\ &= \sum_j H_{ij} \cdot \gamma_j x_j \cdot p_j(t_d)\, dt_d + r_i\, dt_d, \end{aligned} \tag{2.12}$$

where r_i is the random triple-coincidence density in channel i, taken to be uniform in t_d. Subtracting $r_i\, dt_d$ and applying $\mu_n^{(\sigma^2 s)}\{\cdot\}$ to both sides yields a linear system in the pixel-wise

weighted moments,

$$\tilde{Y}_i^{(n)}(s) = \sum_j H_{ij}\, f_j^{(n)}(s), \tag{2.13}$$

where

$$\tilde{Y}_i^{(n)}(s) = \mu_n^{(\sigma^2 s)}\{e^{-st_d} y_i(t_d) - r_i\}, \qquad f_j^{(n)}(s) = \gamma_j x_j \cdot \mu_n^{(\sigma^2 s)}\{e^{-st_d} p_j(t_d)\}. \tag{2.14}$$

Crucially, Eq. (2.13) has the same form as the ordinary emission-tomography model Eq. (2.11), with the weighted-moment data $\tilde{Y}_i^{(n)}(s)$ in place of the counts and the moment image $f_j^{(n)}(s)$ in place of the activity. The moment images can therefore be recovered with an existing PET or TOF-PET reconstruction algorithm. In this work we use the maximum-likelihood expectation-maximization (MLEM) algorithm. Reconstructing at orders n and $n+1$ and forming the voxel-wise ratio gives

$$r_j^{(n)}(s) \equiv \frac{f_j^{(n+1)}(s)}{f_j^{(n)}(s)} = \frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st_d} p_j(t_d)\}}{\mu_n^{(\sigma^2 s)}\{e^{-st_d} p_j(t_d)\}} \approx \frac{n+1}{s + \lambda_{3,j}}, \tag{2.15}$$

where the approximation is Eq. (2.8) applied pixel-wise, since p_j has the same three-component EMG structure as p. Inverting exactly as in Eq. (2.9), with $r^{(n)}(s)$ replaced by $r_j^{(n)}(s)$, yields the o-Ps lifetime (or rate-constant) image. Thus, the presented tomographic o-Ps lifetime imaging method requires only two ordinary moment-image reconstructions followed by a voxel-wise ratio and inversion. No per-voxel fitting is needed. In addition, for adequate choices of n and s the method is not sensitive to the presence of the DA and p-Ps components.

2.5 Estimation of the moment data from list-mode TOF-PET data

Applying the estimator requires, for each triple-coincidence event, its lifetime t_d, an estimate of the random-coincidence density r_i in each channel, and the weighted-moment data $\tilde{Y}_i^{(n)}(s)$. We treat each in turn.

2.5.1 Per-event lifetime

Consider one triple coincidence event. The annihilation photons are detected at positions $\vec{r}_1, \vec{r}_2$, at measured times t_1', t_2', and the prompt gamma at $\vec{r}_\gamma$ at measured time t_γ', where primes denote measured (arrival) times. Figure 2.1 shows the geometry of a triple coincidence event. Let c be the speed of light, $\ell_{12} = \|\vec{r}_1 - \vec{r}_2\|$ the LOR length, and $\vec{r}_e$ the annihilation point, taken to lie on the LOR. Writing each measured time as the corresponding emission or annihilation time plus the relevant photon time of flight gives

$$\frac{t_1' + t_2'}{2} - t_\gamma' = t_d + \frac{\ell_{12}}{2c} - \frac{\|\vec{r}_\gamma - \vec{r}_e\|}{c}, \tag{2.16}$$

subject to the constraint, since $\vec{r}_e$ lies between $\vec{r}_1$ and $\vec{r}_2$,

$$\|\vec{r}_1 - \vec{r}_e\| + \|\vec{r}_2 - \vec{r}_e\| = \ell_{12}. \tag{2.17}$$

The annihilation point on the LOR can be estimated by the TOF difference $\delta t_{12} \equiv t_1' - t_2' = (\|\vec{r}_1 - \vec{r}_e\| - \|\vec{r}_2 - \vec{r}_e\|)/c$. Combined with Eq. (2.17), this gives $\|\vec{r}_1 - \vec{r}_e\| = (c\,\delta t_{12} + \ell_{12})/2 = \ell_{12}(1+\alpha)/2$ with $\alpha = c\,\delta t_{12}/\ell_{12}$, so that

$$\vec{r}_e = \vec{r}_1 + \|\vec{r}_1 - \vec{r}_e\| \frac{\vec{r}_2 - \vec{r}_1}{\ell_{12}} = \beta_1 \vec{r}_1 + \beta_2 \vec{r}_2, \tag{2.18}$$

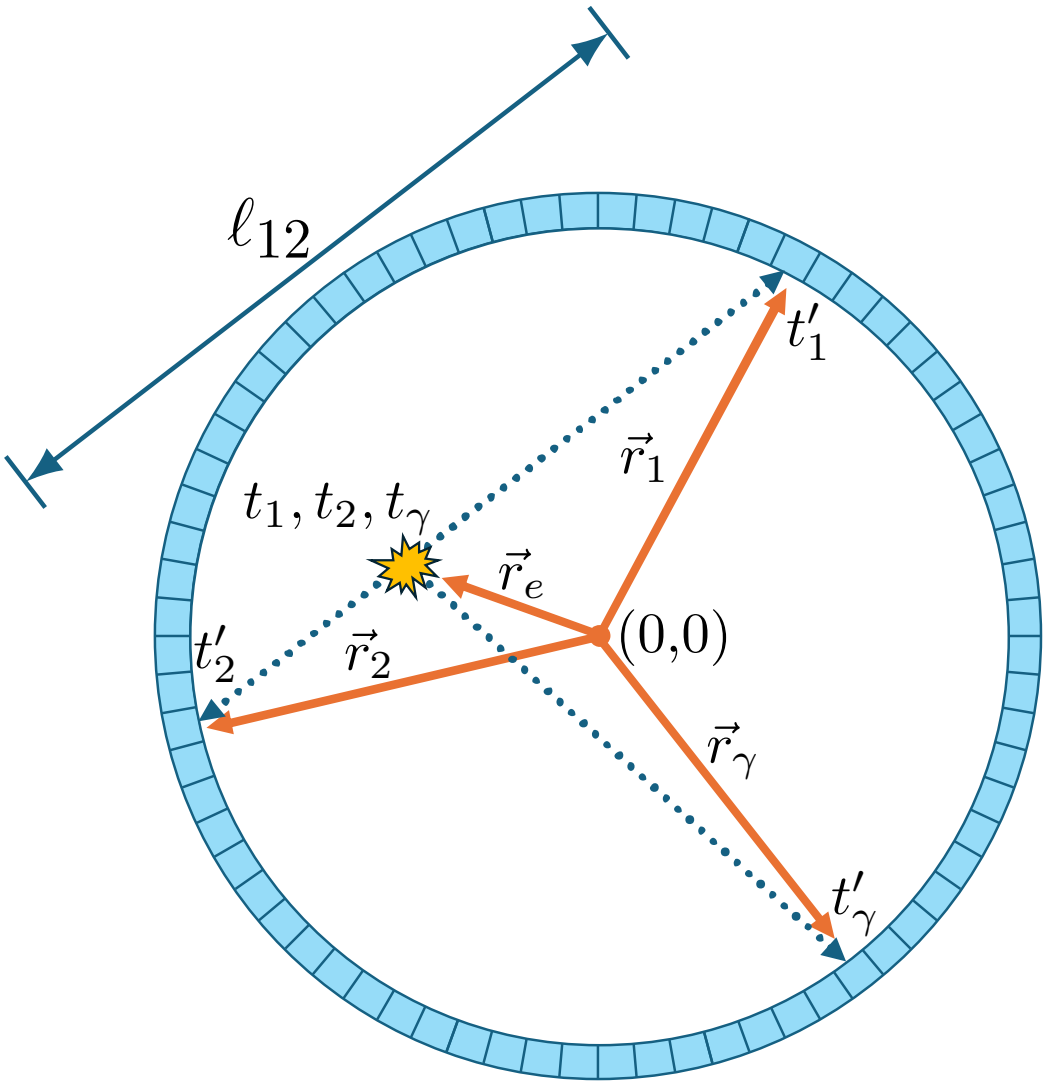

Figure 2.1: Geometry of a triple coincidence event. The annihilation point $\vec{r}_e$ lies on the line of response between the two annihilation photon detection positions $\vec{r}_1$ and $\vec{r}_2$.

where

$$\beta_1 = \frac{1-\alpha}{2}, \quad \beta_2 = \frac{1+\alpha}{2}. \tag{2.19}$$

Substituting Eq. (2.18) into Eq. (2.16), and using $\beta_1 + \beta_2 = 1$ to write $\vec{r}_\gamma - \vec{r}_e = \beta_1(\vec{r}_\gamma - \vec{r}_1) + \beta_2(\vec{r}_\gamma - \vec{r}_2)$, gives the per-event estimate

$$t_d = \frac{t_1' + t_2'}{2} - \frac{\ell_{12}}{2c} - t_\gamma' + \frac{\|\beta_1(\vec{r}_\gamma - \vec{r}_1) + \beta_2(\vec{r}_\gamma - \vec{r}_2)\|}{c}. \tag{2.20}$$

Intuitively, Eq. (2.20) can be understood as the difference of two corrected timestamps,

$$t_d = t_a - t_0, \tag{2.21}$$

the where t_a is the annihilation time and t_0 is the decay time, each recovered from a measured arrival time by subtracting the corresponding photon travel. The first and second terms give the annihilation time. Because $\vec{r}_e$ lies on the line segment joining the two detectors, Eq. (2.17) says that the two annihilation-photon path lengths sum to ℓ_{12} regardless of where along the LOR the annihilation occurred. The mean of the two arrival times is therefore delayed from the annihilation by $\ell_{12}/2c$, and so the annihilation time can be calculated by subtracting that delay from the mean arrival time,

$$\frac{t'_1 + t'_2}{2} - \frac{\ell_{12}}{2c} = t_a. \tag{2.22}$$

The third and fourth terms give the decay time. The prompt gamma travel time cannot be averaged with another photon, so its travel time must be calculated explicitly. The calculation requires $\vec{r}_e$ since,

$$t'_\gamma - \frac{\|\vec{r}_\gamma - \vec{r}_e\|}{c} = t_0. \tag{2.23}$$

The fourth term of Eq. (2.20) is the travel correction, written with β_1 and β_2. The geometric correction for the annihilation time does not depend on where along the LOR the event occurred, however the correction for the decay time does. An error in localizing $\vec{r}_e$ therefore propagates into t_d from the fourth term. A rough estimate of the error in t_0 is given. Consider a representative clinical TOF-capable PET scanner with a detector-ring diameter of approximately 80 cm [75]. The distance from an annihilation occurring near the center of the scanner to any given detector on is therefore approximately $L_\gamma \sim 40\,\text{cm}$. The corresponding

prompt-photon flight time is

$$t_{\gamma,\mathrm{TOF}} = \frac{L_\gamma}{c} \sim \frac{0.40\,\mathrm{m}}{c} \approx 1.33\,\mathrm{ns}. \tag{2.24}$$

Therefore, the geometric correction to the measured prompt-photon arrival time is on the order of 1 ns for a typical clinical PET scanner geometry.

More importantly, uncertainty in the annihilation position produces uncertainty in the geometric correction to t_0. The annihilation position along the LOR is determined by the measured TOF difference $\delta t_{12} = t_1' - t_2'$. The corresponding positional uncertainty is

$$\sigma_{r_e} = \frac{c}{2}\sigma_{\delta t}, \tag{2.25}$$

where $\sigma_{\delta t}$ is the standard deviation in the measured TOF difference. Taking $\sigma_{\delta t}$ to be the full width at half maximum (FWHM) of that distribution, we have $\sigma_{\delta t} = \mathrm{CTR}/2\sqrt{2\ln 2}$ and

$$\sigma_{r_e} = \frac{c\,\mathrm{CTR}}{4\sqrt{2\ln 2}} \approx 0.212\,c\,\mathrm{CTR}. \tag{2.26}$$

Let ξ be the annihilation position along the LOR with $\xi = 0$ at the midpoint of the LOR, $\vec{v}(\xi) = \vec{r}_\gamma - \vec{r}_e(\xi)$ from Eq. 2.23, and $\hat{u} = (\vec{r}_2 - \vec{r}_1)/\|\vec{r}_2 - \vec{r}_1\|$ be the unit vector along the LOR. Then,

$$\frac{d}{d\xi}\|\vec{v}\| = \hat{v} \cdot \frac{d\vec{v}}{d\xi}. \tag{2.27}$$

Because $\vec{r}_e(\xi) = (\vec{r}_1 + \vec{r}_2)/2 + \xi\hat{u}$,

$$\frac{d\vec{r}_e}{d\xi} = \hat{u} \tag{2.28}$$

and therefore

$$\frac{d}{d\xi}\|\vec{v}\| = -\hat{v} \cdot \hat{u} = -\cos\theta, \tag{2.29}$$

where θ is the angle between the LOR and the prompt photon direction. Uncertainty in $\vec{r}_e$

then propagates to the geometric correction as

$$\sigma_{t_0} = \frac{\sigma_{r_e}}{c}|\cos\theta| = \frac{\mathrm{CTR}}{4\sqrt{2\ln 2}}|\cos\theta|. \tag{2.30}$$

The maximum uncertainty therefore occurs for $\theta = 0$, giving

$$\sigma_{t_0}^{\max} = \frac{\mathrm{CTR}}{4\sqrt{2\ln 2}}. \tag{2.31}$$

Thus, a 200 ps CTR corresponds to a maximum contribution of approximately 42 ps to the uncertainty in t_0.

2.5.2 Random density and weighted moments

Consider list-mode acquisition, and let i_k and $t_{d,k}$ be the channel number and lifetime of the kth event. Let $S_i = \{k : i_k = i\}$ index the events in channel i. The lifetime histogram of all acquired events has been observed to fall to a constant level for $t_d \geq T_c$ with $T_c \approx 15$ ns [59]. Let $\mathcal{T} = [T_1, T_2)$ with $T_1 = -2$ ns and $T_2 = 30$ ns be the interval over which t_d is measured. These are conservative values that encompass the important structure of the spectrum and may be optimized in future work. Additionally, $\mathcal{T}$ is chosen empirically in Chapters 3 and 4. The random density is estimated by the average event rate seen in channel i over $\mathcal{T}$,

$$\hat{r}_i = \frac{\left|S_i \cap \left\{k : t_{d,k} \in [T_c,\, T_2)\right\}\right|}{T_2 - T_c}. \tag{2.32}$$

Approximating the moment integral by a bin sum and collapsing the bin sum to a sum over events (the derivation is shown in Appendix A.5), the weighted, random-corrected modified moment sinogram of channel i is

$$\tilde{Y}_i^{(n)}(s) \approx \sum_{k \in S_i} (t_{d,k} + \sigma^2 s)^n \, e^{-s t_{d,k}} - \hat{r}_i \cdot \alpha_n(s; \mathcal{T}), \tag{2.33}$$

where the random-subtraction weight

$$\alpha_n(s;\mathcal{T}) = \mu_n^{(\sigma^2 s)}\{e^{-st};\mathcal{T}\} = \int_{T_1}^{T_2} (t+\sigma^2 s)^n e^{-st}\,dt \tag{2.34}$$

has the closed form (Lemma A.7)

$$\alpha_n(s;\mathcal{T}) = \begin{cases} -\dfrac{n!}{s^{n+1}}\left[e^{-st}\displaystyle\sum_{k=0}^{n}\frac{\big(s(t+\sigma^2 s)\big)^k}{k!}\right]_{T_1}^{T_2}, & s>0, \\ \dfrac{T_2^{n+1}-T_1^{n+1}}{n+1}, & s=0. \end{cases} \tag{2.35}$$

Eq. (2.33) gives a simple single-pass accumulation. To do this, initialize $\tilde{Y}_i^{(n)}(s)$ and $\hat{r}_i$ to zero for all i. For the kth event, add $(t_{d,k}+\sigma^2 s)^n e^{-st_{d,k}}$ to $\tilde{Y}_{i_k}^{(n)}(s)$ and, if $t_{d,k}\in[T_c,T_2]$, increment $\hat{r}_{i_k}$ by $1/(T_2-T_c)$. Once all events are processed, subtract $\hat{r}_i\,\alpha_n(s;\mathcal{T})$ from $\tilde{Y}_i^{(n)}(s)$ for each channel i.

The following two chapters test method using the estimator, its tomographic extension, and the list-mode estimators. The first applies it to one-dimensional PALS measurements, characterizing its bias, variance, and dependence on the order of moment, the event count, the background level, and the truncation window. The second applies it to two-dimensional tomographic reconstruction, where the moment images of Eq. (2.13) are reconstructed by MLEM and combined through Eq. (2.15) to form o-Ps rate-constant maps.

CHAPTER 3

POSITRON ANNIHILATION LIFETIME SPECTROSCOPY

3.1 Introduction

While positronium lifetime imaging (PLI) is emerging for the purpose of producing full three-dimensional images of the orthopositronium (o-Ps) lifetime, this chapter is concerned with only the single-dimensional case represented by Eq. (2.5). For distinction, we therefore refer to this method as one for positron annihilation lifetime spectroscopy (PALS) measurements.

Presented here is a novel analytic method for estimating the o-Ps lifetime from $S(t_d)$. Current methods for this task include fitting a single exponential distribution to the time-difference histogram [59], and fitting a reduced full model given by Eq. (2.5) by assuming known values for certain parameters to the histogram [56]. Additionally, Shibuya et al. has proposed an inverse Laplace transform method to distinguish between positronium lifetimes while merging voxels for better statistics [76]. Their method is able to discern similar lifetime values, but still employs curve-fitting. In comparison with these curve-fitting methods, the method derived in Chapter 2 is an analytic method and is more computationally efficient, which is an important consideration for future application of the method for PLI in which lifetimes need to be obtained for a large number of voxels in a clinical setting. In addition, under mild and realistic conditions the analytic method is not sensitive to the unknown lifetimes for direct-annihilation (DA) and parapositronium (p-Ps), nor to the time-measurement uncertainty σ. Generally, this is not the case with curve-fitting methods.

3.2 Materials and Methods

Throughout this chapter the elapsed time is written t_d and moments are evaluated without translation, $\mu_n\{\cdot\} = \mu_n^{(0)}\{\cdot\}$ in the notation of Eq. (2.6). For the constant values we used in Eq. (2.5), we reference Moskal *et al.* [56]. These values are $\tau_1 = 1/\lambda_1 \approx 0.388\,\text{ns}$ and

$c_1 \approx 0.6$ for DA, $\tau_2 = 1/\lambda_2 \approx 0.125\,\text{ns}$ and $c_2 \approx 0.1$ for p-Ps, and $\tau_3 = 1/\lambda_3 \approx$ 1-10 ns and $c_3 \approx 0.3$ for o-Ps. As for the timing resolution, three values must be distinguished. Let σ_1 be the timing uncertainty for the detection time of a single photon, σ_{TOF} the uncertainty in the difference of two annihilation-photon arrival times, and σ the uncertainty in the elapsed time t_d, which is the width of the Gaussian in Eq. (2.4).

Let t_1', t_2', and t_γ' be the measured arrival times of the two annihilation photons and the prompt gamma, respectively. Assuming all channels have the same uncertainty σ_1, then for any fixed coefficients a_k the linear combination $\sum_k a_k t_k'$ has variance $\sigma_1^2 \sum_k a_k^2$. The coincidence timing resolution (CTR) value for a positron emission tomography (PET) scanner is measured from the coincidence time-difference distribution, $t_1' - t_2'$. By inspection, we may set $a_1 = 1$ and $a_2 = -1$, so $\sigma_{\text{TOF}}^2 = 2\sigma_1^2$ and $\sigma_{\text{TOF}} = \sqrt{2}\,\sigma_1$. Since CTR is the full width at half maximum (FWHM), $\sigma_{\text{TOF}} = \text{CTR}/2\sqrt{2\ln 2}$.

The elapsed time is formed as $t_d = (t_1' + t_2')/2 - t_\gamma'$, for which $a_1 = a_2 = \frac{1}{2}$ and $a_3 = a_\gamma = -1$, so $\sigma^2 = \sigma_1^2\left(\frac{1}{4} + \frac{1}{4} + 1\right) = \frac{3}{2}\sigma_1^2$. Eliminating σ_1 using the three relationships ($\sigma_{\text{TOF}} = \text{CTR}/2\sqrt{2\ln 2}$, $\sigma_{\text{TOF}} = \sqrt{2}\,\sigma_1$, and $\sigma = \sqrt{\frac{3}{2}}\,\sigma_1$) gives σ in terms of the CTR of a given PET scanner:

$$\begin{aligned}
\sigma &= \sqrt{\tfrac{3}{2}}\,\sigma_1 \\
&= \sqrt{\tfrac{3}{2}} \cdot \frac{\sigma_{\text{TOF}}}{\sqrt{2}} \\
&= \frac{\sqrt{3}}{2}\,\sigma_{\text{TOF}} \\
&= \frac{\sqrt{3}}{2} \cdot \frac{\text{CTR}}{2\sqrt{2\ln 2}} \\
&\approx 0.368\,\text{CTR}.
\end{aligned} \tag{3.1}$$

Taking CTR = 600 ps for a representative current clinical time-of-flight (TOF)-PET system [56, 77], gives $\sigma_{\text{TOF}} \approx 255\,\text{ps}$ and $\sigma \approx 221\,\text{ps}$. The simulations reported below use

$\sigma = 169\,\text{ps}$, chosen together with other parameters to match the measured histogram of Moskal *et al.* [56], corresponding by Eq. (3.1) to CTR $\approx 460\,\text{ps}$.

In theory, $\tilde{S}(t_d)$ decreases exponentially to zero as t_d increases, which allows $\mu_n\{\tilde{S}(t_d)\}$ to be well-defined. In reality, noise in $\tilde{S}(t_d)$ does not necessarily decrease with t_d and hence will contribute a substantial statistical error in $\mu_n\{\tilde{S}(t_d)\}$, especially for large n. This instability can be considerably reduced by instead using $\mu_n\{e^{-st_d}\tilde{S}(t_d)\}$, $s > 0$ as the term e^{-st_d} attenuates the contribution from data at large t_d. Although a large s is favored for alleviating the effects of noise, it also diminishes the differences among $s+\lambda_i$ and thereby requires a large n for Eq. (2.8) to hold true. However, calculation of higher-order moments is more susceptible to noise so the value of s needs to be chosen with care. In this work, in consideration of the numerical values of λ_i and c_i, we choose $s = 1$, which is empirically justified based on the results to be reported below. Future work will consider s more extensively.

From Eq. (2.6), $\mu_n\{f(t)\}$ can be regarded as a soft threshold that removes an increasingly wider range of small t data as n increases. By inspection of Fig. 3.1 it can be stated that by choosing an n which makes Eq. (2.8) valid, a "soft" low cutoff has been introduced. This avoids using data where DA and p-Ps are dominant.

3.2.1 Simulated data

The proposed method is evaluated by using computer generated simulation data. All computation codes were implemented in Python version 3.11, using specified values for b, c_i, λ_i, and σ, and the desired total number of events for the histogram N. The simulation program first computed $S(t_{d,i})$ according to Eq. (2.5) at discrete times in $[-5\,\text{ns},\ 25\,\text{ns}]$ and at a regular spacing of 60 ps. Next, it scaled $S(t_{d,i})$ to yield $\sum_i S(t_{d,i}) = N$. Then, the scaled $S(t_{d,i})$ was replaced by an integer drawn by a Poisson random number generator (from the numpy.random module) whose mean equals the scaled $S(t_{d,i})$. Fig. 3.1 shows an example

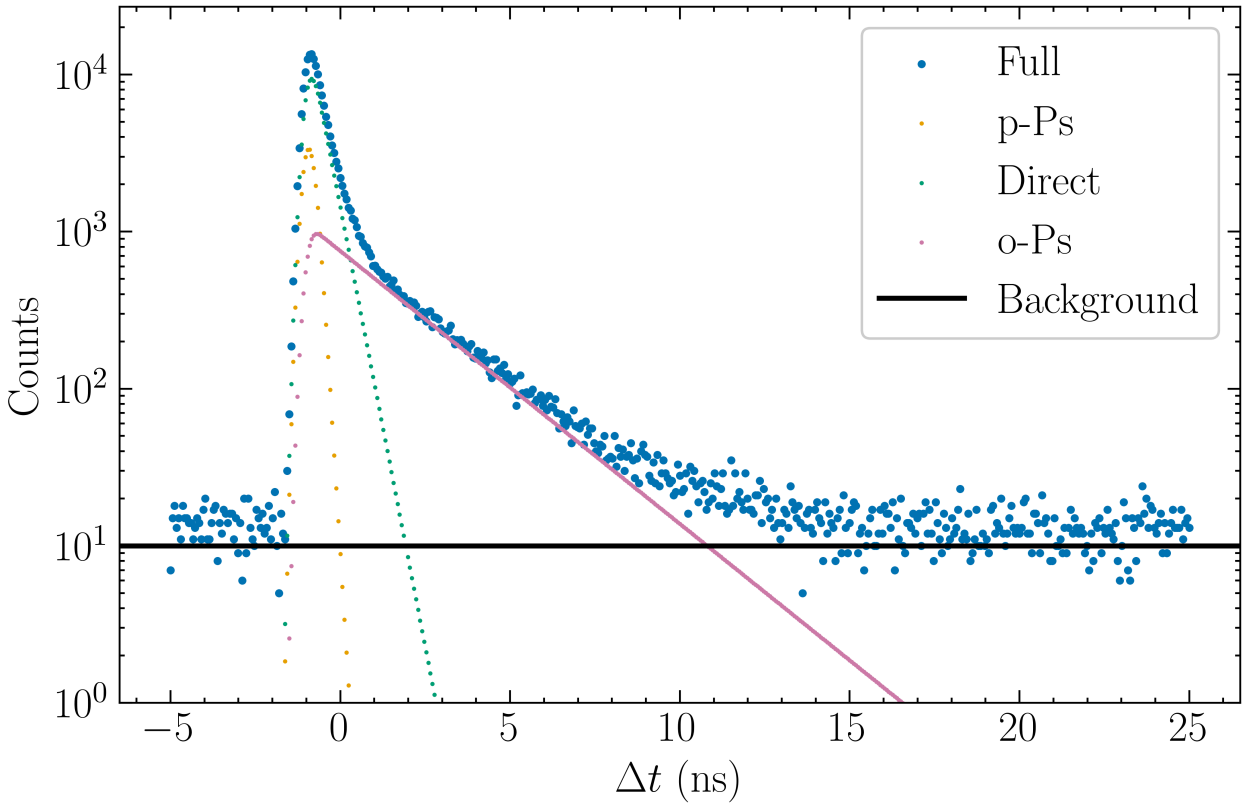

Figure 3.1: A simulated histogram of t_d having 2×10^5 events and using a 60 ps bin size. It was generated by using the following parameters: $\lambda_1 c_1$, $\lambda_2 c_2$, $\lambda_3 c_3 = 1.548$, 0.8, 0.12, τ_1, τ_2, $\tau_3 = 0.388$ ns, 0.125 ns, 2.5 ns, $\sigma = 0.169$ ns, $b = 10$, and the Gaussian rise position was offset by 9.97 ns so the peak was near-centered after the convolution. $\lambda_1, \lambda_2, \lambda_3 = 2.58\,\text{ns}^{-1}$, $8\,\text{ns}^{-1}$, $0.4\,\text{ns}^{-1}$. The individual contributions from the DA, p-Ps, and o-Ps components are shown in green, yellow, and magenta, respectively.

of the generated histogram where the parameters were so chosen that it was similar to the measured histogram reported by Moskal et al. [56]. Each simulated histogram, except where noted, contained $N = 2 \times 10^5$ total events.

3.2.2 Implementation and numerical studies

The background term b was estimated using the average of the histogram in the $t_d < 0$ ns region. The estimate was then subtracted from each histogram data point and λ_3 was calculated according to Eq. (2.8) and Eq. (2.9). For computing moments, Eq. (2.6) was

approximated by naive discrete summation: $\mu_n\{f(t)\} \approx \delta t \sum_i t_i^n \tilde{S}(t_i)$ where t_i are the time points where $f(t)$ is available and δt is the spacing of t_i.

The proposed method was evaluated for accuracy and precision against a number of parameters, including 1) the order of moment n used, 2) the upper truncation of the histogram data, 3) the number of counts N, and 4) the background level b. At present, the range of *in vivo* o-Ps lifetime values has not been precisely established. However, in a recent paper Moskal et al. observed that cardiac myxoma and adipose tissue had mean lifetimes of 1.912 ns and 2.613 ns, respectively [56]. Therefore, we performed evaluations for o-Ps lifetimes of 1.0, 1.5, 2.0, 2.5, and 3.0 ns to cover the likely *in vivo* lifetime range. As mentioned above, lifetimes of the DA and p-Ps components were held fixed at 388 ps and 125 ps respectively, and their population weights were also fixed at 0.6 and 0.1, respectively.

3.3 Results

The results are presented by Fig. 3.2 through Fig. 3.9. Each figure contains individual points which are estimates of τ_3, denoted by $\hat{\tau}_3$ below, derived from our method. Each data point is the mean of the results obtained from 1×10^4 histograms simulated by using the same parameters (1×10^5 in the case of Fig. 3.2), and the shaded regions in the plots give the ± 1 standard deviations (std) about the means. The horizontal lines, when present, indicate the true o-Ps lifetimes that are used to produce simulation data.

3.3.1 Lifetime estimate versus order of moment n

Fig. 3.2 shows the estimated o-Ps lifetime when the order of moment n employed by Eq. (2.9) is varied. Four general trends can be observed. First, all curves show a plateau where the estimated lifetime has essentially zero bias. This plateau occurs between $n \approx 5$ and $n \approx 16$ depending on the o-Ps lifetime. The standard deviations of the estimates are also sufficiently small to allow for discrimination of all the lifetimes examined. Second, the standard deviation

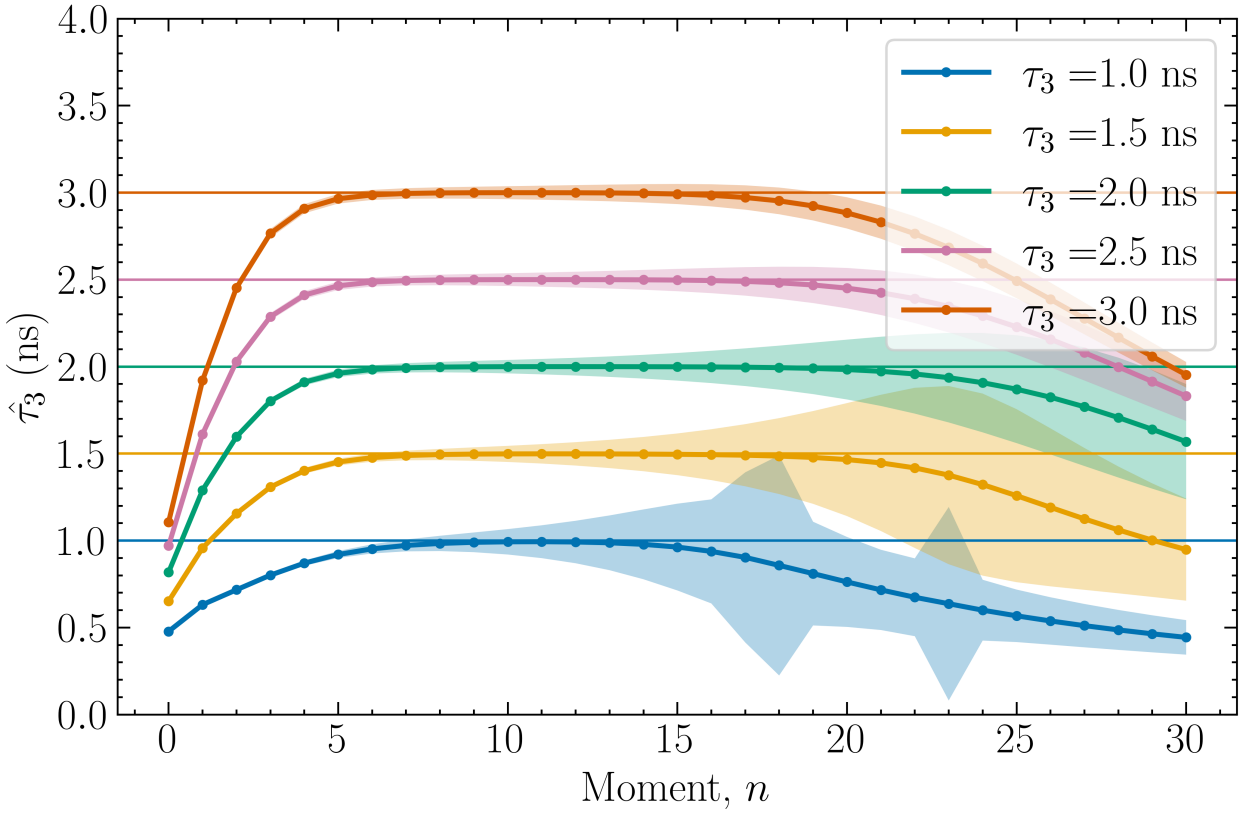

Figure 3.2: Estimated o-Ps lifetime as a function of the order of moment n. The shaded areas for each curve represent the ± 1 standard deviation on each data point. The true τ_3 values shown in the legend are those used for producing simulation data.

increases with n, which is consistent with the observation made above that higher-order moments are more sensitive to data noise. Third, as τ_3 increases the plateau occurs at a lower n. This is because the differences between $s+\lambda_3$ and $s+\lambda_i$, $i=1,2$ increases, allowing the $i=3$ term to dominate the sum in Eq. (2.5) at a smaller n. Fourth, all curves decrease toward zero as the order increases. This is because higher-order moments are increasingly more contributed by data at larger t_d while data is simulated only for $-5\,\text{ns} \leq t_d \leq 25\,\text{ns}$. In practice, the measured histogram is necessarily truncated.

It is also noted that all curves converge at small n. This reflects the situation that with a small n the $i=1,2$ terms in fact dominate the $i=3$ term in Eq. (2.5). As a result, Eq. (2.9) yields some fixed value since c_k, and τ_k, $k=1,2$ are constants. Based on the plot,

Table 3.1: Bias-minimizing estimates of Fig. 3.2, shown with the moment n which minimized the bias. We again emphasize that this moment-based method requires n to be sufficiently large.

τ_3 (ns)	order n	$\hat{\tau}_3$			
		mean (ns)	std (ns)	std/mean	% error
1.0	11	0.99	0.094	0.095	0.654
1.5	11	1.50	0.056	0.037	0.096
2.0	12	2.00	0.051	0.025	0.008
2.5	10	2.50	0.041	0.016	0.001
3.0	10	3.00	0.037	0.012	0.009

we propose that in general $n \lesssim 5$ should not be used for PALS o-Ps estimations. Table 3.1 summarizes the bias-minimizing order of moment in the plateau region and the statistics of the estimates obtained when using this order. As shown, the largest percentage error ratio is only 0.654%, obtained for $\tau_3 = 1.0$ ns.

Fig. 3.3 shows the result obtained when the τ_3 values measured for cardiac myxoma and adipose tissue by Moskal et al. [56] were used to produce simulation data. In this case, the bias-minimizing orders of the moments were found to be $n = 11$ for $\tau_3 = 1.912$ ns (myxoma tissue), which yielded $\hat{\tau}_3 = 1.91 \pm 0.05$ ns, and $n = 11$ for $\tau_3 = 2.613$ ns (adipose tissue), which yielded $\hat{\tau}_3 = 2.61 \pm 0.04$ ns. The means of these estimates are within 0.02% and 0.005% of their respective true values.

3.3.2 Lifetime estimate versus number of events N

The results discussed were obtained by using simulated data with $N = 2 \times 10^5$. Generally, a histogram derived from a larger number of events has better statistics and can lead to better estimates for λ_3. Fig. 3.4 shows the results for N ranging from 1×10^3 and 1×10^8. The order of moment was fixed to $n = 11$. A tabulated summary of these results are shown in Table 3.2.

It can be seen that, as expected, for all the τ_3 examined the standard deviations of the

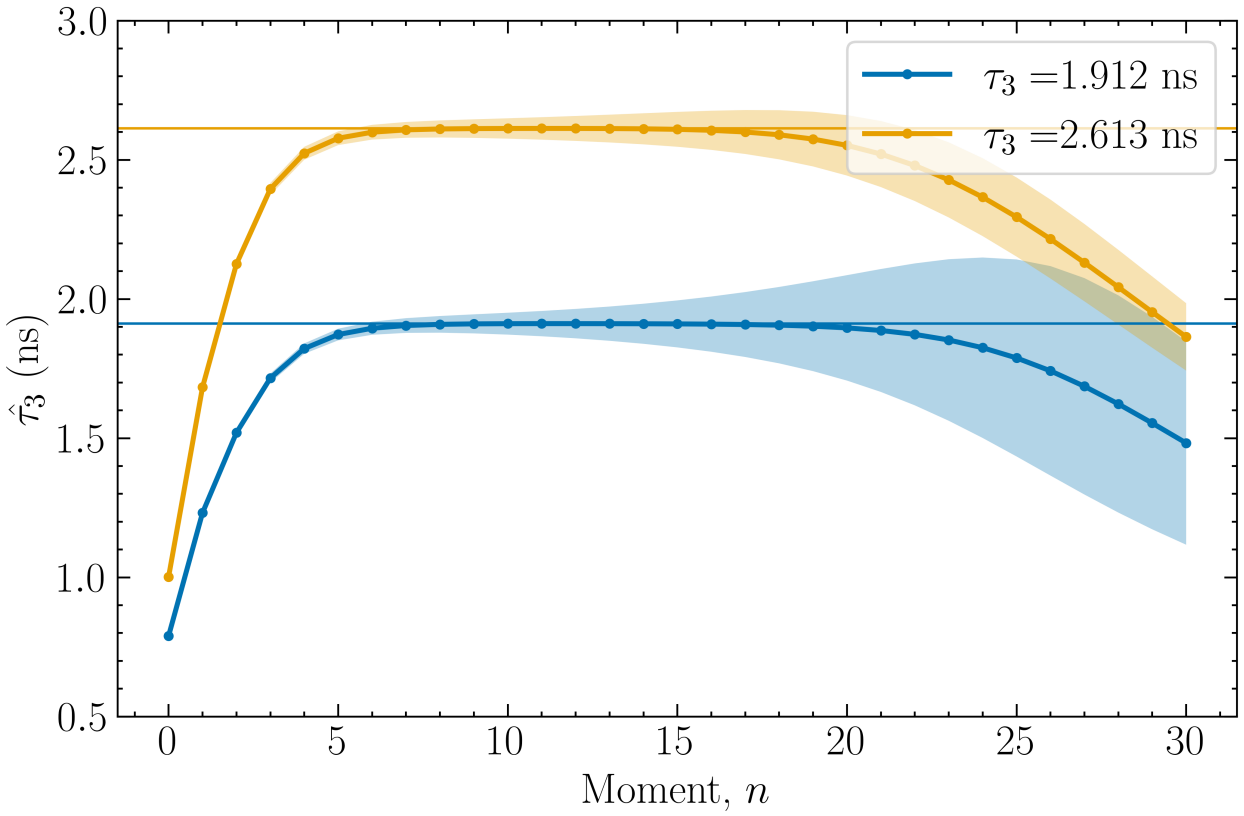

Figure 3.3: Similar to Fig. 3.2, this shows two practical cases for o-Ps lifetimes which have been measured recently *in vivo*

Table 3.2: Selected estimates from Fig. 3.4.

τ_3 (ns)	$\hat{\tau}_3$ (ns)			
	$N = 10^3$	$N = 10^4$	$N = 10^5$	$N = 10^6$
1.0	0.6 ± 0.6	0.9 ± 0.4	1.0 ± 0.1	0.997 ± 0.004
1.5	2 ± 16	1.4 ± 0.3	1.49 ± 0.08	1.498 ± 0.002
2.0	2.0 ± 0.8	2.0 ± 0.2	2.00 ± 0.06	2.000 ± 0.002
2.5	2.5 ± 0.6	2.5 ± 0.2	2.50 ± 0.06	2.500 ± 0.001
3.0	3.0 ± 0.6	3.0 ± 0.1	3.00 ± 0.06	3.000 ± 0.001

estimates decrease continually as N increases. The estimates asymptotically approach the true values as N increases and the statistics are seen to improve. The occurrences of some exceptionally large standard deviations reflect the instability of the ratio estimate given by Eq. (2.9) when the denominator is erroneously small with respect to the numerator due to

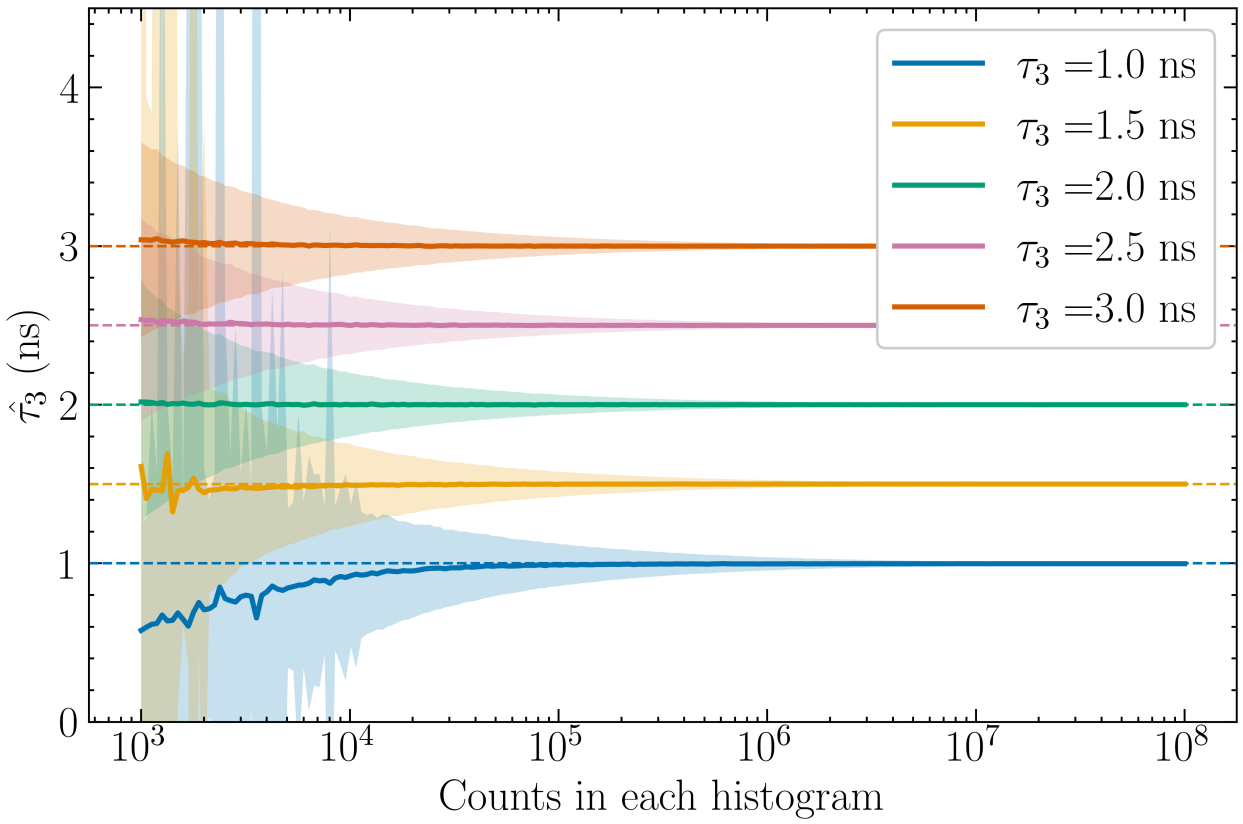

Figure 3.4: Estimated o-Ps lifetimes as a function of the number of events N in the histogram. The shaded areas again indicate ± 1 standard deviation about the mean.

data noise. Fig. 3.5 shows a simulated histogram for $N = 1.5 \times 10^3$ and $\tau_3 = 1.5\,\text{ns}$ when one such case occurs. It can be seen that the $t_d > 4\,\text{ns}$ region of the histogram where o-Ps is assumed to dominate is sparsely populated, leading to large and inconsistent errors in the numerator and denominator of Eq. (2.9). Controlling this variability will be commented on in Sect. 3.4.3.

3.3.3 Lifetime estimate versus cutoff and background

In an attempt to further alleviate the deleterious effects of noise, we also examined removing data above certain t_d. Fig. 3.6 shows that, for the case of $n = 11$, as the upper truncation threshold T_2 is decreased, the standard deviation decreases, which is particularly evident for small τ_3, whereas the bias increases in the form of underestimation. However, as shown in

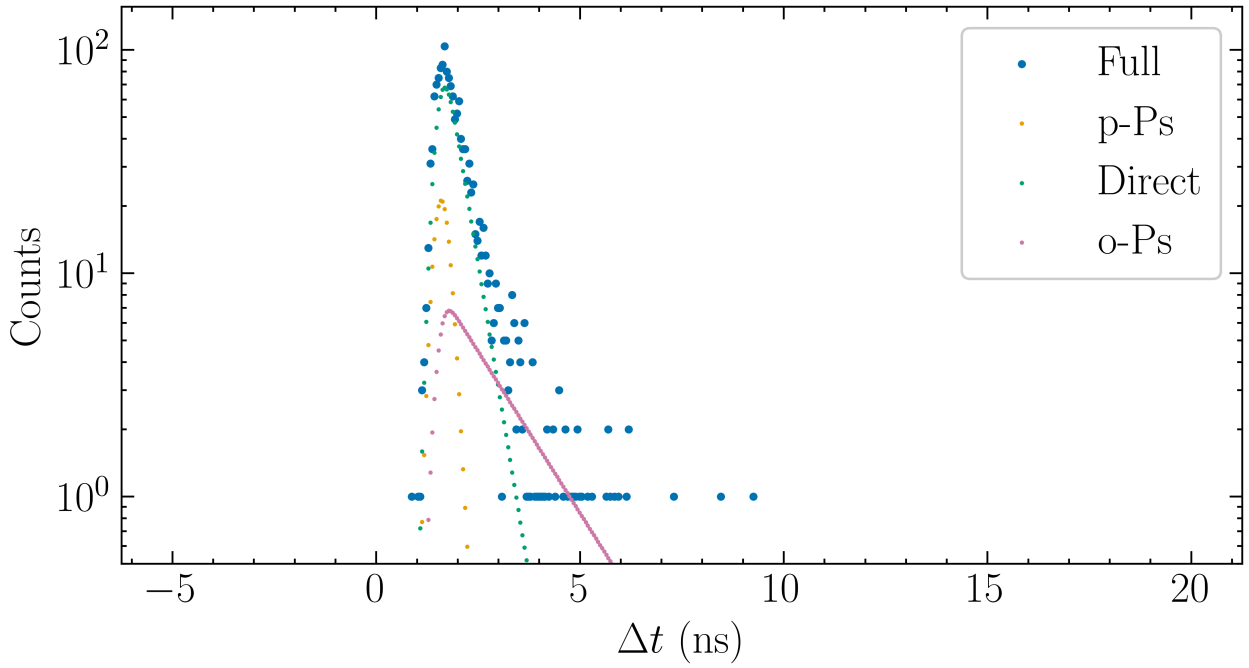

Figure 3.5: A simulated histogram for $\tau_3 = 1.5\,\text{ns}$ and $N = 1,500$, including the noise-free contributions of DA, p-Ps, and o-Ps.

Fig. 3.7, the standard deviations relative to the distances between means decreases as T_2 is lowered, at least for T_2 above $\sim 10\,\text{ns}$. Therefore, at least for discrimination tasks, it is beneficial to apply truncation even though it introduces bias.

The presence of a nonzero background b increases data noise in two manners. First, for a given number of total events the number of events contributing to the signal is lowered. Second, although b may be estimated and subtracted from $S(t_d)$, the noise associated with it remains in the data. Due to its Poisson nature, the variance of the associated noise is proportional to b. Fig. 3.8 shows the relationship between the estimate $\hat{\tau}_3$ and the true τ_3 for three background levels with $n = 11$, $N = 2 \times 10^5$, and no upper truncation. The background is modeled as Poisson noise with means of $b = 0, 10$, and 20 counts.

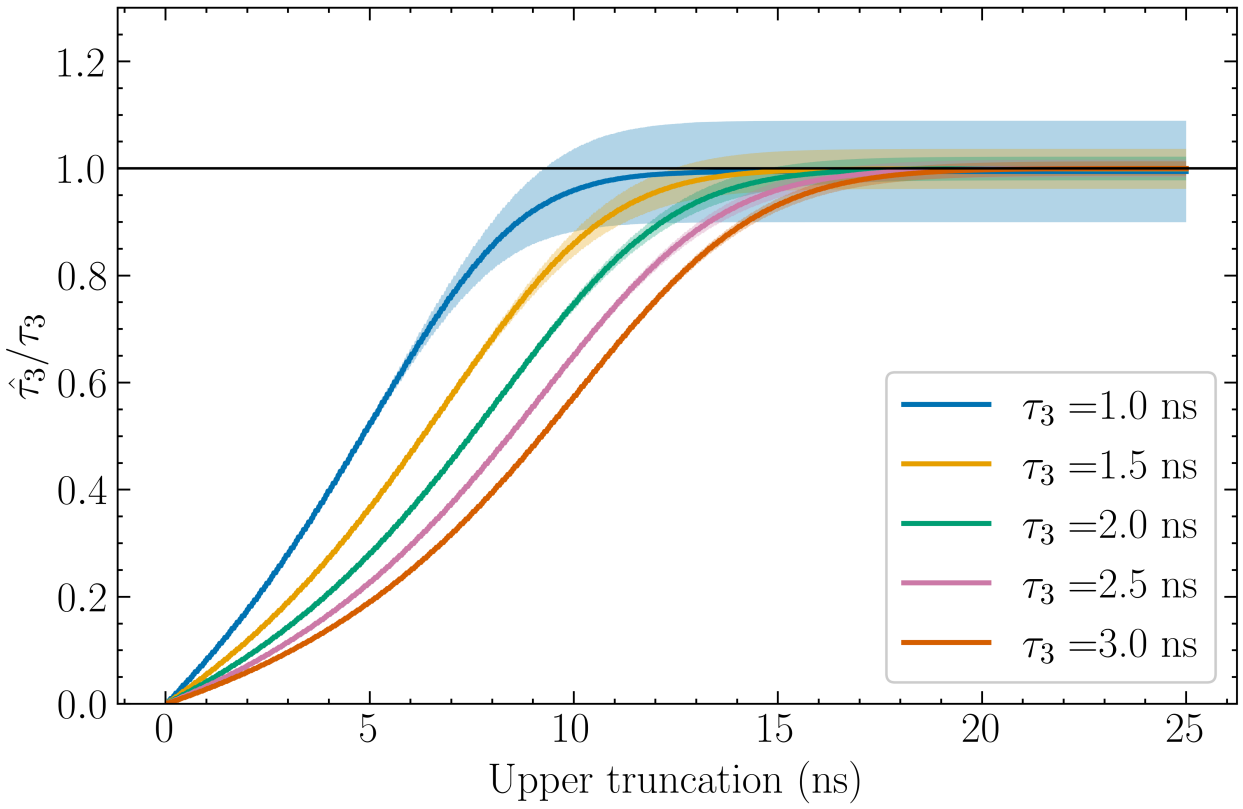

Figure 3.6: Estimated o-Ps lifetime as a function of the upper truncation threshold T_2. Note that here $\hat{\tau}_3/\tau_3$ is plotted. The horizontal line indicates the perfect estimate given by $\hat{\tau}_3/\tau_3 = 1$.

3.3.4 Monotonicity

We have observed that for discrimination tasks it is beneficial to allow for some bias if the std/mean ratio of the estimate decreases. For quantitative tasks, as long as $\hat{\tau}_3$ is related to τ_3 monotonically in the mean, this observation remains true because we can correct for the bias in $\hat{\tau}_3$ by using, for example, a predetermined calibration curve. Such monotonicity is observed in both Fig. 3.6 and Fig. 3.8. Since curves in Fig. 3.2 do not cross one another this monotonicity is also true when varying n. For $b = 0, 10$, and 20 mean counts, $\hat{\tau}_3$ remains monotonic with τ_3 in the mean and is approximately linear for $\tau_3 \gtrsim 1.5\,\text{ns}$. Below $\sim 0.5\,\text{ns}$, monotonicity may not hold for $n = 11$, however this is well below the range for currently-known biological values of τ_3.

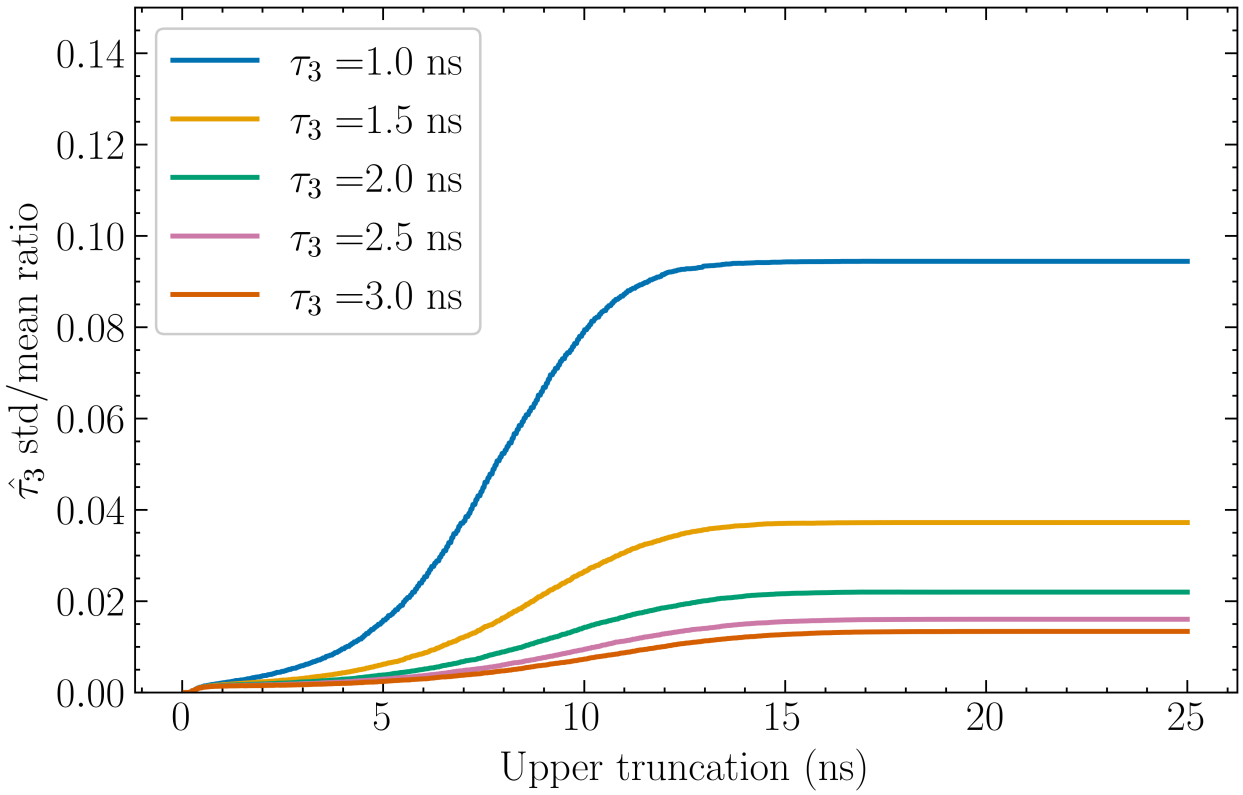

Figure 3.7: The ratio of the standard deviation to the mean of the curves from Fig. 3.6. The ratios are seen to decrease as the upper truncation decreases.

3.4 Discussion

3.4.1 Effect of the number of events

It was observed that our moment-based method is significantly affected by Poisson noise, which increases with decreasing event counts. However, in the mean, the method remains monotonic with the true value for PALS. For an illustrative estimate, consider a typical PET scan with approximately 185 MBq (5 mCi) of activity and a voxel volume of $64\,\mathrm{mm}^3 = 0.064\,\mathrm{cm}^3$. Assuming the human body consists primarily of water with a total mass of 80 kg, the total internal volume is then approximately $8 \times 10^4\,\mathrm{cm}^3$. Under the additional assumption of a uniform activity distribution, this corresponds to an activity concentration of approximately $2.3 \times 10^3\,\mathrm{Bq/cm}^3$ and an expected 148 Bq/voxel. If an effective de-

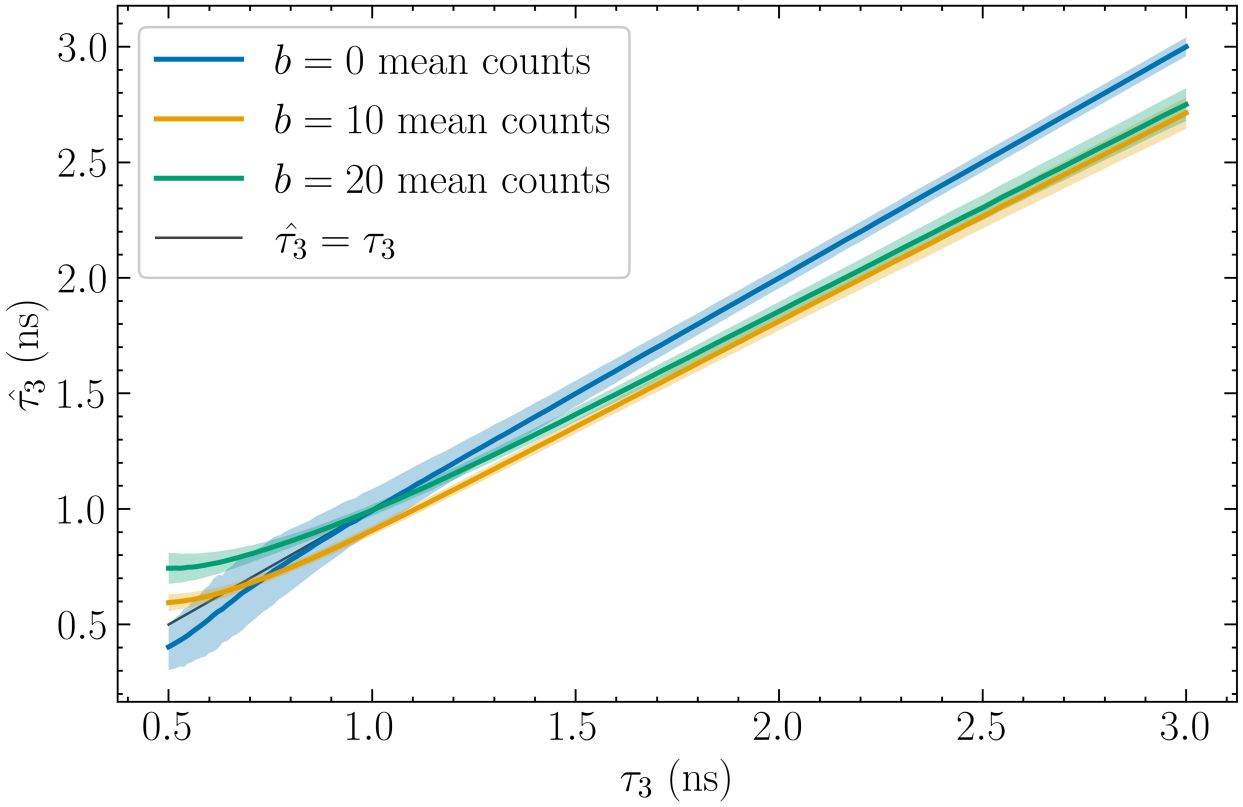

Figure 3.8: Estimated o-Ps lifetime versus the true value for three background levels b with $n = 11$, $N = 2 \times 10^5$ and no upper truncation. The region close to $\tau_3 = 0.5$ ns may not be monotonic for some cases due to the curvature.

tection efficiency of 1% is assumed for usable PLI events, approximately 1.48 counts/s/voxel would be collected. Thus, acquiring 1×10^3 counts per voxel would require approximately 11.3 min. For reference, Moskal *et al.* reported human brain PLI measurements using a 10 min acquisition, with 342 counts in healthy brain tissue, 547 counts in tumor tissue, and 1119 counts in the salivary glands [65]. The method presented in this chapter is applied to simulated PET data in the next chapter, and the effects of noise are further discussed when reconstruction is used. We note however that developments in long-axial field-of-view PET scanners, including total-body scanners, will increase the sensitivity and reduce the time needed to collect sufficient counts for PLI [78, 79]. Additionally, above we have assumed a uniform radiotracer distribution. When using tumor-specific radiotracers, the radioactivity can be concentrated to tumor regions which will improve the statistics.

3.4.2 Short o-Ps lifetimes in vivo

This moment-based method has been shown to produce viable results for lifetimes above approximately 1 ns, depending on noise, moment, and truncation. For lifetimes that are on the order of 1 ns, standard deviations for $n > 15$ (seen in Figure 3.2) grow rapidly, and estimations may no longer be monotonic.

3.4.3 Drawbacks and future work

A main drawback of this method is the response to noise. The standard deviations in Figure 3.4 are a significant fraction of the estimate itself, even though the mean is stable. However, this method will be computationally less expensive to execute over an entire PET dataset than fitting-based methods, which may outweigh its precision for low-count cases. Shibuya et al. has previously calculated that approximately 3×10^8 counts will be needed to estimate oxygen concentration to a precision of 10 mmHg [59]. This would correspond to three times as many counts as the extent of Figure 3.4. At this point, the method has essentially reached its asymptote, minimum bias, and minimum standard deviation.

To control variations in the standard deviation at low counts, a small positive number δ may be introduced, as shown in Eq. (2.9). This results in an increased bias and decreased standard deviation seen in Figure 3.9. Despite the increased overall bias in the asymptotic region (37.6% error), the curve remains monotonic and asymptotic in the mean, allowing for the bias to be easily corrected. For comparison, with 1×10^3 histogram counts, the standard deviation decreased from 0.62 ns for $\delta = 0$ to 0.19 ns for $\delta = 0.3$, while the percentage error increased from 42.9% to 62.4% for the two cases, respectively.

Future work on this estimation method will include its testing on full three-dimensional reconstruction, testing its computational speed with respect to current estimation techniques, and further optimization of fitting parameters, such as the lower truncation point, the e^{-st} term, smoothing of the histogram, and further investigation into noise handling modifica-

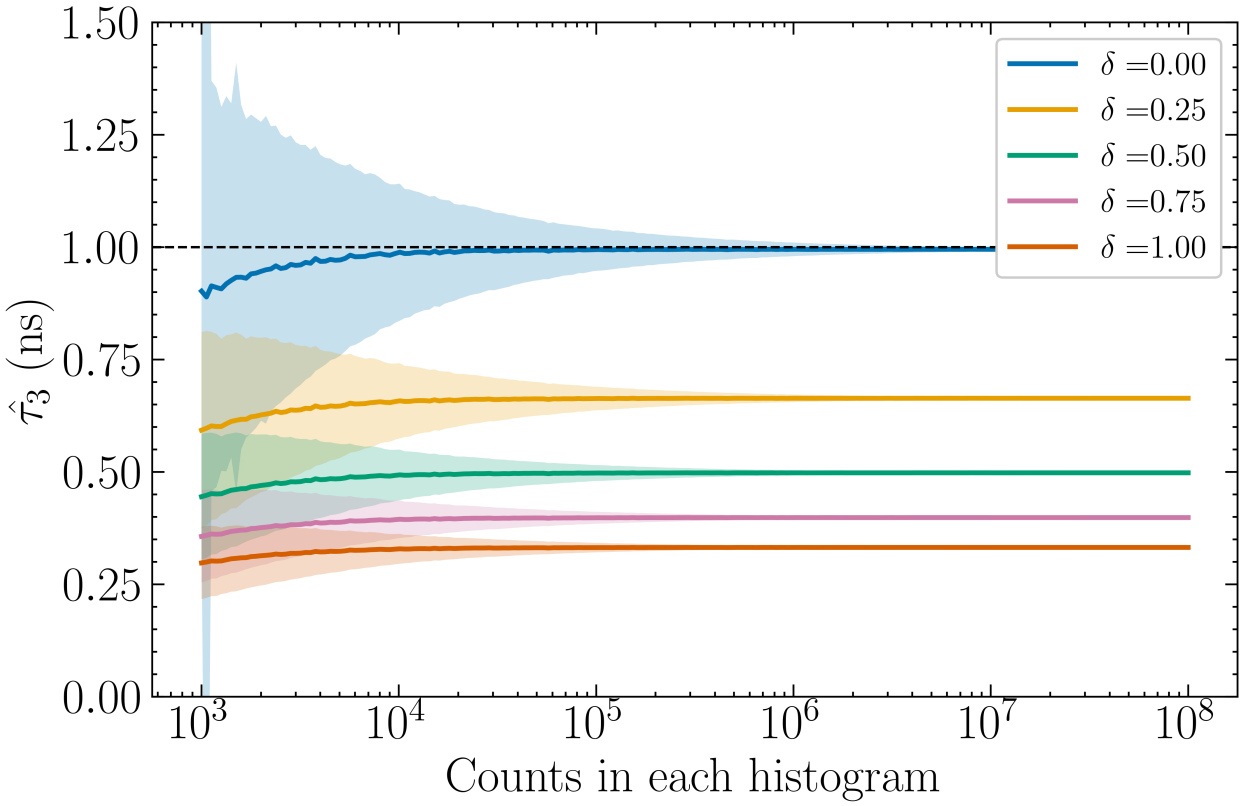

Figure 3.9: Several estimates of the $\tau_3 = 1$ ns case from Figure 3.4, however, with δ varied to demonstrate how the standard deviation variability can be handled for small τ_3 values. Note that increasing the value of δ increases the bias, but decreases the standard deviation for any given number of counts.

tions, such as the δ term. A known limitation of the proposed method is small denominators will produce large fluctuations in the estimate, the δ term attempts to mitigate this effect, however further investigation into this issue will be performed.

3.5 Conclusion

In this chapter we present an analytical method to estimate the orthopositronium lifetime in PALS measurements. This method uses moments of the histogram of arrival time differences, and employs an exponential weighting to mitigate numerical instability in calculation of moments from noisy data. The moment-based method was characterized in this work, and it was shown to be a stable, monotonic estimate in most cases. For cases in which the

standard deviation was large, modifications to the method may be employed.

CHAPTER 4

POSITRONIUM LIFETIME TOMOGRAPHY

4.1 Introduction

Chapter 2 developed an analytic, moment-based estimator of the orthopositronium (o-Ps) rate constant from a positron lifetime histogram and showed that it extends to tomography as a pair of ordinary emission-tomography reconstructions followed by a voxel-wise ratio. This chapter evaluates it on simulated two-dimensional time-of-flight (TOF)-positron emission tomography (PET) data. Because the method in this chapter recovers an image rather than a single lifetime, we refer to it here as positronium lifetime tomography (PLT) to emphasize its tomographic capability. PLT is of interest for hypoxia imaging because it may allow for the o-Ps lifetime to be estimated as a metric for local oxygen concentration without a dedicated hypoxia-targeting tracer, and so avoids the perfusion and delivery limitations described in Chapter 1 for ^{18}F-based hypoxia tracers. Whether it is effective for *in vivo* hypoxia imaging remains an open question, however.

One main difficulty in PLT is that the elapsed time must be referenced to the annihilation position along the line of response (LOR), which can only be localized with finite timing precision according to the coincidence timing resolution (CTR). Consider an annihilation occurring a distance ξ from the midpoint of an LOR of length ℓ_{12}. The two 511 keV photons travel in opposite directions, so their path lengths are $\ell_{12}/2 + \xi$ and $\ell_{12}/2 - \xi$, and the difference in their arrival times is

$$\delta t_{12} = \frac{2\xi}{c}, \qquad \xi = \frac{c\,\delta t_{12}}{2}. \tag{4.1}$$

A factor of two arises because displacement of the annihilation position changes both photon path lengths. An uncertainty in the arrival-time difference therefore produces a positional

uncertainty of

$$\mathrm{FWHM}_{\xi} = \frac{c}{2}\,\mathrm{CTR} \approx 0.015\,\mathrm{cm/ps} \times \mathrm{CTR}, \tag{4.2}$$

where the CTR is the full width at half maximum (FWHM) of the coincidence time-difference distribution. Thus, an improvement of 100 ps in CTR corresponds to an approximately 1.5 cm improvement in localization accuracy [80]. The corresponding standard deviation is $\sigma_{\xi} = c\,\sigma_{\mathrm{TOF}}/2$, with $\sigma_{\mathrm{TOF}} = \mathrm{CTR}/(2\sqrt{2\ln 2})$.

Eq. (4.2) shows how the timing resolution limits the spatial localization. The corresponding localization widths are 3.0 cm for 200 ps, 9.0 cm for 600 ps, and 45 cm for 3000 ps. In the phantom used for the stimulation study presented in this chapter, the two regions of interest are 2.8 cm in diameter, separated by 5.23 cm, within a 13.4 cm field of view. Thus, the localization width exceeds the region diameter even at 200 ps, exceeds the region separation at 600 ps, and exceeds the field of view by more than a factor of three at 3000 ps. Each event therefore samples a spatial region that may contain multiple tissue types, and the o-Ps lifetimes can differ across the region [80]. In direct positioning methods, which assign each event to its TOF-estimated position along the LOR, events originating at different locations may be assigned to the same reconstructed position. This assignment error could cause the resulting lifetime estimate to be influenced by multiple tissue regions. Eq. (4.2) describes localization uncertainty along the LOR, however transverse localization is determined by the detector and crystal geometries and is not directly set by the CTR. Additionally, we also use the estimated annihilation position to correct the prompt-photon travel time in Eq. (2.20). Mislocalization therefore introduces both spatial assignment error and error in the corresponding travel-time correction, compounding the effect of finite timing resolution.

With these considerations, we characterize the developed estimator against four parameters, the order of moment n, the total number of detected events, the CTR, and the number of maximum-likelihood expectation-maximization (MLEM) iterations. Performance is quantified by the normalized mean squared error of the reconstructed rate-constant map and by

the contrast recovery coefficient (CRC) of two regions of interest embedded in a uniform phantom background. The results establish a practical operating range for n, show that the method remains viable at the timing resolution of previous-generation scanners, and demonstrate robustness to variation in the direct-annihilation (DA) and parapositronium (p-Ps) lifetimes.

4.2 Simulation studies

Simulations were run in Python 3 using two parts: a Monte-Carlo data generation program previously reported by Huang *et al.* [81], and a reconstruction program developed to test the method described above. The data generator built a phantom consisting of an elliptical background containing two circular region-of-interest (ROI)s, shown in Fig. 4.1. The background ellipse had semi-axes of 6.5 cm and 3.3 cm along x and y. The two circular ROIs each had a radius of 1.4 cm and were centered on the $y = 0$ axis at $x = \pm 2.62$ cm (± 8 pixels at a pixel size of 0.327 cm), giving a center-to-center separation of 5.23 cm. The 41×41 image grid spanned a 13.4 cm field of view. The phantom simulated o-Ps rate constants of $0.4\,\mathrm{ns}^{-1}$ (right ROI), $0.5\,\mathrm{ns}^{-1}$ (phantom nonzero background), and $0.6\,\mathrm{ns}^{-1}$ (left ROI). The DA and p-Ps rate constants were held fixed at $2.5\,\mathrm{ns}^{-1}$ (giving DA = 400 ps instead of 388 ps in Chapter 3) and $8.0\,\mathrm{ns}^{-1}$ respectively, for the reasons given in Chapter 1. The fractional intensities for the DA, p-Ps, and o-Ps were 0.6, 0.1, and 0.3, respectively.

For each simulated triple coincidence, the elapsed time t_d was estimated from the measured photon detection time and positions using Eq. (2.20). The weighted-moment sinograms $\tilde{Y}^{(n)}$ and $\tilde{Y}^{(n+1)}$ were accumulated via Eq. (2.33). Each moment sinogram was reconstructed by MLEM on Eq. (2.13) using a Siddon-computed TOF system matrix H, and the o-Ps rate-constant map was obtained voxel-wise from the ratio of the two reconstructed moment images via Eqs. (2.9) and (2.15). Random coincidences were not included in the present proof-of-concept simulation, even though the subtraction term in Eq. (2.33) was calculated.

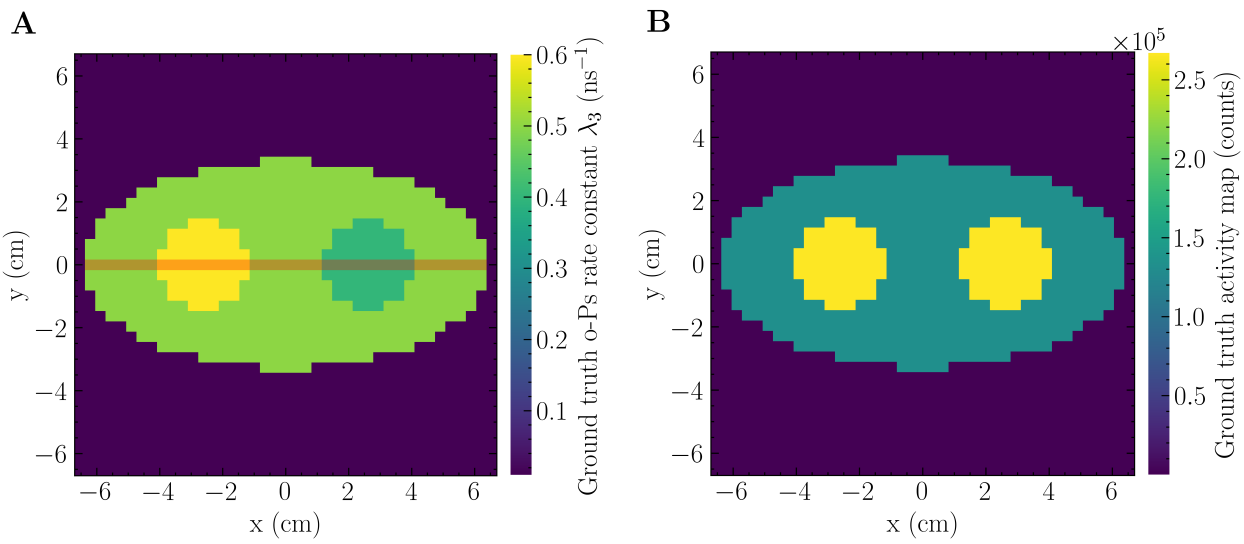

Figure 4.1: Ground truth maps used in this study. A) o-Ps rate constant map with a linear overlay (red) indicating the rate-constant profile taken for each simulated phantom. B) Ground truth activity distribution for the 1×10^8 total events case.

Evaluation of the method under realistic random fractions will be included in future work.

The total number of simulated events, the moment used, and the simulated CTR were varied across a parameter space consisting of $n \in \{0, 2, 4, \ldots, 40\}$, total events $\in \left\{10^5, 10^6, 10^7, 10^8\right\}$, CTR $\in \{200, 600, 3000\}$ ps, and MLEM iterations $m \in \{1, 2, \ldots, 40\}$. The CTR values of 200 ps, 600 ps, and 3000 ps were chosen to represent a state-of-the-art TOF PET system, a previous-generation TOF PET system, and a PET system that produces TOF measurements even though it is not considered TOF capable (referred to here as "non-TOF"), respectively.

For each swept parameter set, $K = 10$ simulation runs were conducted, differing only in their Poisson noise realization, and each was reconstructed for $M = 40$ MLEM iterations. Two things are derived from the resulting $K \times M$ images: an error metric that selects the operating iteration, and mean and standard deviation images at that iteration which separate the bias of the estimator from its run-to-run variability.

Error metric and iteration selection. Let $\hat{\lambda}_{3,ij}^{(k,m)}$ denote the estimated o-Ps rate constant at pixel (i, j) from run k at iteration m, and let $\lambda_{3,ij}$ denote the ground truth. The

normalized mean squared error (NMSE) of one reconstructed image is

$$\mathrm{NMSE}^{(k,m)} = \frac{\sum_{i=1}^{H}\sum_{j=1}^{W}\left(\hat{\lambda}_{3,ij}^{(k,m)} - \lambda_{3,ij}\right)^2}{\sum_{i=1}^{H}\sum_{j=1}^{W}\lambda_{3,ij}^2}, \tag{4.3}$$

where H and W are the image height and width in pixels. Averaging over the K runs gives a single value for each iteration,

$$\overline{\mathrm{NMSE}}^{(m)} = \frac{1}{K}\sum_{k=1}^{K}\mathrm{NMSE}^{(k,m)}. \tag{4.4}$$

Because MLEM amplifies noise with iteration, the reconstruction result is noted at the iteration

$$m^\star = \arg\min_m \ \overline{\mathrm{NMSE}}^{(m)}, \tag{4.5}$$

and $\overline{\mathrm{NMSE}}^{(m^\star)}$ is reported as the performance of that parameter set. Fig. 4.2 plots this minimum against the order of moment, and Fig. 4.3 shows the $\overline{\mathrm{NMSE}}^{(m)}$ curves from which each minimum was taken.

Sample mean and standard deviation images. All images reported below are taken at the selected iteration $m^\star$. Writing $\hat{\lambda}_{3,ij}^{(k)}$ for $\hat{\lambda}_{3,ij}^{(k,m^\star)}$, the voxel-wise mean and standard deviation across the K runs are

$$\bar{\lambda}_{3,ij} = \frac{1}{K}\sum_{k=1}^{K}\hat{\lambda}_{3,ij}^{(k)}, \tag{4.6}$$

and

$$s_{3,ij} = \sqrt{\frac{1}{K-1}\sum_{k=1}^{K}\left(\hat{\lambda}_{3,ij}^{(k)} - \bar{\lambda}_{3,ij}\right)^2}, \tag{4.7}$$

respectively. The mean image represents the expected value of the estimator at each pixel. The standard deviation image gives the run-to-run variability at the same pixel.

Profiles. A profile is taken along the horizontal line $y = 0\,\text{cm}$, drawn in red in Fig. 4.1a, which passes through the centres of both ROIs and the background on either side of them. For a mean image the profile is $\bar{\lambda}_{3,ij}$ along that line, plotted with a shaded band of $\pm s_{3,ij}$ read from the same pixels. These profiles are referred to as the "rate-constant profiles".

Single realizations. To show the effect of noise on one acquisition rather than in the mean, a single run is also selected arbitrarily from the same set of K for each scenario. A standard deviation across runs is not defined for one image, so these rate-constant profiles carry no band. The pixel-wise absolute deviation from ground truth, $|\hat{\lambda}_{3,ij} - \lambda_{3,ij}|$, is shown as a separate panel in its place.

4.3 Results

4.3.1 Minimum $\overline{NMSE}$

For each moment n, the minimum $\overline{\text{NMSE}}$ for all MLEM is shown in Fig. 4.2. Each point is the minimum of the $\overline{\text{NMSE}}$ over all MLEM iterations. The iteration at which this minimum occurs is not shown here. A broad trend is observed in which the minimum $\overline{\text{NMSE}}$ decreases with increasing event-counts. For any given event-count level, the minimum $\overline{\text{NMSE}}$ increased with the CTR, with the lowest values at CTR = 200 ps, followed by 600 ps, and the highest at 3000 ps. The exception is the result from the 1×10^5 counts-case in which the minimum $\overline{\text{NMSE}}$ for CTR = 3000 ps was lower than for the other two cases. Additionally, we re-emphasize that the presented method assumes a sufficiently large n and, accordingly, the results obtained at small n should be interpreted with caution.

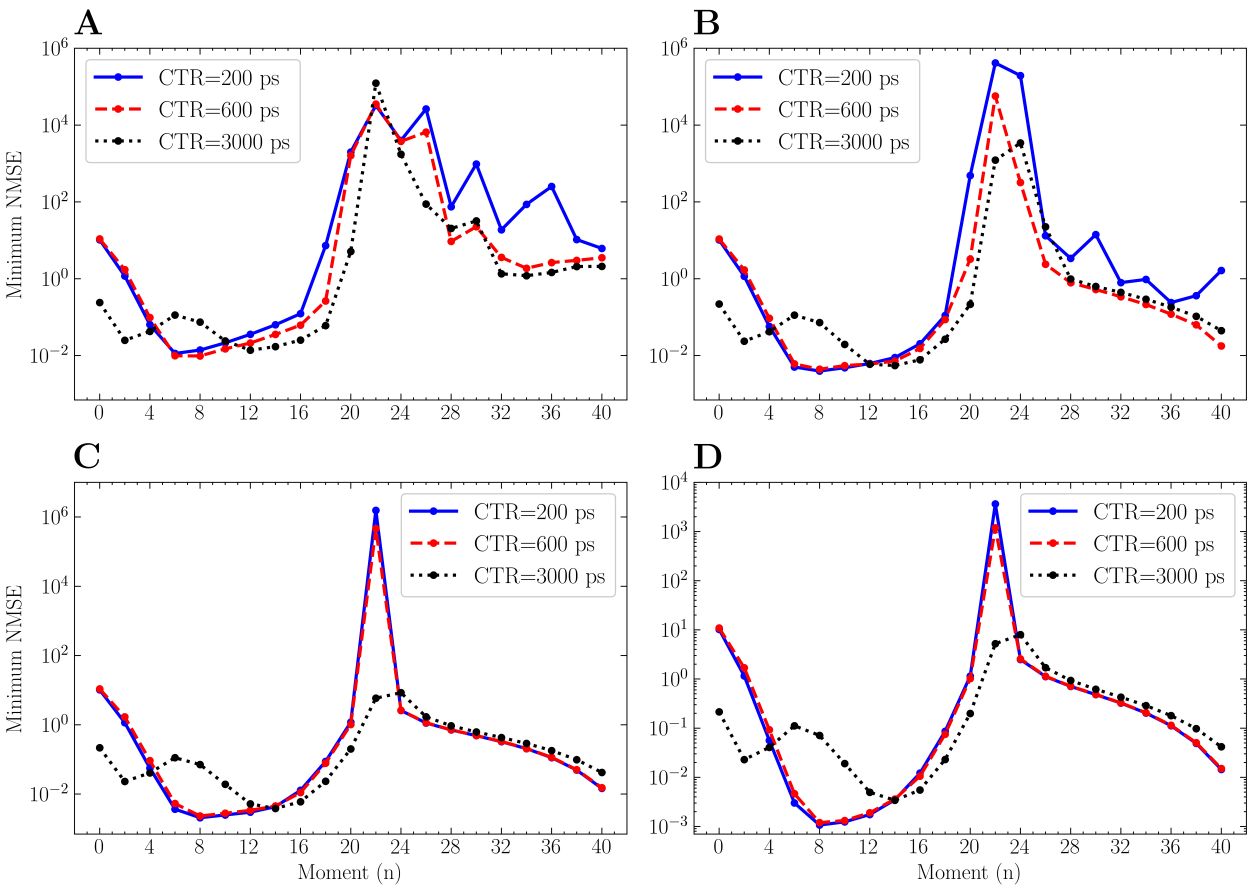

Figure 4.2: Mean minimum $\overline{\text{NMSE}}$ observed for each even moment for CTRs of 200 ps, 600 ps, and 3000 ps, for A) 1×10^5 counts, B) 1×10^6 counts, C) 1×10^7 counts, and D) 1×10^8 counts, across all MLEM iterations.

4.3.2 $\overline{NMSE}$ and MLEM iteration

In Fig. 4.3, the results shown are the mean $\overline{\text{NMSE}}$ which resulted in the minimum $\overline{\text{NMSE}}$ shown in Fig. 4.2. These results evaluate the effect of MLEM reconstruction iterations and provide an estimate for the number of iterations to be used for various CTRs and event-counts. Here it is observed that with the 1×10^5 counts and 1×10^6 counts cases the minimum $\overline{\text{NMSE}}$ occurs at $m = 1$, resulting from faster amplification of noise with the MLEM iteration at low event-count cases.

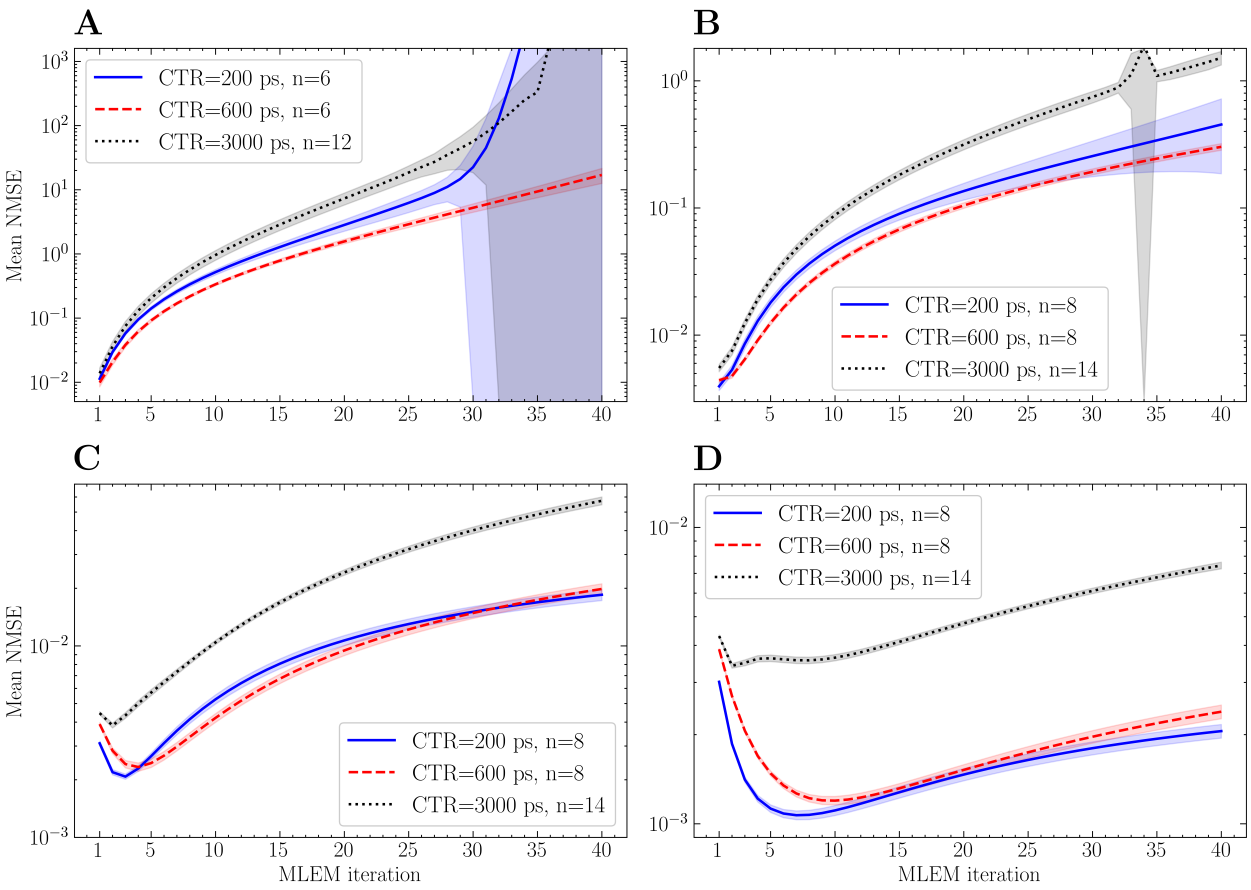

Figure 4.3: Specific cases of the data from Fig. 4.2. The plots are those where the minimum $\overline{\text{NMSE}}$ points for A) 1×10^5 counts, B) 1×10^6 counts, C) 1×10^7 counts, and D) 1×10^8 counts were observed. The shaded areas represent the ± 1 standard deviations.

4.3.3 CRC

The CRC results for the highest and lowest event-count cases are shown in Figs. 4.4 and 4.7, respectively. Each curve corresponds to an NMSE-minimizing moment. At 1×10^8 counts (Fig. 4.4), the CRC for all three cases of both ROIs rises with iteration (except for brief decrease for the 3000 ps case from iterations 6-10) and approaches unity, reaching roughly 0.95 to 0.97 by the final iteration and showing that the estimator recovers nearly the full contrast in the low-noise limit. The left ROI rises faster than the right for every CTR, and the two TOF-capable CTRs (200 ps and 600 ps) converge fastest, with the 200 ps case reaching the plateau within ~ 15 iterations, whereas the non-TOF case (3000 ps) rises more

slowly and reaches the same plateau only after roughly 35 to 40 iterations.

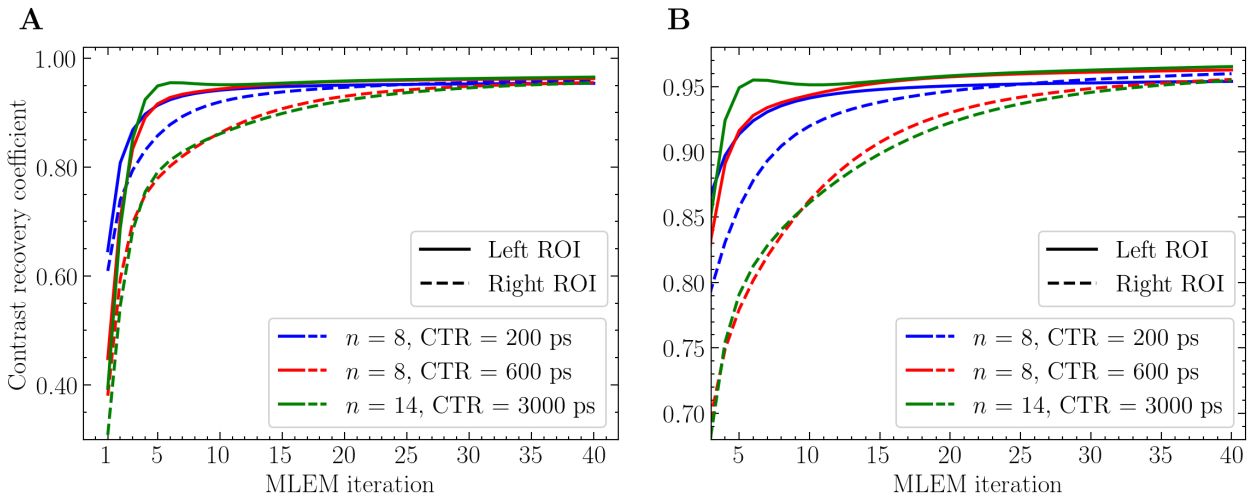

Figure 4.4: A) The CRC for 1×10^8 counts versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A)

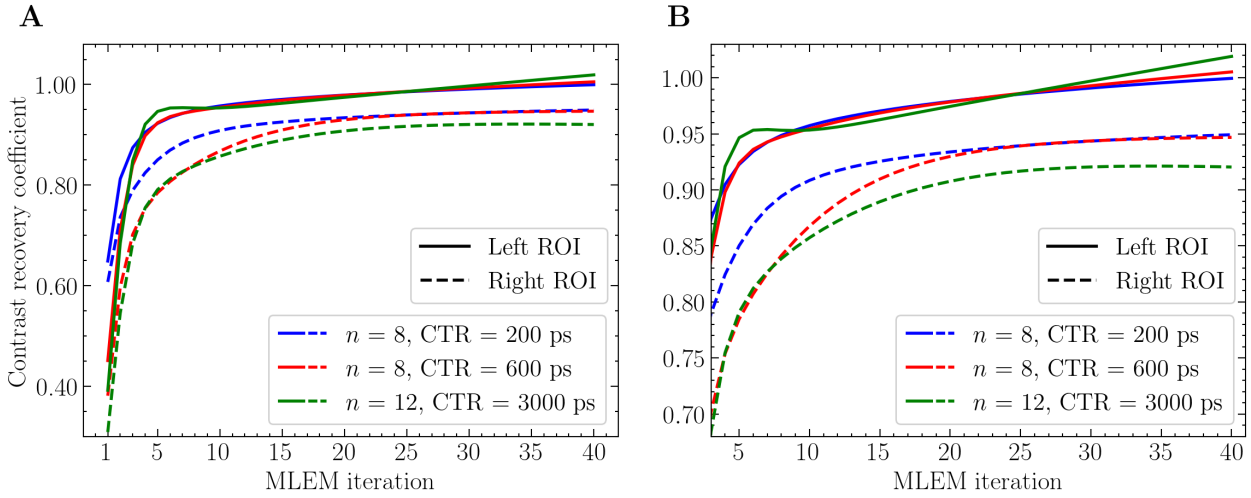

Figure 4.5: A) The CRC for 1×10^7 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A)

At 1×10^5 counts (Fig. 4.7), the curves instead exhibit noise-dominated behavior. The left ROI rises past unity by the fifth iteration for all three CTRs and continues to climb, while the right ROI reaches a local maximum well below unity (0.7147, 0.6572, and 0.6109 for CTR = 200 ps, 600 ps, and 3000 ps, at iterations 3 to 4) before then declining, and it never recovers the full contrast in the low-iteration regime. As shown in Section 4.3.2, the

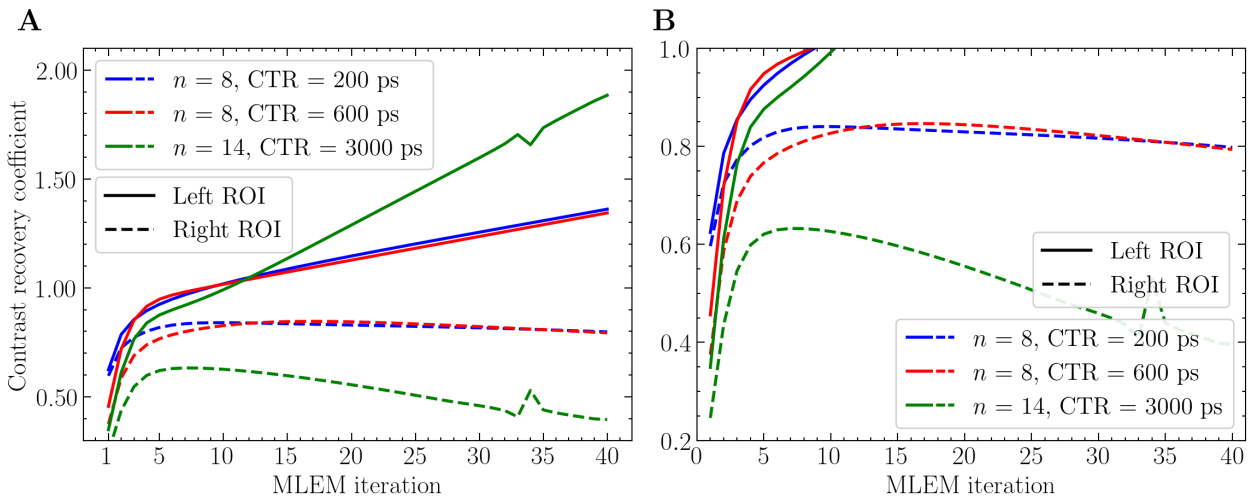

Figure 4.6: A) The CRC for 1×10^6 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A)

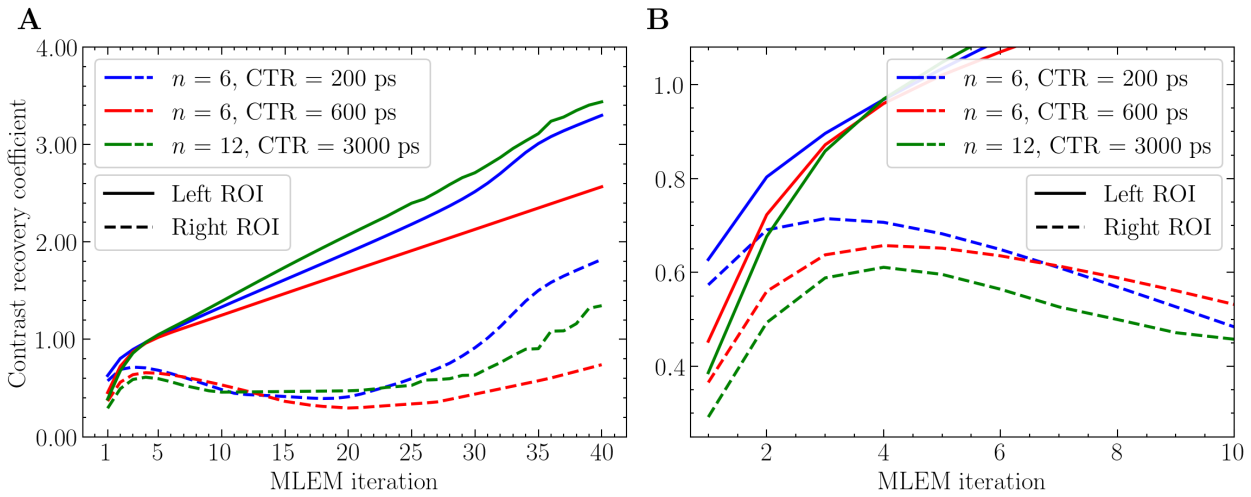

Figure 4.7: A) The CRC for 1×10^5 events versus the MLEM iteration for the three $\overline{\text{NMSE}}$-minimizing cases for each CTR. B) A zoomed-in plot of (A)

NMSE-minimizing iteration in this regime occurs very early, so the estimate must be done before noise dominates.

A CRC exceeding unity, seen here for the left ROI at low counts and at large iteration number generally, is a consequence of MLEM semiconvergence, the behavior of the reconstruction algorithm to initially show a decrease in reconstruction error followed by an increase in error as later iterations begin to fit statistical noise [82,83]. As iterations proceed, MLEM

amplifies high-frequency noise, which inflates the apparent ROI-to-background contrast beyond its true value. The effect is most pronounced for the left ROI, whose positive contrast overshoots, whereas the right ROI under-recovers in the early iterations where the estimate is taken. A CRC above one therefore reflects noise rather than improved contrast recovery, consistent with the early NMSE minima and the subsequent rise of NMSE in Fig. 4.3. Because all reported estimates are taken at the NMSE-minimizing iteration, this overshoot does not affect them. Additionally, the NMSE-minimizing moment does not necessarily correspond with the CRC-maximizing moment. It was observed that instead of $n = 14$, using $n = 12$ for CTR $= 3000$ ps resulted in a larger maximum CRC of approximately 0.96.

4.3.4 Mean and single-realization images

Using the $\overline{\text{NMSE}}$-minimizing parameters identified above, we present mean and single-realization images for three scenarios spanning the parameter sweep. For each scenario the mean and standard deviation images of Eq. (4.7) are shown together with the $y = 0$ cm rate-constant profile and its ± 1 standard deviation band. Because the mean images average over $K = 10$ realizations, some departure of the mean rate-constant profile from ground truth could be evidence of bias instead of noise, although the sample size is limited. A single realization drawn arbitrarily from the same set is shown alongside each, with the absolute deviation from ground truth $|\hat{\lambda}_3 - \lambda_3|$ in place of the standard deviation image.

We define the best case as the parameter set with the smallest minimum $\overline{\text{NMSE}}$ over the entire sweep, namely $n = 8$, a CTR of 200 ps, 1×10^8 counts, and 7 MLEM iterations. Fig. 4.8 shows the sample mean and standard deviation images and Fig. 4.9 a single realization.

The middle-ground case uses mid-tier parameters representative of a currently available clinical scanner, a CTR of 600 ps and 1×10^7 counts, for which the $\overline{\text{NMSE}}$-minimizing moment and iteration were $n = 8$ and 4. Fig. 4.10 shows the mean and standard deviation images and Fig. 4.11 a single realization.

The worst case is defined as the best-performing parameters at the extremes of the sweep, the lowest event count of 1×10^5 and the largest CTR of 3000 ps, which gave $n = 12$ and the first MLEM iteration. Fig. 4.14 shows the mean and standard deviation images and Fig. 4.15 a single realization.

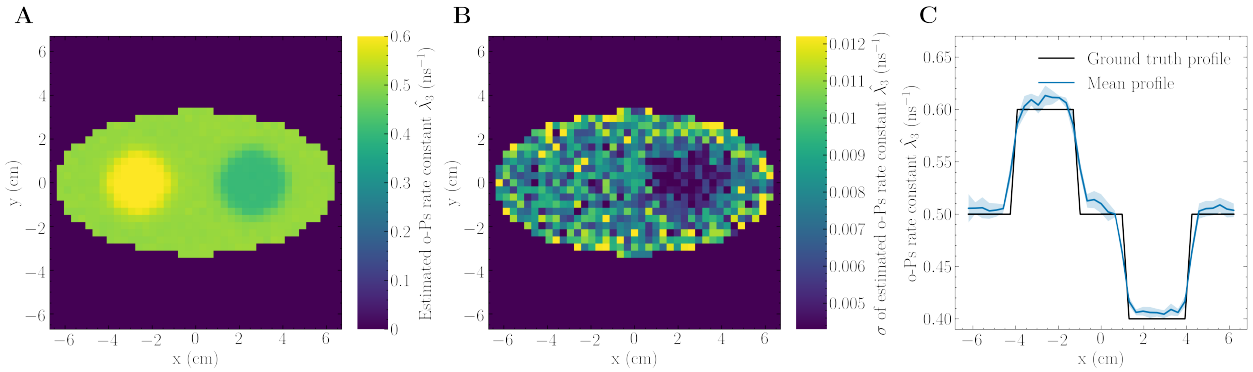

Figure 4.8: The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 200 ps, $n = 8$, 1×10^8 total counts, and taken from the 7th MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean.

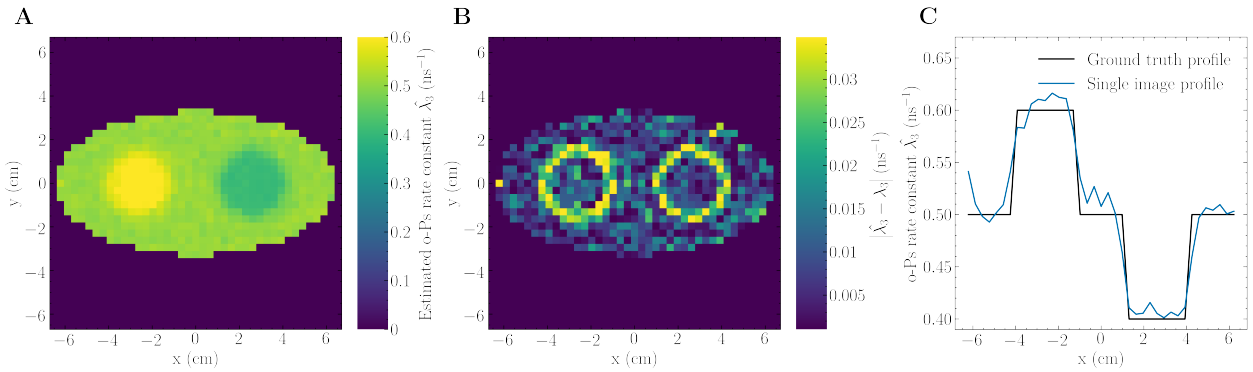

Figure 4.9: A) A single realization of an image that was used to produce the data in Fig. 4.8. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm.

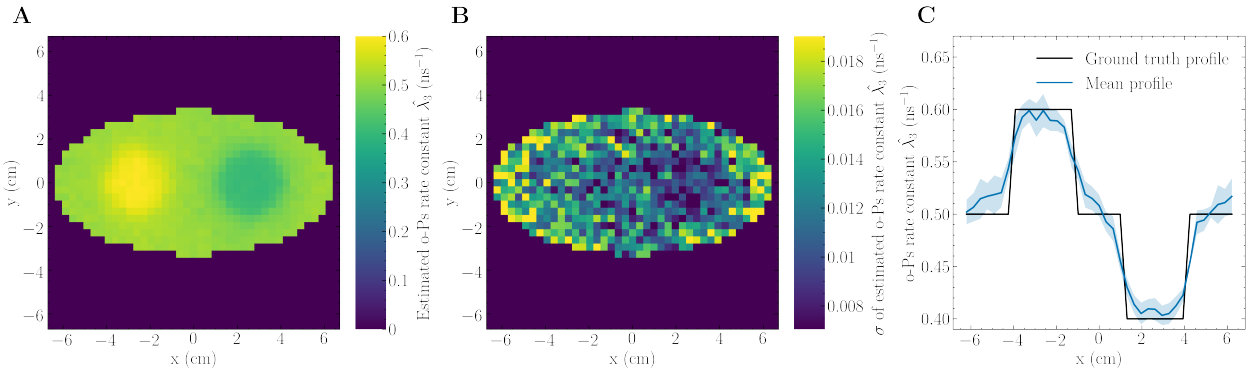

Figure 4.10: The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^7 total counts, and taken from the 4th MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean.

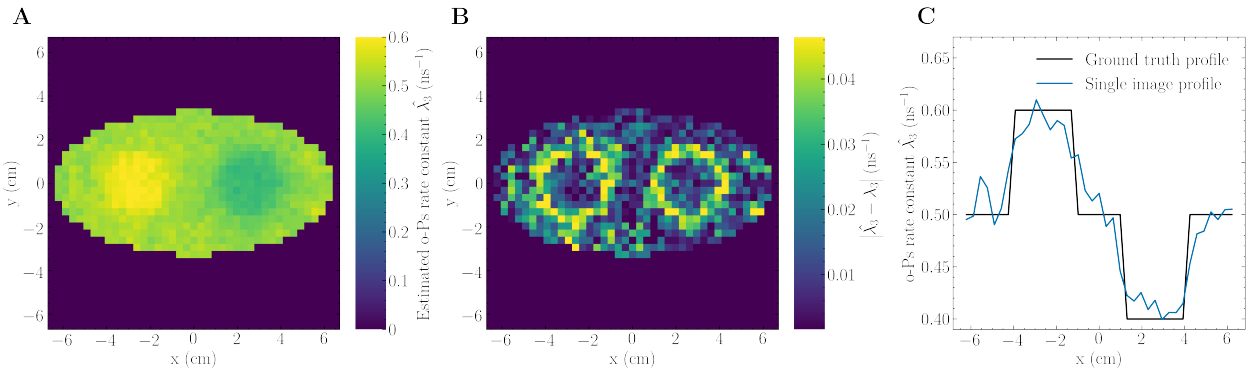

Figure 4.11: A) A single realization of an image that was used to produce the data in Fig. 4.10. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm.

4.3.5 Non-o-Ps component perturbations

To characterize the dependence of the method on the components it is designed to suppress, the DA and p-Ps lifetimes were perturbed randomly by ± 10 %, using the same perturbation for all voxels, and the reconstruction repeated. Fig. 4.16 shows the result for the middle-ground parameters of Fig. 4.11, that is a CTR of 600 ps, $n = 8$, 1×10^7 counts, and the 4th MLEM iteration.

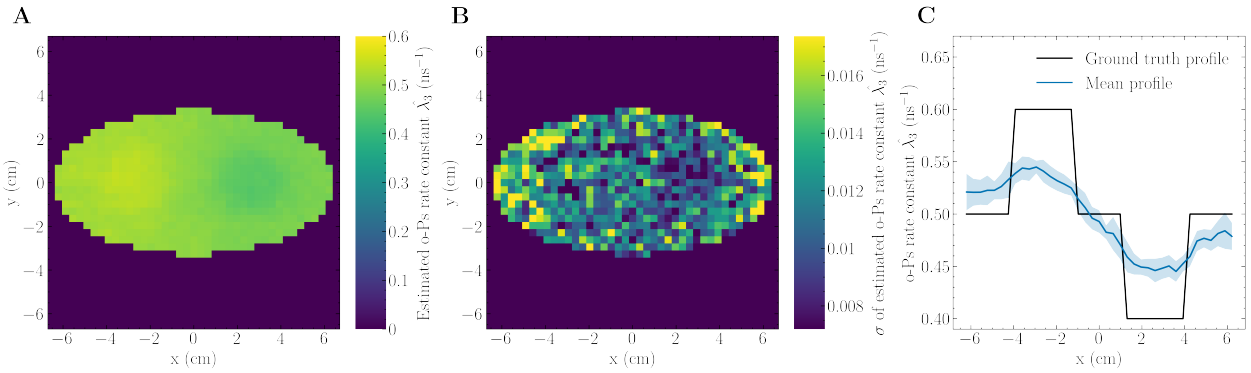

Figure 4.12: The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^6 total counts, and taken from the 1st MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean.

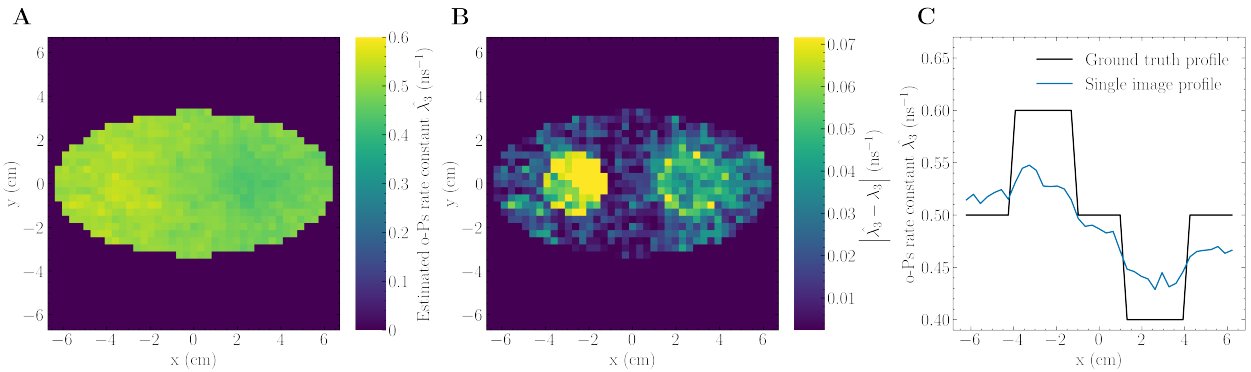

Figure 4.13: A) A single realization of an image that was used to produce the data in Fig. 4.12. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm.

4.4 Discussion

4.4.1 Influence of moment and CTR

The moment of choice (n) is the critical parameter for optimizing this method. In theory, larger n will selectively filter out the shorter-lived p-Ps and DA components, thereby improving the o-Ps lifetime estimation. However, increasing n also exacerbates the impact of noise as the signal-to-noise of the time-difference-histogram of prompt photon and coincidence

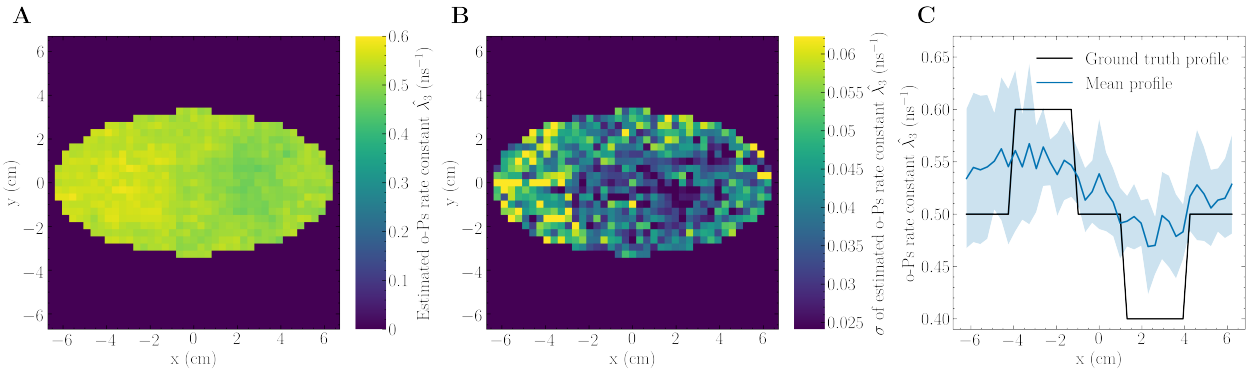

Figure 4.14: The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 3000 ps, $n = 12$, 1×10^5 total counts, and taken from the 1st MLEM iteration. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean.

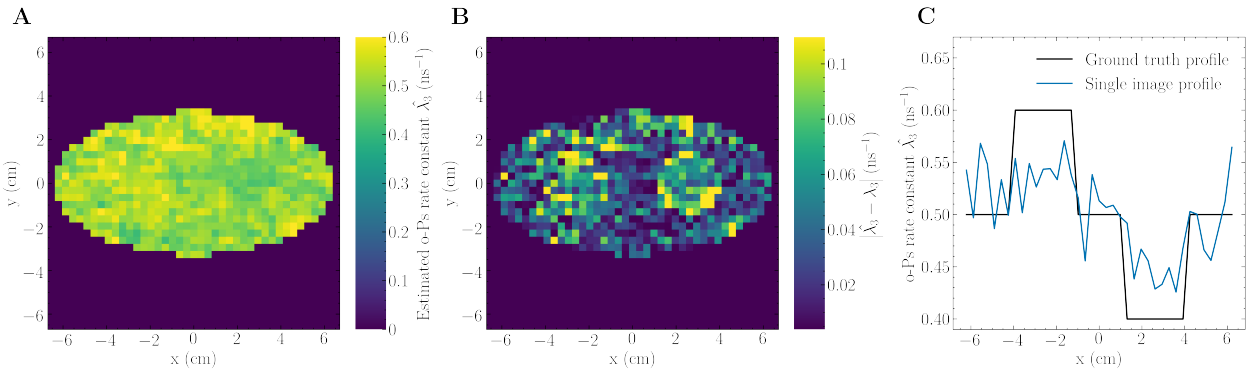

Figure 4.15: A) A single realization of an image that was used to produce the data in Fig. 4.14. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm.

photon arrival times degrades at large t_d. For the simulated phantom, the range of n from $8 \leq n \leq 16$ provided a sufficiently large moment to reduce bias while controlling the noise. The minimum $\overline{\text{NMSE}}$ occurred at CTR = 200 ps in all cases except 1×10^5 counts where large variances shifted it to 600 ps ($\overline{\text{NMSE}} = 9.8 \times 10^{-3}$ for 600 ps, $\overline{\text{NMSE}} = 1.1 \times 10^{-2}$ for 200 ps). Even though there are cases with sufficient ($> 1 \times 10^5$) counts where, for a given n, the larger CTRs produce smaller $\overline{\text{NMSE}}$, the $\overline{\text{NMSE}}$-minimizing n should be chosen for each event-counts case.

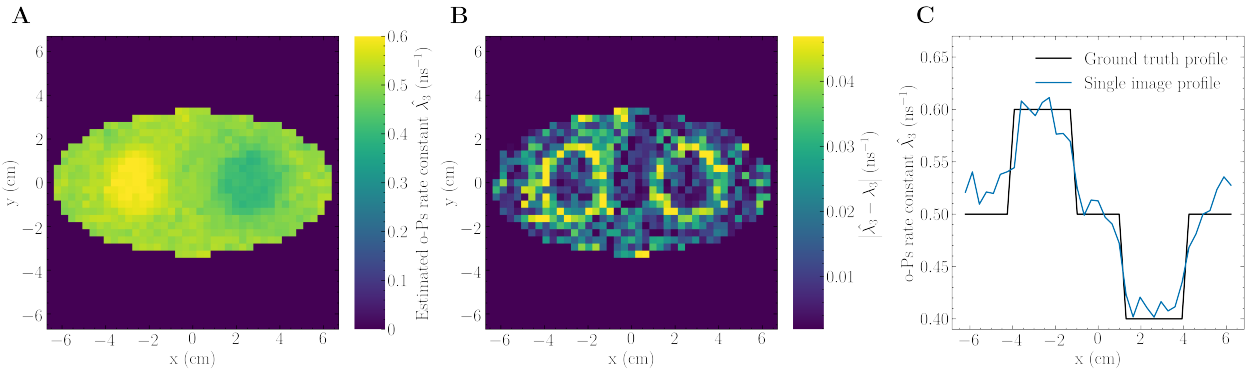

Figure 4.16: The same reconstruction data as Fig. 4.11, but with the p-Ps and DA lifetimes allowed to be perturbed randomly (same perturbation for all voxels) by ±10%. This data produced using CTR = 600 ps, $n = 8$, 1×10^7 counts, and taken from the 4th MLEM iteration. A) A single realization of the mean data. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0$ cm.

Two mechanisms degrade the estimate as n grows past its useful range. The first is truncation. The windowed moment of Eq. (2.33) weights each event by $(t_d + \sigma^2 s)^n e^{-st_d}$, which is maximized at $t_d = n/s - \sigma^2 s$. With $s = 1\,\text{ns}^{-1}$ the maximum remains near $t_d \approx n$ ns, so it lies outside the measurement window $\mathcal{T}$ once $n \gtrsim s\,T_2$. Since the $(n+1)$th moment is weighted more than the nth, truncation suppresses $f^{(n+1)}$ more than $f^{(n)}$, inflating the ratio $f^{(n)}/f^{(n+1)}$ and biasing $\hat{\lambda}_3$ upward, with the effect growing as n approaches $s\,T_2$.

The second mechanism involves the random-subtraction term. As n grows past the useful window, the $(t_d + \sigma^2 s)^n$ weighting amplifies the noise in the histogram tail faster than it suppresses non-o-Ps components, causing the moment images to become increasingly dominated by tail fluctuations. Due to the subtraction term in Eq. (2.33), the $(n+1)$th moment image is driven negative before the nth moment image, which remains predominantly positive. This sign mismatch causes the voxel-wise ratio in Eq. (2.15) to fluctuate erratically, leading to an unstable rate estimate and the observed peak in the minimum $\overline{\text{NMSE}}$ near $n \approx 22$ in Fig. 4.2. At even larger values of n, both moment images become negative and increasingly proportional, causing their ratio to converge toward an approximately uniform rate map. This map is close in magnitude to the mean o-Ps rate constant so the resulting

low $\overline{\text{NMSE}}$ is a consequence of the NMSE metric definition instead of a true improvement in the estimation.

The CTR is another important consideration. The method was shown to be viable for non-TOF capable systems (CTR = 3000 ps), with reasonable $\overline{\text{NMSE}}$ to TOF capable systems at CTRs of 200 ps and 600 ps. Even so, the temporal resolution gain given by TOF capability reduces the $\overline{\text{NMSE}}$ across most of the tested cases. Considering the histogram $S(t_d) = A \cdot p(t_d) + b$, a larger CTR will act as smoothing of the exponentially modified Gaussian (EMG) and a larger moment will be required to extract a similar contrast in signal when the ratio is taken. The cost of this is noise amplification. Additionally, from Fig. 4.3, there is evidence of the method struggling at low event-counts. The $\overline{\text{NMSE}}$-minimizing MLEM iteration for both 1×10^5 and 1×10^6 counts was the first iteration. Limited event rates are a known challenge for positronium lifetime imaging rather than a shortcoming specific to our reconstruction approach [84–87], due to the reliance on the relatively rare triple coincidence event. This scarcity means that when counts are low the data are dominated by Poisson noise, and MLEM noise amplifies with iteration [88], so additional iterations amplify that noise rather than resolve structure. The limitation is expected to ease with detector hardware, however. Long-axial-field-of-view and total-body PET systems provide substantially higher sensitivity for image-forming triple coincidences, with reported gains of $\sim$13.5$\times$ (LYSO) and $\sim$5.2$\times$ (plastic scintillator) relative to a conventional 20 cm axial-FOV scanner [85], and have already enabled *in vivo* positronium imaging in humans [65]. We note, however, that at clinical activity levels the dominant constraint is the random triple-coincidence fraction admitted by the wide prompt-γ timing window instead of sensitivity [87].

4.4.2 Bias and lookup table correction

The method exhibits a residual bias in most cases, arising from the trade-off between choosing a bias-reducing moment n and controlling noise amplification. Additionally, as can be seen in

Fig. 4.12, the $\overline{\text{NMSE}}$-minimizing case is not always the one with the most contrast recovery. There is a tradeoff also between $\overline{\text{NMSE}}$ and CRC, as seen when considering both Fig. 4.3 and Fig. 4.6. Therefore, the tradeoff extends to the choice of reducing bias in the estimate versus reducing noise in the image. For example, Figs. 4.17 and 4.18 show the mean and single-realization images for the same parameters as Figs. 4.12 and 4.13, but taken from the 4th MLEM iteration instead of the 1st. The mean image in Fig. 4.17 A shows a less biased estimate of the ground truth rate constant. However, this improvement comes at the cost of increased noise, with a mean standard deviation of $0.0634\,\text{ns}^{-1}$ for the 4th iteration compared to $0.0339\,\text{ns}^{-1}$ for the 1st iteration.

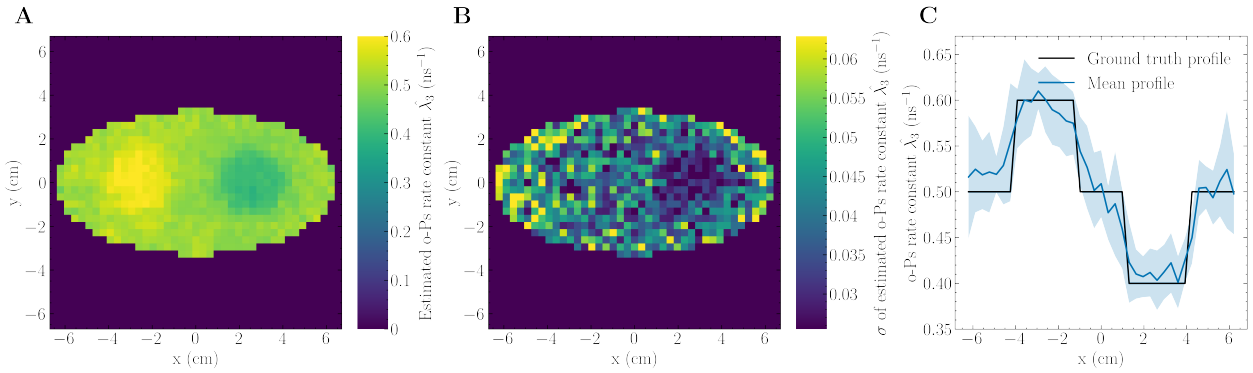

Figure 4.17: The voxel-wise mean and standard deviation calculated from 10 simulation runs which produced the minimum $\overline{\text{NMSE}}$ results from Fig. 4.2 for CTR = 600 ps, $n = 8$, 1×10^6 total counts, and taken from the 4th MLEM iteration, as opposed to the 1st like in Fig. 4.12. A) is the mean image, B) is the standard deviations image, and C) is the rate-constant profile taken for voxels along $y = 0$ cm. The solid blue curve in (C) is the mean and the shaded region is the ± 1 standard deviation about the mean.

A possible correction of the bias may be achieved by using a look-up table (LUT) to map the estimated rate constant to the true value. To test this, a LUT was constructed by simulating a calibration phantom with a ROI of the same dimensions as the ROI in the main phantom, except that it was placed in the middle of the phantom, as shown in Fig. 4.19 A. The background phantom lifetime was held fixed at $\lambda_3 = 0.5\,\text{ns}^{-1}$, and the activity map was uniform. We attempted to correct the 1×10^6 counts case with $n = 8$ at 600 ps and at the 1st MLEM iteration. The ROI rate constant was swept over 20 values

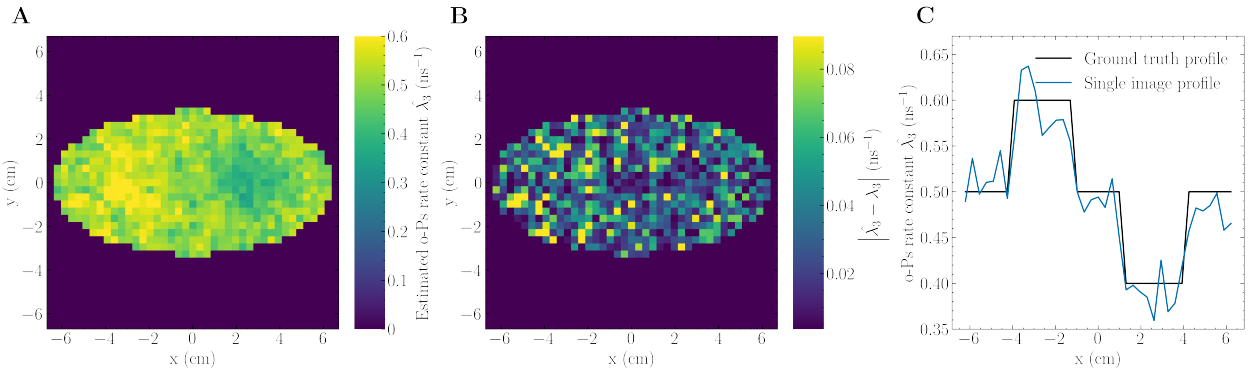

Figure 4.18: A) A single realization of an image that was used to produce the data in Fig. 4.17. B) The absolute difference between the estimate and true o-Ps rate constant, $|\hat{\lambda_3} - \lambda_3|$. C) The rate-constant profile taken for voxels along $y = 0\,\text{cm}$.

spanning $0.3\,\text{ns}^{-1}$ to $0.7\,\text{ns}^{-1}$ with 30 independent simulations at each value to produce a mean estimate for each. The resulting calibration curve is shown in Fig. 4.20 A. Although 20 rate constant values were simulated, one pair of adjacent reconstructed values was non-monotonic (indistinguishable) within the measured uncertainty and was therefore pooled by isotonic regression using pool-adjacent-violators algorithm (PAVA), which replaced the two values with their weighted average and produced a monotonic calibration curve with 19 unique LUT nodes.

Overall the LUT decreased the bias, as shown in Fig. 4.20 B. The LUT had a slope of 0.339 at iteration 1, corresponding to a gain of approximately 2.95. Leave-one-out (LOO) validation yielded a root mean squared (RMS) residual of $0.0085\,\text{ns}^{-1}$ for the calibration points. When applied to the main phantom, the LUT shifted the reconstructed ROI means from $0.454\,\text{ns}^{-1}$ to $0.379\,\text{ns}^{-1}$ for the $0.40\,\text{ns}^{-1}$ region, from $0.500\,\text{ns}^{-1}$ to $0.491\,\text{ns}^{-1}$ for the $0.50\,\text{ns}^{-1}$ region, and from $0.535\,\text{ns}^{-1}$ to $0.642\,\text{ns}^{-1}$ for the $0.60\,\text{ns}^{-1}$ region. However, the voxel-level standard deviations also increased following calibration.

Unfortunately the LUT is dependent on at least the reconstruction iteration and geometry of the phantom it is built from. This means that to theoretically build a practical LUT, many individual cases would need to be simulated. Other methods to correct the bias may possible, since an ideal estimator would combine the bias-reducing properties of a large

MLEM iteration with the noise-reducing properties of a small MLEM iteration, while being independent of the phantom and ROI geometries. This will be a topic for future work.

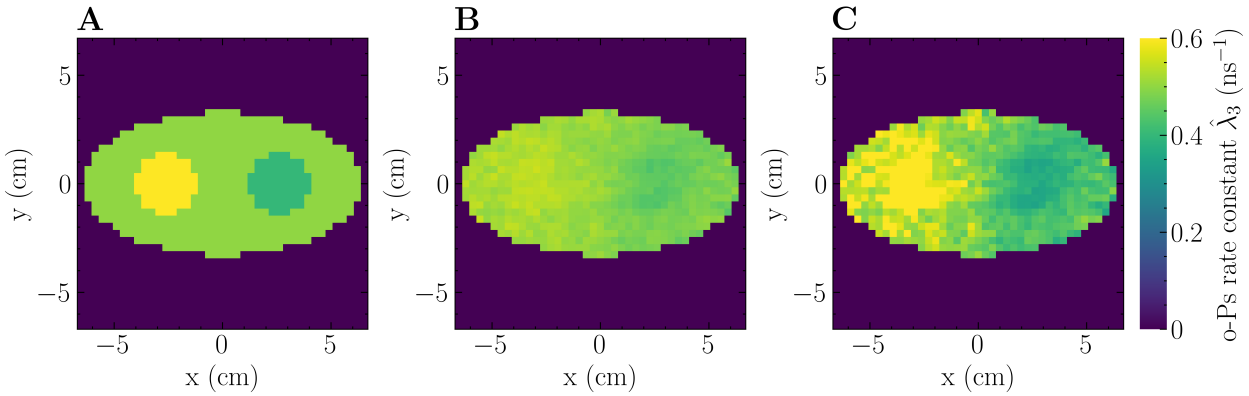

Figure 4.19: Ground-truth, uncorrected, and LUT-corrected o-Ps rate-constant maps for the structured phantom at the 1st MLEM iteration. A) Ground-truth map, B) uncorrected reconstructed map, and C) LUT-corrected reconstructed map.

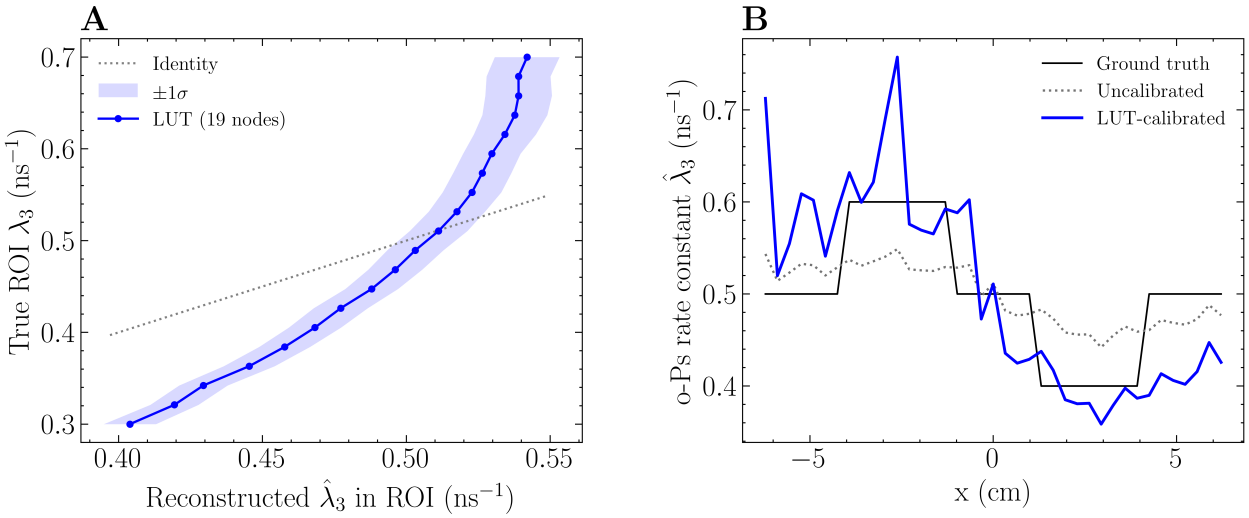

Figure 4.20: A) The LUT, shown via the true versus estimated o-Ps rate constant values inside the ROI. The shaded region denotes the ± 1 standard deviation about the mean. Although 20 nodes are plotted (one for each rate constant), 19 are used as two were degenerate. B) The resulting line profile from applying the LUT to the case from Fig. 4.12.

4.5 Conclusions

We have presented a fast, moment-based method for two-dimensional PET data. The results demonstrate that the proposed method provides a robust and efficient estimation for the o-Ps rate constant when a sufficient number of event-counts is exceeded. By utilizing the ratio of successive weighted moments, the long lived o-Ps lifetime is successfully isolated from the signal containing other components including the p-Ps and DA lifetimes. This is achieved using an analytic technique, making it computationally efficient to perform. After reconstruction of subsequent weighted moment images, taking the ratio immediately gives the rate constant without the need for voxel-level calculations. The range $n \in [8, 16]$ provided a practical operating range for the tested phantom geometries, with higher moments trading higher noise for lower bias.

Future work will investigate the optimization of the parameter s, extended phantom geometries, post-processing techniques such as Gaussian smoothing, and validation on empirical three-dimensional PET data. Current efforts in other areas of nuclear medicine will also help to enable PLT, such as deep learning-based denoising [89], which may be applicable here. As PET hardware continues to improve through investigations in full body PET, next-generation scintillators, and other means to maximize photon detection and minimize CTR, positronium lifetime imaging (PLI) capabilities and image quality will improve in parallel, supporting its eventual clinical translation.

Part II

An Analysis of Experimental Results for Positron Emission Tomography-Based Hypoxia Imaging

CHAPTER 5

ADDITION OF CONFORMAL RADIATION THERAPY TO A PRECLINICAL MULTI-MODAL IMAGING PROTOCOL

5.1 Introduction

Intensity-modulated radiation therapy (IMRT) was developed in 1982 [90], and has since become a vital external-beam radiation therapy (EBRT) modality and ubiquitous in the radiation oncologist's toolbox [91]. In our research group, we have developed a protocol to conduct oxygen-guided conformal radiation therapy (CRT) using conformal beams [24,25,92] on preclinical murine models. The eventual implementation of IMRT from the current oxygen-guided CRT experiments is the overall goal of these projects. CRT for small preclinical animal models has been enabled by the development of rapid 3D printed tungsten loaded blocks, which can deliver more conformal treatments to limit radiation to organs at risk [24, 93].

At its core, the CRT protocol uses multi-modal imaging to 1) produce a pO_2 image, and 2) contour the hypoxic region-of-interest (ROI) to deliver EBRT with conformal modulation of two opposing beams. This protocol can be extended or modified to include more imaging modalities, multiple mice in a single experimental day, and various CRT treatment planning options. In this chapter, we focus on the modified protocol developed specifically for this dissertation work, which forms the foundation for the data collection and analysis presented in the following chapters. This involved two types of hypoxia experiments. The first was a multi-modal CRT experiment in which the treatment planning was performed with either dynamic contrast-enhanced (DCE)-magnetic resonance imaging (MRI)-based correction of ^{18}F-fluoromisonidazole (FMISO)-positron emission tomography (PET), or a control. The second was an experiment which utilized only the multi-modality portion, along with the radioisotope ^{44g}Sc. Both of these experiments and their results will be detailed in Chapters 6, 7, and 8.

The overall experimental workflow is shown in Fig. 5.1. The contribution from this dissertation work is the development of a CRT treatment scheme using the two-animal multi-modal imaging protocol with the electron paramagnetic resonance (EPR)-PET dual scanner system.

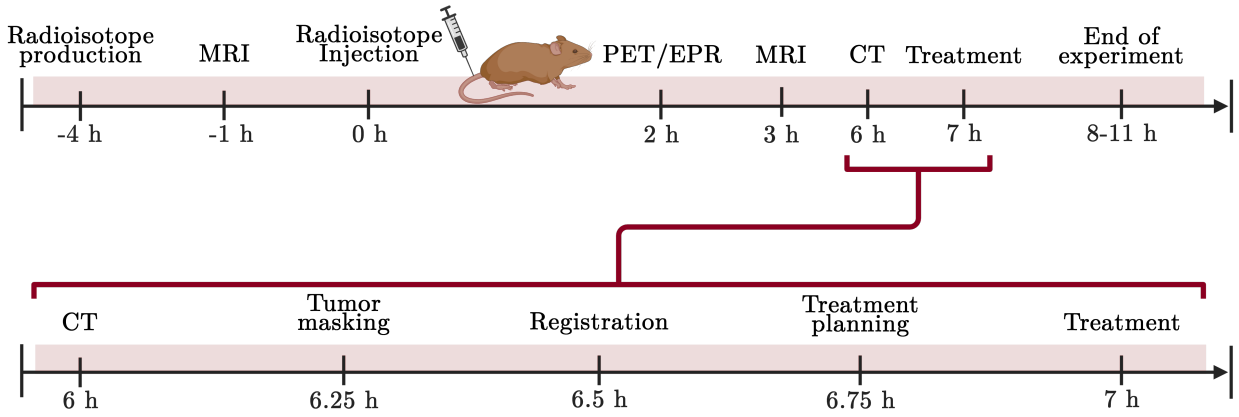

Figure 5.1: Timeline of a single imaging and treatment protocol for a general experiment. Times are given relative to FMISO injection at 0 h. Radioisotope production begins at −4 h and MRI is acquired at −1 h for tumor sizing. Following a 2 h tracer uptake period, PET and EPR imaging are performed, followed by a second MRI acquisition for tumor contouring and DCE. Next is treatment-planning computed tomography (CT) at 6 h and irradiation at 7 h. The lower axis expands the 6 h to 7 h planning-to-treatment window. After CT, the tumor is contoured (6.25 h), all modalities are registered to the common EPR-based fiducial space (6.5 h), the boost plan is generated (6.75 h), and treatment is delivered (7 h). The animal is maintained under anesthesia throughout. Times shown are approximate.

5.2 Radioisotope production

FMISO was produced on-site by the Cyclotron Facility at the University of Chicago, which houses an 18 MeV medical cyclotron (Ion Beam Applications, Louvain-la-Neuve, Belgium). Bombardment of a proton beam on an enriched ^{18}O target produces the ^{18}F atoms which are then transferred to shielded hot cells. The synthesis procedure used is well-known [94]. After FMISO synthesis, the radioactive solution is allowed to decay either for several hours if the

cyclotron production occurred in the morning, or overnight if it occurred the day previous to the animal portion of the protocol. Before distribution for injection, the FMISO is dosed to $\gtrsim 200\,\mathrm{mCi}$ or $\lesssim 1\,\mathrm{mL}$ (whichever has less volume to minimize shock to the animal), placed in a leaded container, and transported to the animal preparation station.

5.3 Animal preparation and minimization of distress

The entirety of the protocol, including all imaging and radiotherapy, is performed while the animals are under general anesthesia. Isoflurane is used for induction and maintenance of unconsciousness. A soft, rubber-like dental molding material (vinyl polysiloxane) is used to restrain the tumor-bearing leg in a bed to minimize movement. The bed contains fiducials for use in image registration. Breathing rate and rectal temperature are monitored closely and recorded every five minutes to track both immediate changes in physiological state and changes that may occur over hours. Visual confirmation of breathing is also done regularly, especially when transporting the animals from one station to another. All animal procedures follow the recommendations from the University of Chicago Institutional Animal Care and Use Committee (IACUC).

Euthanasia is conducted when needed using two methods to verify death. In our case, we use isoflurane overdose followed by cervical dislocation to minimize pain. This is consistent with the recommendations of the Panel on Euthanasia of the American Veterinary Medical Association and the University of Chicago IACUC.

5.4 Preclinical MRI

The preclinical MRI used in these studies is a 9.4 T, 30 cm bore diameter preclinical MRI scanner which is capable of T1-weigted (T1W), T2-weigted (T2W), DCE, diffusion weighted imaging, and other parametric imaging modalities with and without contrast (Bruker, Et-

tlingen, Germany). It is housed in the Integrated Small Animal Imaging Research Resource (iSAIRR) MRI facility. This is approximately 5 minutes away from the EPR-PET combined scanner and animal preparation areas by walking, which is important as the animals are maintained on anesthesia during this transportation using a mobile isoflurane cart. Mobile temperature and breathing rate monitoring probes, anesthesia flow controls, and a heated pad are transported with the cart.

The resolution of the in-plane T2W images that are typically used in the protocol is approximately 0.1 mm with a 0.75 mm slice thickness. These images are utilized for image registration, anatomical segmentation, and measuring the tumor volume prior to choosing the animal candidates. The tumor, leg, and muscle regions in the tumor-bearing leg are segmented manually using the T2W images. The embedded fiducial markers are then used to register all modalities.

Fig. 5.1 shows MRI occurring before the radioisotope injection. Although the timeline shows MRI before injection, MRI is often repeated after EPR-PET to attain DCE images and more accurate tumor contours. The T-1 hour MRI scan is mainly for choosing, based on tumor size, which two of the four candidate mice will be used on the given experiment day.

The DCE sequence is acquired to characterize tumor perfusion and vascular permeability. A bolus of gadolinium contrast agent is injected intravenously while a series of T1W images is acquired continuously, capturing the wash-in and wash-out of contrast in each voxel. The resulting signal-time curves are converted to contrast-agent concentration and fit using the Tofts two-compartment pharmacokinetic model [95, 96], which describes the exchange of contrast between blood plasma and the extravascular extracellular space,

$$C_t(t) = K^{\text{trans}} \int_0^t C_p(\tau)\, e^{-(K^{\text{trans}}/v_e)(t-\tau)}\, \mathrm{d}\tau, \tag{5.1}$$

where C_t is the tissue concentration, C_p is the arterial plasma concentration (the arterial

input function), K^{trans} is the volume transfer constant between plasma and the extravascular extracellular space, and v_e is the fractional volume of that space. Fitting Eq. (5.1) voxel-wise yields parametric maps of K^{trans} and v_e. K^{trans} reflects the combined effects of blood flow and capillary permeability and is elevated where vasculature is leaky, while v_e measures the volume fraction accessible to the contrast agent. Because poorly perfused regions deliver both oxygen and tracer inefficiently, these maps may provide a perfusion context for correcting the FMISO images [7].

5.5 Tumor growth

fibrosarcoma (FSa) cells are thawed and injected into the flanks of a C3H donor mouse and grown to approximately 1 cm. Viable tissue is separated from necrotic debris and minced. The resulting brei is added to 0.025 % trypsin in 80 mL of phosphate-buffered saline containing 0.01 % DNase. The suspension is stirred for 20 min with a magnetic stirrer, after which the undigested tissue is allowed to settle and the supernatant is collected and washed by centrifugation at 1200 rpm for approximately 5 min. The pellet is resuspended in Dulbecco's Modified Eagle Medium with 10 % fetal bovine serum to a concentration of up to 40×10^6 cells/mL. For tumor induction, 0.1×10^6 cells to 4.0×10^6 cells in 10 µL to 50 µL of suspension are injected intramuscularly into the leg, and tumors are grown to approximately 1 cm in length before the experiment protocol day.

5.6 Injections

All injections to the mice are done via a tail vein cannula, which is prepared prior to the beginning of the protocol. Ground truth pO_2 is acquired with electron paramagnetic resonance imaging (EPRI) using an injected triarylmethyl (trityl) spin probe, whose single unpaired electron produces a narrow resonance and slow electron-spin relaxation that the EPR imager

can measure [34, 97]. Molecular oxygen itself relaxes too rapidly to be imaged directly. Instead, it is sensed through Heisenberg spin exchange with the probe, whereby collisions with O_2 increase the trityl relaxation rate in linear proportion to the local oxygen concentration, so that the measured relaxation rate maps directly to absolute pO_2 [34, 98].

Oxygen imaging used the triaryl methyl (trityl) radical OX071, a partially deuterated tris(8-carboxyl-2,2,6,6-tetra(2-(1-hydroxy-2,2-d_2-ethyl))benzo[1,2-d:4,5-d′]bis(1,3)dithiol-4-yl)-methyl sodium salt (molecular weight 1451 g/mol), obtained from GE Healthcare (London, UK) and dissolved in trisodium citrate [99]. The probe is well tolerated at imaging doses and is administered through a tail-vein cannula. Because trityl clears from the tumor over roughly 20 min to 30 min, a maintenance infusion is used to hold the concentration steady through the EPR acquisition. Each mouse receives an initial 70 µL bolus followed by a continuous infusion over 30 min, for a total dose of approximately 70 µmol/kg.

For the radioisotopes (either ^{18}F in FMISO or ^{44g}Sc), the dose is usually maximized since the PET acquisition is usually limited to 30 minutes. The same procedure is used regardless of isotope. The dose is injected via the tail vein cannula 2 hours before PET imaging. This allows the compound to biodistribute through the mice and into the tumor. ^{44g}Sc has a half-life of 4.04 h while ^{18}F is 1.83 h. These are relatively short-lived and allow for the mice and working areas (in the case of spills or contamination) to be near-background levels within 1-3 days, assuming there are no longer-lived species from the radioisotope production.

5.7 Preclinical EPR-PET combined scanner

The combined EPR-PET scanner was developed jointly by the Halpern and Kao laboratories together with The University of Chicago Cyclotron Group [100–102], and is shown in Fig. 5.2. A preclinical small-animal PET scanner is mounted inside a EPR imaging system so that the same animal can be imaged by both modalities simultaneously without repositioning. In this work it is used to acquire co-registered PET and oxygen EPR images of tumors in mice.

The EPR subsystem is a pulsed imager operating at 715 MHz to 730 MHz, with a static field supplied by a 27 mT permanent magnet together with gradient and offset coils for spatial encoding [100]. Oxygen imaging uses the injectable trityl spin probe [99]. Under these conditions the subsystem resolves oxygen to approximately 1.0 mmHg at low oxygen concentrations. The PET subsystem, characterized previously [102], has a spatial resolution of approximately 1.6 mm for FMISO imaging.

Imaging follows a fixed animal-handling protocol. Each mouse is anesthetized and cannulated through the tail vein for tracer and spin-probe administration, then placed on a bed with the tumor-bearing leg immobilized against built-in fiducial markers. The markers are visible on T2W MRI, EPR, and CT, and a fixed transformation derived from them registers the PET and EPR volumes into a common frame.

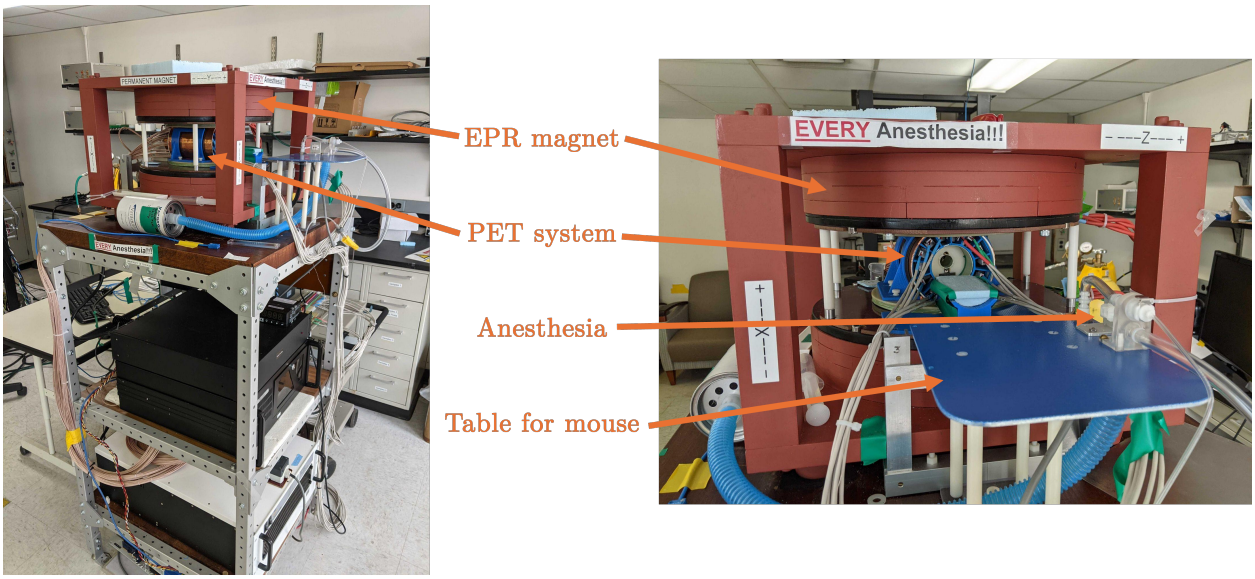

Figure 5.2: The combined EPR-PET small-animal scanner. A preclinical PET insert is mounted coaxially within the bore of the EPR imaging magnet, so that FMISO PET and oxygen EPR images can be acquired simultaneously and without repositioning the animal. Left: the full system on its mobile rack. Right: detail of the imaging bore showing the EPR magnet, the PET seated inside it, the mouse table, and the isoflurane anesthesia line.

5.8 Preclinical CT and CRT

Irradiation follows the 3D-printed aperture (block)-based preclinical CRT methodology of [103], which focuses a prism of radiation on hypoxic tumor regions using the X-RAD 225Cx orthovoltage irradiator (Precision X-Ray, North Branford, CT) used in these studies. Rather than a multileaf collimator, intensity modulation is achieved with 3D-printed compensators of spatially varying thickness. They are composed of tungsten-doped polylactic acid plastic. Two equally spaced coplanar beams are inverse-planned, using the CT function of the X-RAD 225Cx for final registration and alignment. Two compensators are printed (one for each opposed beam angle) with consumer-grade 3D printers. They are cylindrical plugs that seat in a collimator mount in the beam output.

Following [24], the base prescription is set to the previously reported tumor control dose 15% (TCD15) for this tumor model, 22.5 Gy, delivered whole-field to the tumor, while protocol-defined subvolumes (identified on EPR, corrected FMISO, or uncorrected FMISO, depending on the protocol) are boosted to the tumor control dose 95% (TCD95), 35.5 Gy.

5.9 Conclusion of experimental protocol

The protocol yields a set of co-registered images of the tumor-bearing leg, examples of which are shown in Fig. 5.3. The T2W MRI and CT provide the anatomical references. The tumor is contoured on the T2W image (magenta). The EPR image supplies a voxel-wise map of oxygen partial pressure. The FMISO PET image, shown here as tumor-to-muscle ratio (TMR), reports tracer retention over the same volume, and the DCE-derived K^{trans} and v_e maps characterize vascular permeability and extravascular extracellular volume, respectively. Together these provide, within a single registered frame, an oxygen reference standard from EPR, the tracer distribution to be evaluated, and the perfusion information for correction.

Depending on the specific protocol and experiment, subsets of these images may be

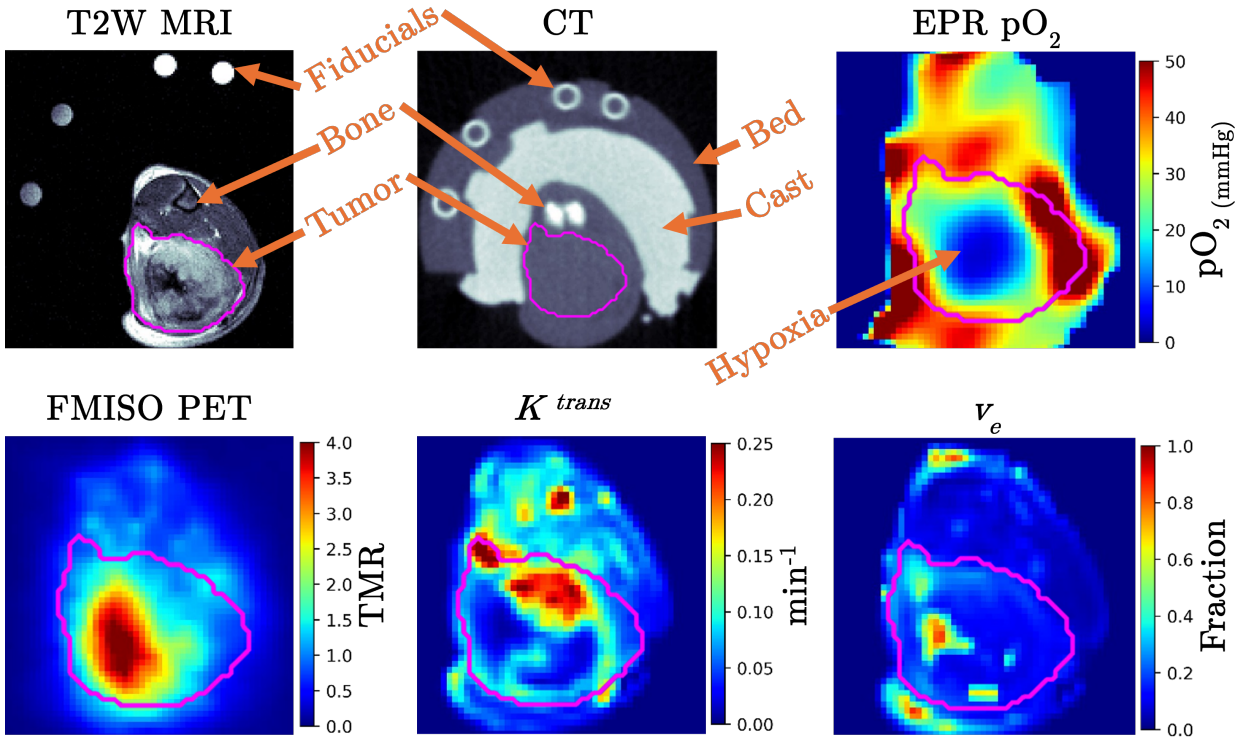

Figure 5.3: Representative co-registered images of a tumor-bearing leg. All panels show the same axial slice, with the tumor contour (magenta) drawn on the T2W MRI and propagated to every modality through the fiducial-based registration. Top row: T2W MRI and CT provide anatomy and landmarks, with fiducial markers, bone, bed, and immobilizing cast labeled. The EPR map gives oxygen partial pressure in mmHg, with the hypoxic tumor core appearing as a dark-blue low-pO_2 region. Bottom row: FMISO PET shown as tumor-to-muscle ratio (TMR), and the DCE-derived K^{trans} (min^{-1}) and v_e (fractional) maps. Elevated FMISO uptake spatially corresponds to the EPR-defined hypoxic region in this example.

acquired and used for treatment planning. When treatments are done using FMISO, however, the treatment planning is performed completely blinded to the EPR images. The EPR fiducials serve as the reference datums during the image registration, however the oxygen images are not used until the post-experiment analysis.

CHAPTER 6

PRELIMINARY SURVIVAL STUDY OF DCE-MRI-BASED CRT CORRECTION IN A PRECLINICAL FIBROSARCOMA MODEL

6.1 Introduction

In a previous study by Gertsenshteyn *et al.*, [104], the optimal threshold to define tumor hypoxia using positron emission tomography (PET) ^{18}F-fluoromisonidazole (FMISO) was identified for preclinical squamous cell carcinoma VII (SCCVII). Later in the dissertation work by Gertsenshteyn [105], another threshold was found for preclinical fibrosarcoma (FSa) models. The FSa threshold was identified to be a tumor-to-muscle ratio (TMR) $\geq$ 1.4. In a different study, the theoretical correction of FMISO images using dynamic contrast-enhanced (DCE) magnetic resonance imaging (MRI) was explored [106]. In this chapter, we detail the work that has been conducted since, keeping in mind Chapter 5 which described the experimental protocol and associated logistics that this chapter's work has employed.

Prior work established that guiding radiation boosts to hypoxic tumor regions identified by electron paramagnetic resonance (EPR) partial pressure of oxygen (pO_2) imaging improves local tumor control in preclinical FSa [24]. Unfortunately, EPR is not yet clinically available [107]. In an effort to work towards oxygen-guided radiation therapy, we have begun conformal radiation therapy (CRT) experiments on preclinical FSa models using "uncorrected" and "corrected" FMISO data. The uncorrected experiment is simply the CRT plan detailed in Chapter 5 that gives radiation boosts to tumor regions of TMR $\geq$ 1.4, while the corrected experiment is the same threshold but applied to DCE-corrected FMISO images.

6.2 Brief reintroduction of DCE

Detailed in Chapter 5, Section 5.4, DCE is a dynamic MRI technique in which images are acquired repeatedly before, during, and after the intravenous injection of a gadolinium-based contrast agent. This allows for tracking the contrast agent's wash-in and wash-out through the tissue [108]. Fitting the resulting contrast-time curves to a pharmacokinetic model gives voxel-wise estimates of tissue perfusion and microvascular permeability. Specifically, we emply the volume transfer constant K^{trans} between the blood plasma and the extravascular extracellular space and the extravascular extracellular volume fraction v_e. In the present work these parameters serve to correct the FMISO signal for the perfusion effects that otherwise confound its interpretation as a marker of hypoxia.

6.3 DCE correction of FMISO images

Low FMISO uptake is ambiguous. A tumor voxel can retain little tracer either because it is well oxygenated or because poor perfusion limits how much tracer reaches it [109]. Additionally, FMISO can under-report hypoxia in intermediate and acutely hypoxic cells [110]. A threshold applied to the raw FMISO image can therefore place the hypoxic boundary away from what would be the true pO_2-defined threshold. Using pO_2 attained from EPR imaging as the reference standard [7, 111], Gertsenshteyn *et al.* found only modest spatial agreement between the hypoxic regions of raw FMISO PET and electron paramagnetic resonance oxygen imaging (EPROI), with a baseline Dice similarity coefficient (DSC) near 0.42 to 0.44 across imaging groups [105].

The correction is trained against a target built from the EPROI image. A logistic function models FMISO retention as a function of pO_2, with its inflection at the hypoxia threshold [106]. We refer to this EPR-derived FMISO image as PET^{M}.

The corrected image is formed by fitting the two DCE maps, K^{trans} and v_e [96, 112]

(Chapter 5), to the difference between this target and the measured FMISO image by linear least squares. The fitted weights are then used to modify the DCE maps, which are added back to the measured PET image. During usage, we modify a given PET into a "corrected" PET map,

$$\widehat{\mathrm{PET}} = \mathrm{PET} + \hat{\beta}_{K^{\mathrm{trans}}} K^{\mathrm{trans}} + \hat{\beta}_{v_e} v_e, \tag{6.1}$$

where $\hat{\beta}_{K^{\mathrm{trans}}}$ and $\hat{\beta}_{v_e}$ weight the K^{trans} and v_e maps. The correction raised the DSC and lowered the 95th percentile Hausdorff distance against EPROI [105].

In the present study the correction is applied to the FSa FMISO images using the group mean weighting constants reported for mammary carcinoma (MCa4) in the proof-of-concept study [106], namely $\hat{\beta}_{K^{\mathrm{trans}}} = 8.62$ and $\hat{\beta}_{v_e} = -2.52$. A single corrected image is formed for each tumor from Equation 6.1 with these fixed constants, and the CRT experiment with the corrected protocol then applies a boost to the regions of TMR ≥ 1.4 measured on $\widehat{\mathrm{PET}}_M$. Because these constants were trained on MCa4 mammary carcinoma rather than FSa, applying them here assumes that the perfusion-to-hypoxia relationship is applicable to different tumor lines.

6.4 Results

6.4.1 Cohort and baseline characteristics

A total of 26 tumors were analyzed, randomly assigned to 12 in the corrected arm and 14 in the uncorrected arm. Baseline hypoxia metrics were available for 25 of these since one uncorrected animal had no covariate data due to no presence of hypoxia. The endpoint was recurrence free survival over a 90 day window, with a recurrence scored as an event and controlled or censored animals treated as censored at their last follow up. Five recurrences were observed. Baseline hypoxia metrics for each arm are summarized in Table 6.2 and defined in Table 6.1. The two arms were well-matched on the metrics Hsurf, HsurfNorm,

Table 6.1: Metrics used in this study for correlations to hypoxia.

Metric	Definition
HRegions	Number of distinct hypoxic regions
Hsurf	Total surface area of hypoxic regions (cm^2)
HsurfNorm	Ratio of hypoxic surface area to that of a sphere of equal volume
TumorRate	Tumor growth rate over three days leading to experiment (µL/day)
HF10	EPR-defined hypoxia (≤ 10 mmHg)
HVolume	Total volume of hypoxic regions (mL)

Table 6.2: Baseline hypoxia metrics by protocol. Continuous metrics are mean ± SD and are compared with Welch's two-sample t-test. The P column compares the corrected and uncorrected arms. Covariate data were unavailable for one uncorrected animal and therefore the first six rows are computed on 12 corrected and 13 uncorrected tumors while the recurrence row uses all 26 animals.

Metric	Protocol		P
	Corrected	Uncorrected	
HRegions	2.5 ± 2.9	7.4 ± 6.3	0.022
Hsurf (cm^2)	0.62 ± 0.55	0.78 ± 0.63	0.482
HsurfNorm	1.14 ± 0.59	1.33 ± 0.75	0.493
TumorRate (µL/day)	91 ± 31	70 ± 24	0.067
HF10	0.089 ± 0.091	0.105 ± 0.087	0.654
HVolume (mL)	0.033 ± 0.037	0.036 ± 0.034	0.861
Recurrence, n (%)	2 (16.7%)	3 (21.4%)	1.000

HF10, and HVolume (all $P > 0.4$), but differed in the number of hypoxic regions, which was lower in the corrected arm (2.5 for corrected versus 7.4 for uncorrected, $P = 0.02$). Tumor growth rate showed a nonsignificant trend (91.4 for corrected versus 70.1 for uncorrected, $P = 0.07$). The recurrence rate did not differ between arms (16.7% versus 21.4%, $P = 1.0$).

6.4.2 Survival by protocol

Recurrence free survival did not differ between the corrected and uncorrected arms. The Kaplan-Meier curves are shown in Fig. 6.1. They overlap throughout the 90 day window, with heavily overlapping confidence intervals. A log-rank test found no significant difference

between the arms (test statistic 0.10, $P = 0.76$). By day 90 the estimated recurrence free probability was similar in the two arms, near 0.83 for the corrected arm and 0.77 for the uncorrected arm.

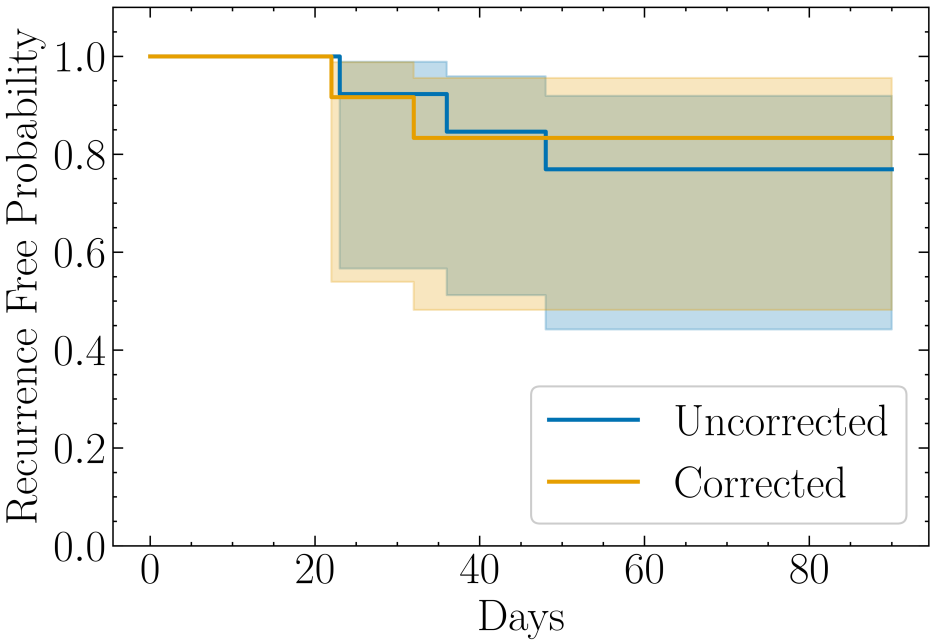

Figure 6.1: Kaplan-Meier recurrence free survival for the corrected and uncorrected arms over the 90 day window. Shaded bands are 95% confidence intervals. The curves overlap throughout, and a log-rank test found no significant difference ($P = 0.76$).

6.4.3 Hypoxia metric structure

Pearson correlations among the six hypoxia metrics are given in Table 6.3. Hsurf, HF10, and HVolume were mutually correlated at $r = 0.88$ to 0.93. HsurfNorm correlated with these three at $r = 0.53$ to 0.74. HRegions correlated with Hsurf and HsurfNorm at $r = 0.52$ and 0.73, and with HF10 and HVolume at $r = 0.27$ and 0.24. All correlations between TumorRate and the five hypoxia metrics were below 0.3 in magnitude. Selected relationships are illustrated in Figs. 6.2, 6.3, and 6.4. Hypoxic surface area rises roughly monotonically with HF10 (Fig. 6.2), consistent with both metrics measuring similar physical manifestations. The number of hypoxic regions rises and then falls across the HF10 range (Fig. 6.3), a weak

Table 6.3: Pearson correlation coefficients among the six hypoxia metrics.

	HRegions	Hsurf	HsurfNorm	TumorRate	HF10	HVolume
HRegions	1.00	0.52	0.73	−0.09	0.27	0.24
Hsurf	0.52	1.00	0.74	0.23	0.88	0.93
HsurfNorm	0.73	0.74	1.00	−0.04	0.56	0.53
TumorRate	−0.09	0.23	−0.04	1.00	0.14	0.28
HF10	0.27	0.88	0.56	0.14	1.00	0.92
HVolume	0.24	0.93	0.53	0.28	0.92	1.00

and nonmonotonic relationship, and tumor growth rate shows no trend with HF10 (Fig. 6.4).

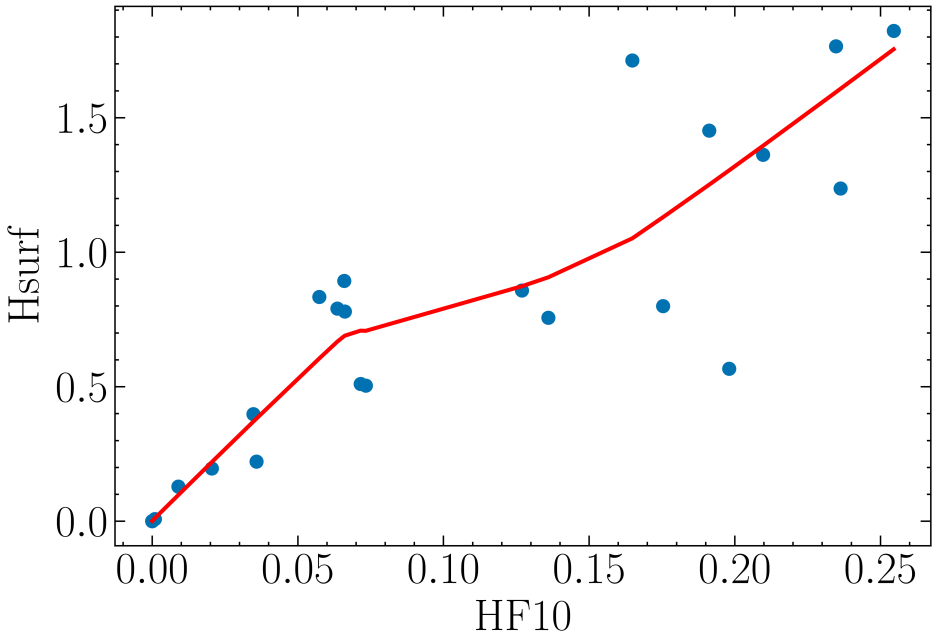

Figure 6.2: locally estimated scatterplot smoothing (LOESS)-smoothed scatter of hypoxic surface area against HF10, with the red curve as the LOESS fit. Hypoxic surface area rises roughly monotonically with HF10, a strong positive relationship.

6.5 Conclusion

We compared hypoxia-guided CRT boosts targeted on corrected and uncorrected FMISO images in a preclinical FSa fibrosarcoma model. Both protocols received whole-tumor doses described in Chapter 5. Over the 90 day follow up, recurrence free survival did not differ

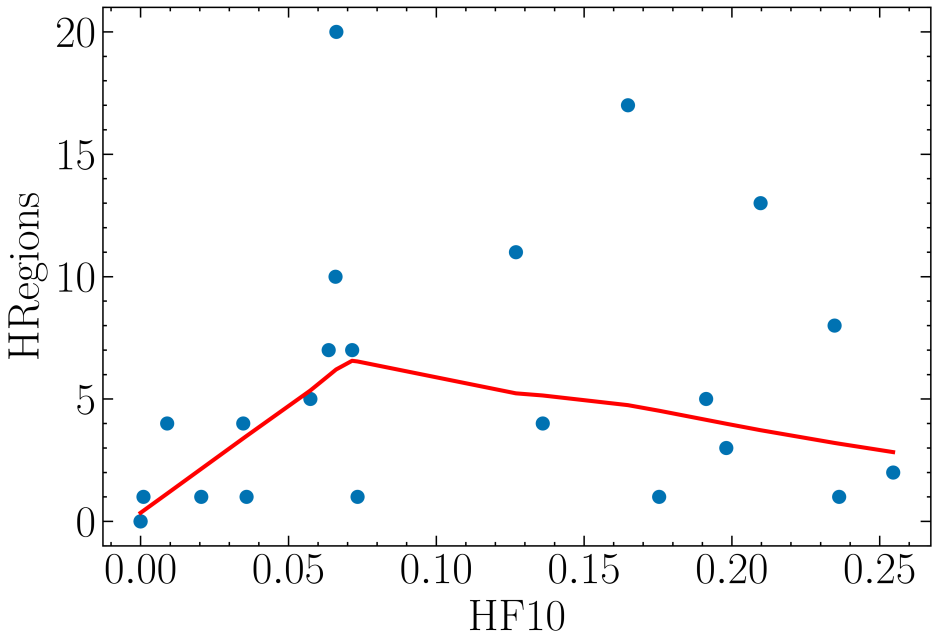

Figure 6.3: LOESS-smoothed scatter of the number of hypoxic regions against HF10, with the red curve as the LOESS fit. The number of hypoxic regions rises then falls across the HF10 range, a weak and nonmonotonic relationship.

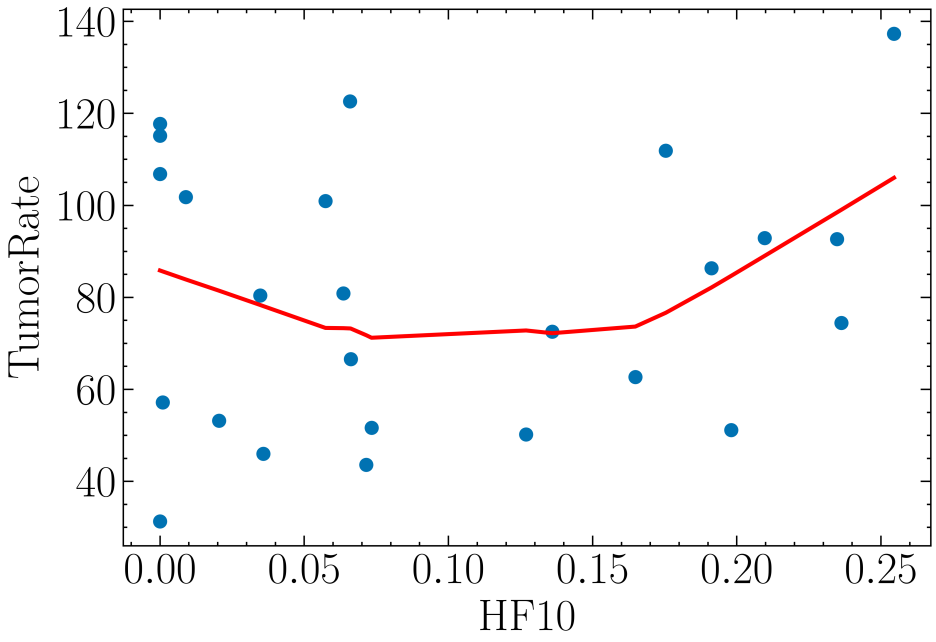

Figure 6.4: LOESS-smoothed scatter of tumor growth rate against HF10, with the red curve as the LOESS fit. Tumor growth rate shows no trend with HF10.

between the two arms.

This null result should be read as inconclusive rather than as evidence that the correction has no effect on outcome. Although the boost doses followed previously established tumor control dose curves [24], considerably fewer recurrences occurred than those curves predicted, yielding only five events among 26 tumors. With so few events, the study has little statistical power to resolve a difference but may suggest that the FMISO-guided boost produced good tumor control in both arms. However, increased statistics will be needed for more conclusive results.

CHAPTER 7

OPTIMAL THRESHOLD ANALYSIS OF FMISO IMAGES IN PRECLINICAL FIBROSARCOMA MODELS

7.1 Introduction

The data collected in Chapter 6, while originally intended for use in the survival analysis of corrected and uncorrected ^{18}F-fluoromisonidazole (FMISO)-based oxygen-guided conformal radiation therapy (CRT), can also be used to re-examine the optimal FMISO threshold identified by Gertsenshteyn *et al.* [104]. For this analysis, we combine fibrosarcoma (FSa) multi-modal datasets from both non-CRT and CRT protocols.

The previous threshold to be evaluated against, tumor-to-muscle ratio (TMR) ≥ 1.4, was calculated for FSa preclinical tumors [105]. That study employed 10 FSa tumors. Since then, we have gathered data on an additional 30 FSa tumors, all imaged with the protocol described in Chapter 5. Each tumor provided spatially co-registered voxel-wise partial pressure of oxygen (pO_2), standardized uptake value (SUV), and TMR maps.

7.2 Methods

7.2.1 Imaging data and region of interest

The region-of-interest (ROI) for a tumor was the set of voxels with valid pO_2, SUV, and TMR, and a voxel was taken to be truly hypoxic when its pO_2 level was at or below 10 mmHg. A tumor was excluded from analysis when fewer than 5% of its ROI voxels were hypoxic by this definition, or when it had no valid ROI.

7.2.2 Definition of SUV and TMR

The analysis uses two scores derived from the same reconstructed positron emission tomography (PET) activity image. For a voxel x with measured activity concentration $C(x)$, the SUV normalizes the uptake to the injected activity A_{inj} and the animal mass m,

$$\text{SUV}(x) = \frac{C(x)}{A_{\text{inj}}/m}. \tag{7.1}$$

The TMR further normalizes each voxel to a reference equal to the mean SUV over a muscle region M drawn from the same image,

$$\text{TMR}(x) = \frac{\text{SUV}(x)}{\overline{\text{SUV}}_M}, \qquad \overline{\text{SUV}}_M = \frac{1}{|M|} \sum_{y \in M} \text{SUV}(y), \tag{7.2}$$

where $|\cdot|$ denotes cardinality, or number of voxels. Since $\overline{\text{SUV}}_M$ is a single positive number for a given image, the two scores are proportional within each tumor,

$$\text{TMR}(x) = c\,\text{SUV}(x), \qquad c = \frac{1}{\overline{\text{SUV}}_M} > 0, \tag{7.3}$$

where the constant c is fixed within a tumor but differs between tumors, because each tumor has its own muscle reference.

7.2.3 Thresholding strategies

Five rules for labeling a voxel hypoxic on the PET side were compared on the same tumors. The first was the fixed threshold TMR ≥ 1.4. The remaining four each contained a single free parameter θ optimized on the data, namely an absolute cutoff TMR $\geq \theta$, a per-image TMR factor TMR $\geq \theta \cdot \text{mean(TMR)}$, an absolute cutoff SUV$\geq \theta$, and a per-image SUV factor SUV $\geq \theta \cdot \max(\text{SUV})$. The two factor rules rescale the threshold to each tumor.

7.2.4 Agreement metrics and the composite score

For each tumor and strategy, the resulting hypoxic mask B, the set of voxels labeled hypoxic by the strategy, was compared with the pO_2 hypoxic mask A, the set of ground truth hypoxic voxels defined by electron paramagnetic resonance (EPR), using three metrics, 1) the Dice similarity coefficient (DSC), 2) the 95th percentile Hausdorff distance ($d_{H,95}$), and 3) the balanced accuracy (BA). Both masks are defined on the image domain $\Omega \subset \mathbb{R}^3$, the bounding box for each tumor's ROI. The DSC is calculated as

$$\mathrm{DSC} = \frac{2\,|A \cap B|}{|A| + |B|}, \tag{7.4}$$

and measures volumetric overlap. The BA, the mean of sensitivity and specificity, measures voxel-wise performance in classifying voxels as hypoxic or non-hypoxic. Because BA is a mean of two rates its chance level is 0.5, so it was rescaled to $\widehat{\mathrm{BA}} = 2\,\mathrm{BA} - 1$. The $d_{H,95}$ is a boundary-agreement measure. For each point on the ground truth hypoxic boundary it takes the distance to the nearest point on the threshold-defined hypoxic boundary, does the same in the reverse direction, and reports the 95th percentile of these distances. The three were combined into a single overall hypoxic similarity (OHS) score,

$$\mathrm{OHS} = \tfrac{1}{3}\left(\widehat{\mathrm{BA}} + \mathrm{DSC} + \hat{\mathrm{d}}_{\mathrm{H},95}\right), \tag{7.5}$$

where

$$\hat{\mathrm{d}}_{\mathrm{H},95} = 1 - \frac{\mathrm{d}_{\mathrm{H},95}}{\mathrm{diam}(\Omega)}, \qquad \mathrm{diam}(\Omega) = \sup_{x,y\in\Omega} \|x - y\|_2, \tag{7.6}$$

and $\|\cdot\|_2$ is the Euclidean norm. The Hausdorff distance is rescaled by the reference distance $\mathrm{diam}(\Omega)$, the tumor ROI bounding box (Ω) diagonal (voxel range of this diagonal is approximately 21.7 to 31.0), so that all three terms lie on a common zero-to-one scale with one denoting perfect agreement. In particular, $\hat{\mathrm{d}}_{\mathrm{H},95} = 1$ when the two boundaries coincide

at the 95th percentile, that is when no more than 5% of the boundary points on either side lie more than zero voxels from the other boundary.

7.2.5 Threshold optimization

Two types of thresholding rule were considered, an *absolute* rule and a *factor* rule, differing in whether the cutoff is shared across tumors or rescaled to each tumor. The absolute rule labels a voxel hypoxic when its TMR or SUV exceeds a single fixed cutoff applied to every tumor,

$$\text{TMR} \geq \theta \qquad \text{or} \qquad \text{SUV} \geq \theta, \tag{7.7}$$

where the cutoff θ carries the same units as the image. This is the form of the established TMR ≥ 1.4 threshold [104], and a single global cutoff is what allows a fixed clinical decision value to be reported, so it is a natural rule. The factor rule instead labels a voxel hypoxic when its value exceeds a per-tumor multiple of that tumor's own summary statistic,

$$\text{TMR} \geq \theta \cdot \text{mean}(\text{TMR}) \qquad \text{or} \qquad \text{SUV} \geq \theta \cdot \max(\text{SUV}), \tag{7.8}$$

where θ is now a dimensionless multiplier and the effective cutoff rescales to each tumor. Each score is paired with the summary statistic appropriate to it. TMR is already normalized to muscle uptake, so its distribution is comparable across tumors and its mean is ideally a stable per-tumor reference. Raw SUV is not normalized and its mean depends on injected activity, so its per-tumor maximum is used as the reference. The cross pairings, a TMR maximum factor and a SUV mean factor, were therefore not included.

Considering Eq. 7.3, for a factor rule the constant c appears on both sides of the inequality

and cancels,

$$\begin{aligned} \mathrm{TMR} &\geq \theta \cdot \mathrm{mean}(\mathrm{TMR}) \\ c\,\mathrm{SUV} &\geq \theta \cdot c\,\mathrm{mean}(\mathrm{SUV}) \\ \mathrm{SUV} &\geq \theta \cdot \mathrm{mean}(\mathrm{SUV}), \end{aligned} \tag{7.9}$$

and identically for the maximum. Therefore, a given factor rule for TMR and the same factor rule for SUV differ by scalar. For this reason we chose to use mean and maximum as the reference statistic. Additionally, we use both TMR and SUV to examine data in both units, though in principle one may be used and the other derived using a unit conversion. For an absolute rule the constant does not cancel,

$$\mathrm{TMR} \geq \theta \iff \mathrm{SUV} \geq \theta/c, \tag{7.10}$$

so a single global TMR cutoff corresponds to a different SUV cutoff in every tumor and vice versa.

This gives the four strategies, absolute TMR, factor TMR, absolute SUV, and factor SUV, which together span the two scores under both a shared and a per-image cutoff. For each strategy, θ was chosen to maximize the mean OHS across the included tumors. The search used 1000 candidate values of θ, spanning the observed range of TMR or SUV across all tumors for the two absolute strategies and the interval from 0 to 5 for the two factor strategies. For the factor rules, the mean(TMR) and max(SUV) references were computed per tumor.

7.2.6 Voxel-level discrimination

Threshold-free discrimination was assessed by a voxel-level receiver operating characteristic curve (ROC) analysis using the PET score as a continuous predictor of pO_2-defined hypoxia.

Let Ω denote the set of all voxels in the tumor volume, and let $A \subseteq \Omega$ denote the subset of voxels classified as hypoxic according to the pO_2 reference. For a voxel-wise PET score $S(x)$ and threshold t, the predicted hypoxic region is defined as $B_S(t) = \{x \in \Omega : S(x) \geq t\}$. The true positive fraction (TPF) and false positive fraction (FPF) are then given by

$$\mathrm{TPF}_S(t) = \frac{|A \cap B_S(t)|}{|A|}, \qquad \mathrm{FPF}_S(t) = \frac{|B_S(t) \setminus A|}{|\Omega \setminus A|}. \tag{7.11}$$

where $|\cdot|$ denotes the number of voxels in a set. Thus, $|A|$ is the total number of hypoxic voxels, $|\Omega \setminus A|$ is the total number of non-hypoxic voxels, $|A \cap B_S(t)|$ is the number of correctly identified hypoxic voxels (true positives), and $|B_S(t) \cap (\Omega \setminus A)|$ is the number of non-hypoxic voxels incorrectly classified as hypoxic (false positives).

The ROC curve is the set of points $\{(\mathrm{FPF}_S(t), \mathrm{TPF}_S(t)) : t \in \mathbb{R}\}$, which forms a piecewise-constant curve as the threshold t sweeps the range of observed scores. Each tumor's ROC curve was computed as the empirical step function returned by the ROC algorithm. For visualization, these curves were then evaluated on a common FPF grid using linear interpolation and averaged to obtain the mean ROC. The reported area under the curve (AUC) is the mean of the individual tumor AUCs.

The TMR and SUV curves are identical within every tumor. By Eq. 7.3 the two scores differ by a positive constant, so for any cutoff t,

$$B_{\mathrm{TMR}}(t) = \{x : c\,\mathrm{SUV}(x) \geq t\} = \{x : \mathrm{SUV}(x) \geq t/c\} = B_{\mathrm{SUV}}(t/c). \tag{7.12}$$

Therefore, since t/c sweeps along with t, the two scores generate same set of $(\mathrm{FPF}, \mathrm{TPF})$ pairs. The AUC is therefore identical for TMR and SUV, however both are shown for completeness.

7.2.7 Statistical analysis

Each thresholding method predicts a binary hypoxic mask from a tumor's PET image, which is compared against the EPR-defined hypoxic mask, so the methods are evaluated on how well the predicted masks agree with ground truth. Across the 20 tumors, the five methods were compared with paired t-tests on the per-tumor metric values. The family-wise type I error rate was controlled with the Holm-Bonferroni step-down correction, which controls the same family-wise rate as the Bonferroni correction but is uniformly more powerful.

For the four thresholding methods the free parameter θ was both estimated and evaluated on the same tumors, so their apparent advantage over the fixed baseline is optimistically biased. To quantify and remove this bias, each tunable method was re-assessed by three tumor-level resampling schemes, in each of which θ was re-estimated on a training set of tumors by maximizing their mean OHS and then scored on the held-out tumors at that fixed θ. In leave-one-out cross-validation, one tumor was held out and θ was estimated on the other 19, repeated so that every tumor served once as the held-out case. In repeated five-fold cross-validation, the 20 tumors were partitioned at random into five folds of whole tumors, θ was estimated on four folds and scored on the fifth, each fold served once as the validation set, and the whole partitioning was repeated 50 times with different random splits and the per-tumor held-out scores averaged.

For the 0.632 bootstrap, 1000 new datasets were created by randomly sampling the 20 tumors with replacement, meaning a tumor could be selected multiple times while others were omitted [113]. For each resampled dataset, the optimal threshold θ was determined and then evaluated on the out-of-bag (OOB) tumors, i.e., the tumors that were not included in that resampled dataset. The final performance estimate was calculated as a weighted average of the in-sample and out-of-bag results:

$$\widehat{\mathrm{OHS}}_{0.632} = 0.368 \times \widehat{\mathrm{OHS}}_{\mathrm{in}} + 0.632 \times \widehat{\mathrm{OHS}}_{\mathrm{OOB}}, \tag{7.13}$$

where $\widehat{\mathrm{OHS}}_{\mathrm{in}}$ is the mean OHS on the resampled dataset and $\widehat{\mathrm{OHS}}_{\mathrm{OOB}}$ is the mean OHS on the omitted or OOB tumors. For each method the optimism was defined as the in-sample mean OHS minus the held-out mean.

7.3 Results

7.3.1 Cohort-level hypoxia statistics

Of the 40 tumors, 20 met the inclusion criteria. 19 tumors were excluded for a pO_2 hypoxic fraction below 5%, including seven with no hypoxic voxels at all, and one for an empty ROI. Fig. 7.1 shows the distribution of per-tumor hypoxic fraction together with the inclusion floor, and Fig. 7.2 shows the pooled voxel-wise pO_2 distribution over the retained tumors.

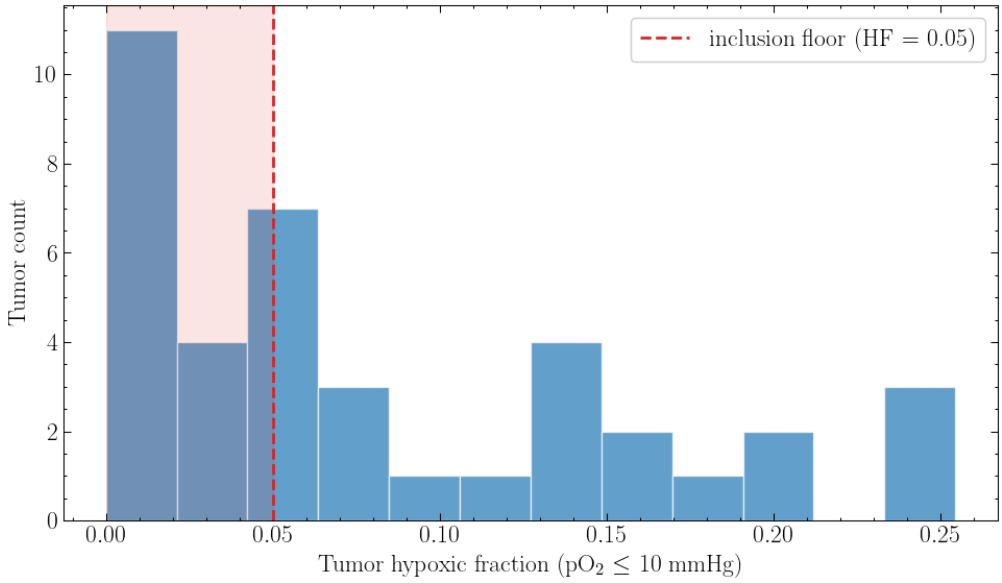

Figure 7.1: Histogram of per-tumor hypoxic fraction, the fraction of ROI voxels at or below 10 mmHg, across all tumors with a valid ROI. The dashed line is the 5% inclusion floor and the shaded region marks the tumors that were excluded.

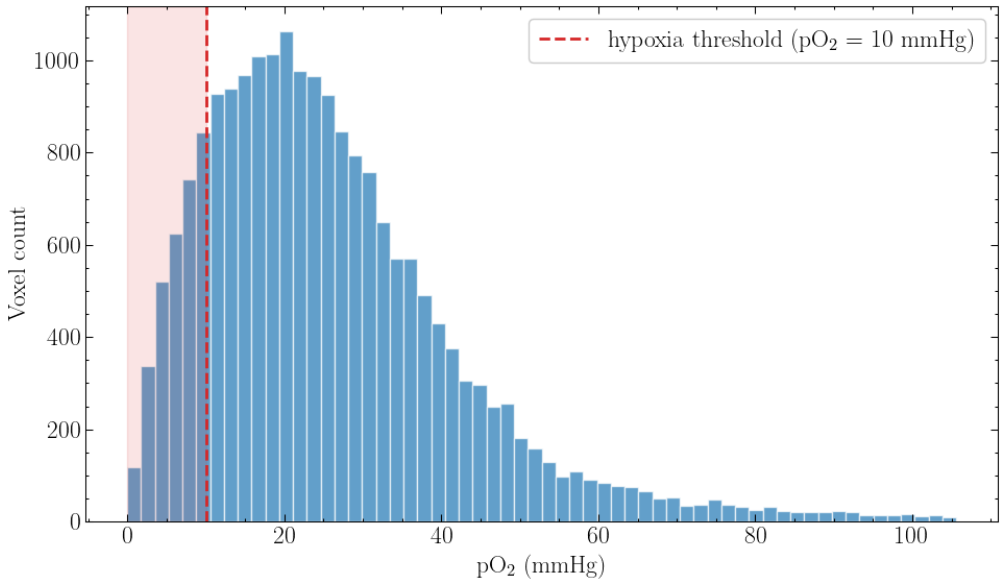

Figure 7.2: Distribution of voxel-wise pO_2 pooled across all included tumors within the ROI. The dashed line marks the 10 mmHg hypoxia threshold and the shaded region the hypoxic range.

7.3.2 Voxel-level discrimination

Across the retained tumors, TMR separated hypoxic from oxygenated voxels with a mean area under the ROC AUC of 0.8 ± 0.2 (0.775 with extra precision ignoring the uncertainty), shown in Fig. 7.3. A single tumor had an AUC of less than chance at 0.363. This single anti-correlated result is not unexpected, and was seen in published results [104]. The voxel-wise relationship is shown in Fig. 7.4. For completeness the identical ROC curve for SUV is also shown.

7.3.3 Optimal thresholds and strategy comparison

The optimal absolute TMR cutoff converged to 1.40, effectively identical to the established TMR $\geq$ 1.4, so optimizing the absolute threshold gave no improvement over the literature value. The factor rules, however, were the strategies with the highest OHS. Table 7.1 reports the performance metrics of each strategy to predict EPR-defined hypoxia. The median and interquartile range (IQR) are used here because the metrics are bounded, skewed, and

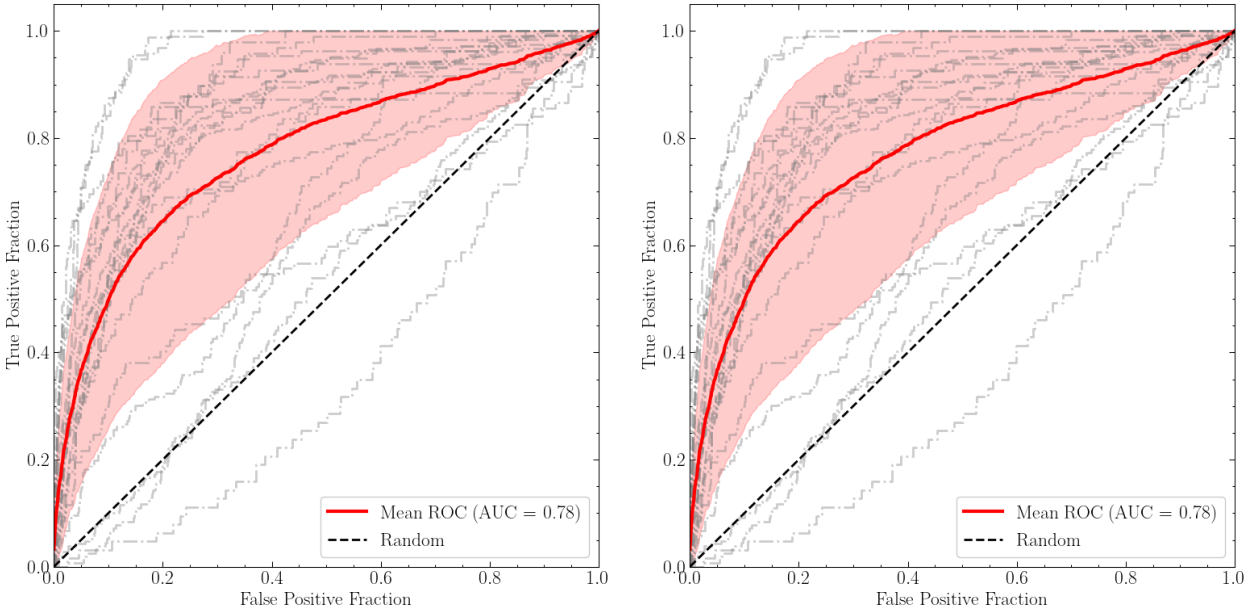

Figure 7.3: Voxel-level ROC curves for TMR (left) and SUV (right) as scores to predict for pO_2-defined hypoxia. Grey curves represent ROC curves for individual tumors, the red curve is the mean across tumors, and the shaded band denotes the ± standard error (SE). The TMR and SUV curves are expected to be identical.

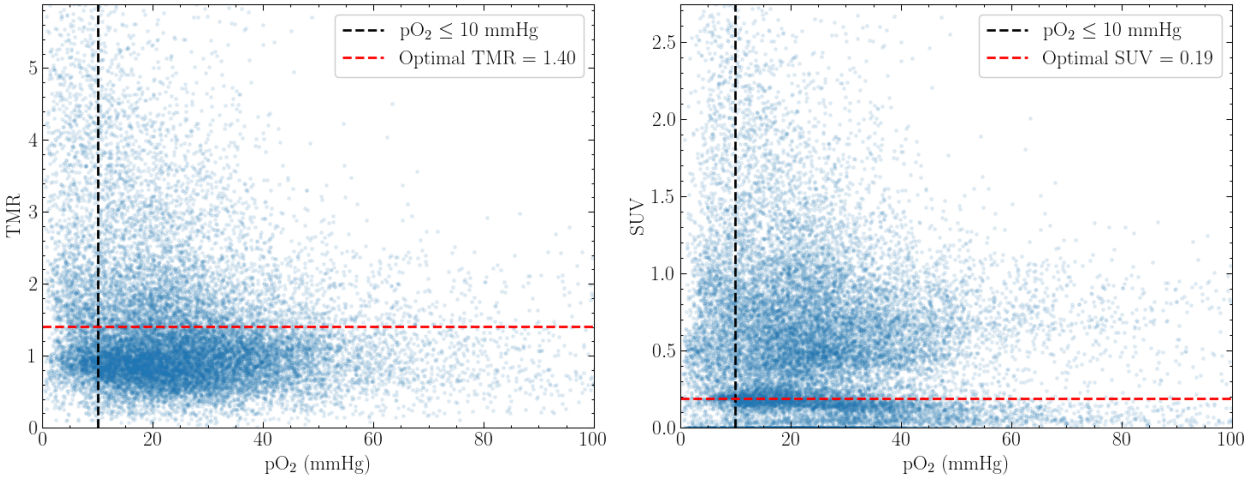

Figure 7.4: Voxel-wise relationship between pO_2 and PET signal, pooled across the included tumors, with TMR on the left and SUV on the right. The vertical dashed line is the 10 mmHg hypoxia threshold and the horizontal dashed line is the optimized θ using the absolute rule.

contain outliers across only 20 tumors. The rules TMR $\geq 1.19 \cdot \text{mean(TMR)}$ and SUV $\geq 0.44 \cdot \max(\text{SUV})$ reached a median OHS of 0.62 and 0.60 and a median DSC of 0.43, compared to the values 0.51 and 0.31 for the fixed baseline, respectively. It was observed that the SUV factor rule matched the TMR factor rule. The corresponding OHS functions versus cutoff values are shown in Fig. 7.5 and the full distributions of each metric in Fig. 7.7. Fig. 7.6 shows each individual metric as a function of the corresponding threshold value, and Fig. 7.8 shows the box plots for each individual metric. The shaded areas in Figs. 7.5 and 7.6 were calculated using the SE of the mean across tumors at each threshold value, estimated by bootstrap resampling of the 20 tumors.

Table 7.1: Performance metrics for each thresholding method predicting EPR-defined hypoxia across the 20 included tumors, reported as median and IQR. All four columns are on a common zero-to-one scale for which higher values indicate better agreement with ground truth. For the SUV$\geq$ 0.19 method the $\hat{d}_{H,95}$ entry is computed from 19 tumors, since in one tumor no voxel exceeded the cutoff and the predicted mask was empty, leaving $\hat{\text{d}}_{\text{H,95}}$ undefined. A fifth method, the optimized absolute TMR cutoff, converged to 1.40 and is not listed separately because it is identical to the fixed TMR$\geq$ 1.4 baseline shown in the first row.

Method	DSC	$\widehat{\text{BA}}$	$\hat{\text{d}}_{\text{H,95}}$	OHS
TMR$\geq$ 1.4	0.31 [0.13, 0.52]	0.38 [0.06, 0.51]	0.84 [0.81, 0.88]	0.51 [0.34, 0.62]
TMR $\geq 1.19 \cdot \text{mean(TMR)}$	0.43 [0.33, 0.55]	0.53 [0.36, 0.63]	0.88 [0.86, 0.90]	0.62 [0.54, 0.67]
SUV$\geq$ 0.19	0.24 [0.13, 0.34]	0.00 [0.00, 0.02]	0.81 [0.76, 0.85]	0.35 [0.32, 0.38]
SUV $\geq 0.44 \cdot \max(\text{SUV})$	0.43 [0.34, 0.51]	0.48 [0.33, 0.62]	0.88 [0.84, 0.90]	0.60 [0.53, 0.67]

7.3.4 Internal validation

Cross-validation showed that the higher OHS of the factor rules generalizes rather than reflects overfitting, and the conclusion was stable across three resampling schemes. Table 7.2 and Fig. 7.9 summarize the results. Each method was compared against the fixed TMR $\geq$ 1.4 baseline with a two-sided paired t-test on the per-tumor OHS values, in which each of the 20 tumors contributes one pair, its held-out score from the thresholding method and its score

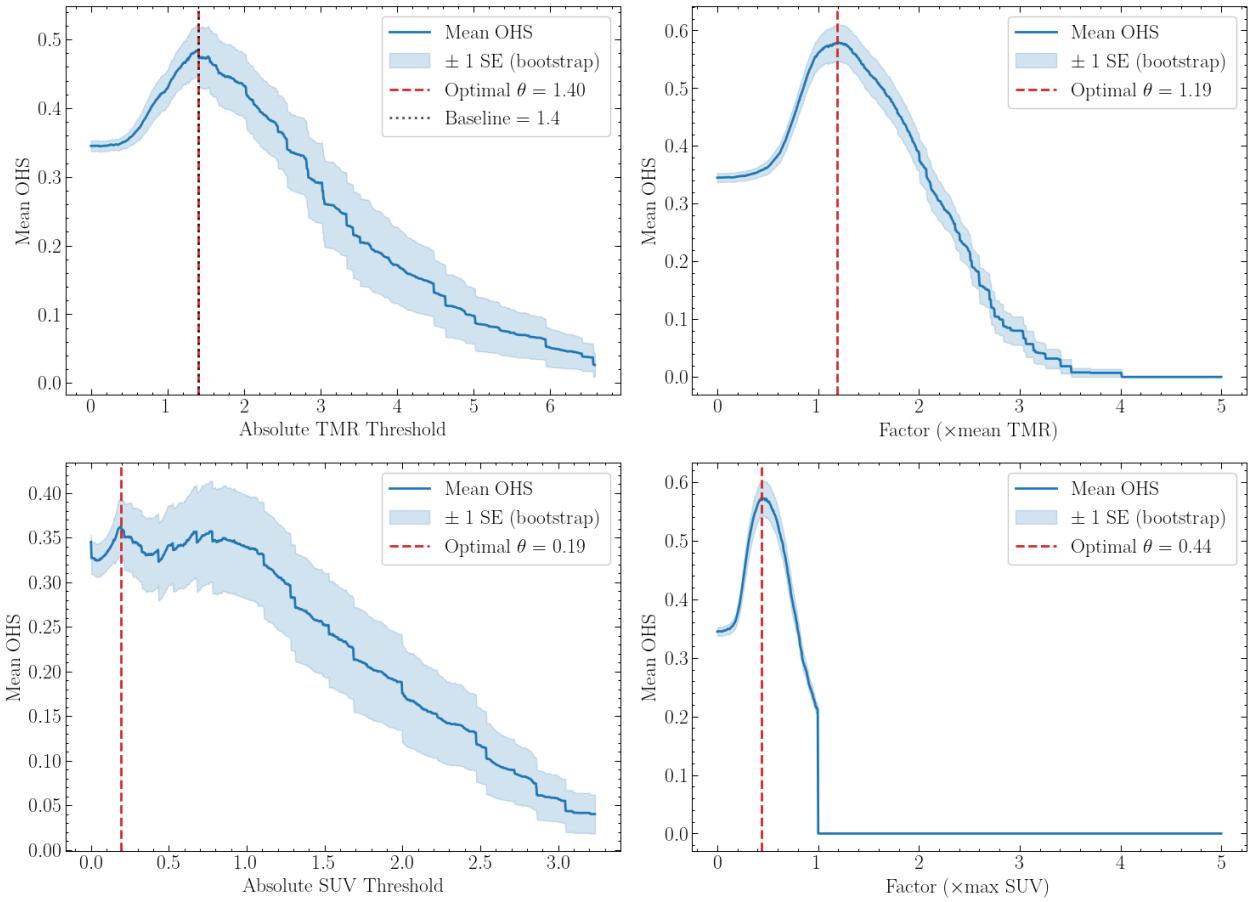

Figure 7.5: Mean OHS as a function of the decision threshold for each thresholding method, with the shaded band showing ± SE across tumors. The top row is the absolute and per-image-factor TMR, and the bottom row the absolute and per-image-factor SUV. The dashed line is the optimal threshold and, on the absolute-TMR panel, the dotted line is the TMR ≥ 1.4 baseline.

under the fixed baseline, giving 19 degrees of freedom. For each metric the four baseline-versus-method comparisons form one group, and the reported p values are Holm-Bonferroni adjusted across that group so they may be compared directly against the significance level $\alpha = 0.05$.

The two factor rules generalized well, with negligible optimism under all schemes (leave-one-out (LOO) 0.001 and 0.012 for the TMR and SUV factors, five-fold 0.001 and 0.008), and their LOO, five-fold, and 0.632-bootstrap held-out OHS agreed to within 0.01. Both methods improved on the fixed baseline on the held-out tumors, with the TMR factor rule by 0.095 in

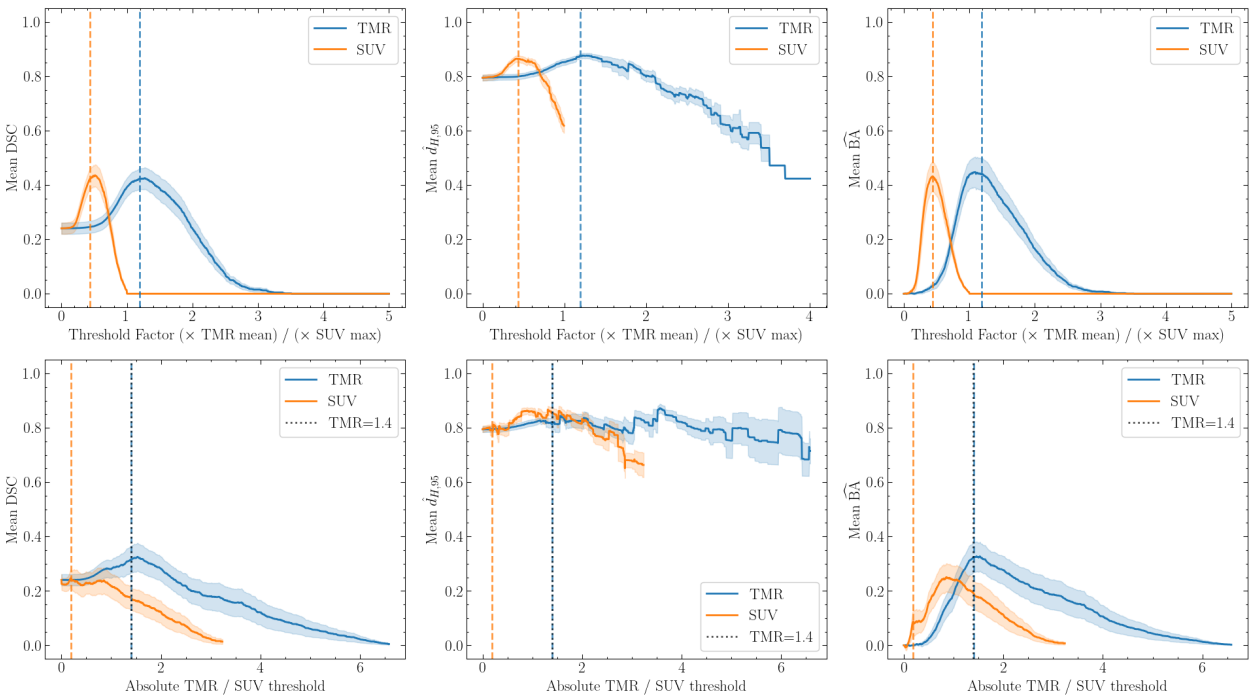

Figure 7.6: Mean DSC, $\hat{d}_{H,95}$, and $\widehat{BA}$, overlaying TMR and SUV, with bands showing $\pm$ SE. The top row shows the per-image factor rules and the bottom row shows the absolute cutoffs. The figures are arranged in columns for the DSC, $\hat{d}_{H,95}$, and $\widehat{BA}$, respectively. The dashed vertical lines are the locations of the respective optimal thresholds, and the dotted line in the bottom-left panel is the baseline threshold scheme. The vertical lines are the OHS maximum locations (not the individual metric maximum locations).

mean OHS and the SUV factor rule by 0.077 in mean OHS, and both differences remained significant after the Holm-Bonferroni correction under both LOO ($p = 0.006$ and $p = 0.026$) and five-fold cross-validation ($p = 0.007$ and $p = 0.019$). By contrast, the optimized absolute TMR cutoff was no better than the baseline after cross-validation ($p = 0.311$), and the absolute SUV cutoff was significantly worse, carrying by far the largest optimism (LOO 0.142) and a held-out OHS of 0.263 below the baseline ($p < 0.001$).

7.4 Discussion

The ROC analysis shows moderate discrimination, with a mean per-tumor AUC of 0.8 ± 0.2. The large uncertainty indicates that FMISO uptake and pO_2 discrepancies vary substantially

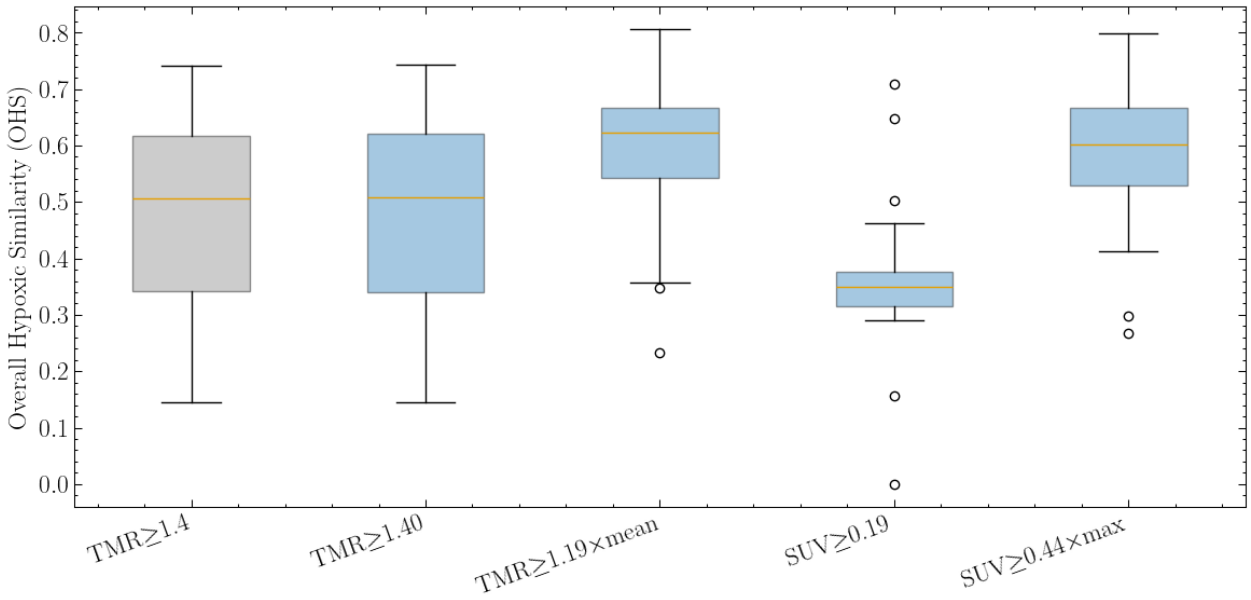

Figure 7.7: OHS for each thresholding strategy across the included tumors after optimization. Boxes show the median and IQR with whiskers to the full range if there are no outliers, or to 1.5 × IQR otherwise.

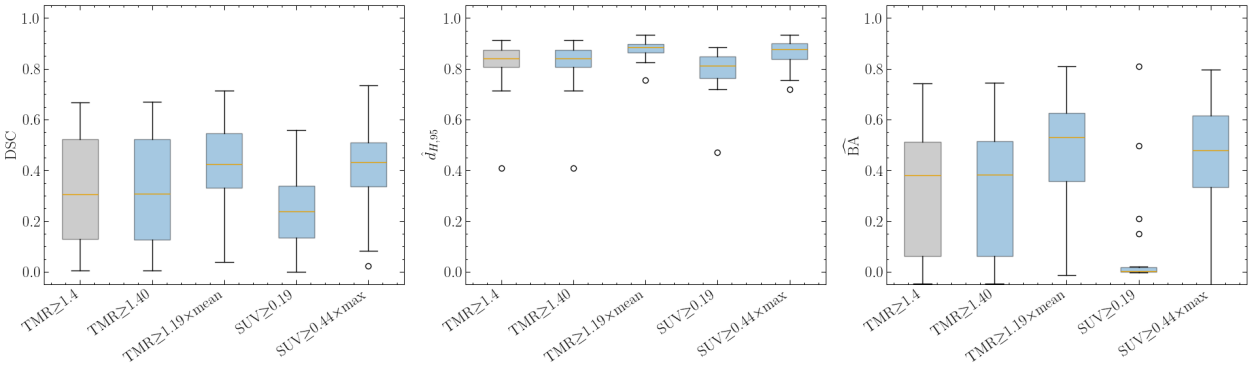

Figure 7.8: Distributions of the enabled component metrics by thresholding strategy, from left to right the DSC, the $\hat{d}_{H,95}$ in voxels, and $\widehat{BA}$. Higher is better for all metrics.

between lesions of a single tumor type, since all 40 starting tumors were FSas. For the OHS thresholding analysis, agreement with the EPR-defined hypoxic regions saturated at a level well below perfect overlap for every method tested. The best rules reached only a median DSC near 0.43, and the overlap could not be raised further by any choice of threshold. This is due to a voxel-level AUC of 0.8 ± 0.2 combined with the low prevalence of hypoxic voxels, which together bound the achievable overlap regardless of where the operating threshold is

Table 7.2: Optimized thresholds and held-out mean OHS under LOO, repeated five-fold cross validation, 0.632-bootstrap validation, and the optimism. Optimal θ maximized the in-sample mean OHS. For reference, the fixed TMR$\geq$ 1.4 baseline had a mean OHS of 0.484.

Method	Optimal θ	In-sample	LOO	Five-fold	0.632	Optimism
TMR$\geq \theta$	1.40	0.484	0.471	0.470	0.468	0.013
TMR$\geq \theta$·mean (TMR)	1.19	0.580	0.579	0.578	0.577	0.001
SUV$\geq \theta$	0.19	0.362	0.220	0.288	0.327	0.142
SUV$\geq \theta$·max (SUV)	0.44	0.573	0.561	0.565	0.569	0.012

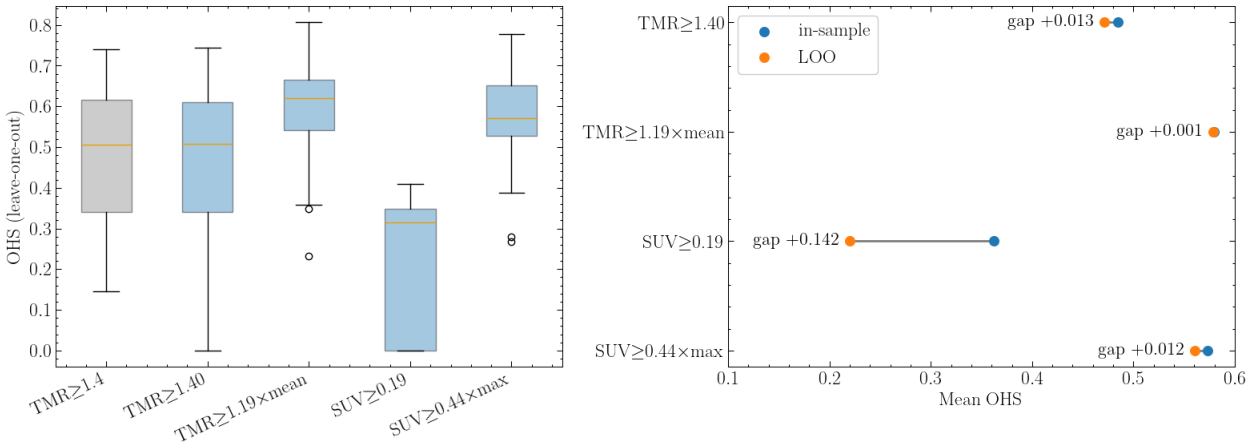

Figure 7.9: LOO cross-validation. On the left the cross-validated OHS for each strategy, and on the right the in-sample versus held-out mean OHS for the four tunable strategies with the optimism gap annotated.

placed. Further improvement in analysis would require more hypoxic voxels to test against the FMISO uptake. Additionally, the choice of threshold based on physical characteristics such as tumor size, hypoxic fraction, or other measurements may be incorporated in the future.

7.5 Conclusions

In this chapter we combined FSa multi-modal datasets from the non-CRT and CRT protocols into a single cohort of 40 tumors, 20 of which met the hypoxic-fraction inclusion criterion, and used it to re-examine the optimal FMISO threshold for defining tumor hypoxia against

EPR-pO_2 ground truth. Optimizing an absolute TMR cutoff reproduced the established value of TMR $\geq$ 1.4 almost exactly, confirming the absolute threshold rule reported by Gertsenshteyn *et al.* [104] on this larger dataset. The two per-image factor rules, based on TMR or SUV, agreed more closely with the EPR-defined hypoxic regions than the fixed absolute baseline. Internal validation was conducted using LOO, five-fold cross-validation, and bootstrap resampling. The factor rules demonstrated minimal optimism bias and confirmed a significant performance improvement for the factor rules, following Holm-Bonferroni adjustments for multiple comparisons. The optimized absolute SUV cutoff was significantly worse than the baseline and carried the largest optimism.

The analysis revealed a ceiling on the achievable agreement. The best strategies reached a median DSC near 0.43, and no choice of threshold raised the overlap further, a limit imposed jointly by the moderate voxel-level discrimination of TMR, with a mean AUC of 0.775, and the low prevalence of hypoxic voxels in these tumors. Threshold selection alone has therefore approached what the FMISO signal may support, and further gains in localizing hypoxia will require either datasets with a larger hypoxic burden or improvements to the FMISO data itself, such as the dynamic contrast-enhanced (DCE)-based correction examined in Chapter 6.

CHAPTER 8

THE FIRST POSITRONIUM LIFETIME TOMOGRAPHY DONE USING ^{44g}Sc WITH GROUND TRUTH EPR pO_2

8.1 Introduction

The experience gained from developing the robust conformal radiation therapy (CRT) protocol detailed in Chapter 5 for the ^{18}F-fluoromisonidazole (FMISO) Kaplan-Meier survival analysis and related dynamic contrast-enhanced (DCE)-based correction work allowed our team to conduct the first positronium lifetime tomography (PLT) using a ground truth reference to oxygen. This proof-of-concept experiment was designed to test the capabilities of the combined electron paramagnetic resonance (EPR)-positron emission tomography (PET) scanner and the moment-based reconstruction method for PLT while being able to compare the resulting orthopositronium (o-Ps) lifetime maps back to ground truth oxygen maps using the EPR capabilities. To date, there are no known systems besides our EPR-PET scanner that allow direct association between measured o-Ps lifetimes and partial pressure of oxygen (pO_2) *in vivo*.

The experiment was conducted in two phases. The first was a phantom imaging acquisition using gas phantoms of varying oxygen levels which were hypothesized to show a large o-Ps lifetime contrast. The second was a single animal imaging acquisition in which the o-Ps lifetime maps were compared between T2-weigted (T2W)-magnetic resonance imaging (MRI)-defined tumor and muscle regions. Since the tumor had expected hypoxic regions, o-Ps lifetime contrast should also be observed.

^{44g}Sc was the radioisotope chosen for this experiment. A conventional PET tracer needs to only emit a positron, since the image is reconstructed from the two coincident annihilation photons alone. positronium lifetime imaging (PLI)/PLT additionally requires a reference for the time at which the positron is emitted, so that the interval between positron emission and

annihilation can be recovered. As explained in Chapter 1, this reference is ideally supplied by a prompt gamma ray emitted within a few picoseconds of the positron [114]. A radioisotope suited to PLI must therefore emit both a positron and a prompt gamma ray, which restricts the candidates to $\beta^+\gamma$ emitters. Within this group, ^{44g}Sc has been identified as the most-promising candidate for PLI since, as discussed below, it has more well-suited properties than competing isotopes [67, 115, 116].

Each annihilation photon carries 0.511 MeV, fixed by the rest mass of the electron and positron, and inside a detector these photons generate a Compton edge at 0.341 MeV where a large share of the counts accumulate [35]. For the prompt gamma ray to be identified reliably by its energy, that energy must sit well away from 0.341 MeV. The prompt gamma ray of ^{44g}Sc lies at 1.157 MeV, above the Compton edge and separated from the annihilation photons [117].

^{44g}Sc is also importantly more efficient at producing triple coincidence events than other $\beta^+\gamma$ emitters. It undergoes β^+ decay in 94.27% of decays and emits a 1157 keV prompt γ ray in approximately 99.9% of decays. Thus, approximately 94.3% of decays produce both a positron and the prompt γ ray [117]. This high combined branching ratio increases the counting statistics that the lifetime estimate relies on.

Finally, ^{44g}Sc has a clinically-viable half-life. The 4.04 hour half-life of ^{44g}Sc is close to the 1.83 hour half-life of ^{18}F, the most widely used PET radioisotope [118], and far shorter than the multi-day half-lives of the longest-lived PET tracers [119], so its dosimetry is reasonably usable.

8.2 Methods

^{44g}Sc was produced at the University of Chicago Cyclotron Facility using the ^{44}Ca(p, n)^{44g}Sc reaction. The radioisotope was radiolabeled with prostate-specific membrane antigen (PSMA)-617 using a similar protocol established for other scandium radioisotopes [120]. The proton

beam was focused onto a ^{44}CaO target which had been calcified from $^{44}CaCO_3$ by heating to 900 deg C for several hours. The beam energy was set to 18 MeV and was degraded to 13.6 MeV.

Gas phantoms were made from glass ampoules containing deionized water through which the chosen gas was bubbled continuously throughout the PET acquisition. Approximately 170 µCi of ^{44g}Sc-PSMA-617 was injected into each phantom. The phantoms were connected to gas cylinders which supplied them with independent sources of air, argon, and carbogen (95% oxygen and 5% carbon dioxide). The phantom acquisition lasted for 4 hours 24 minutes.

The same batch of ^{44g}Sc-PSMA-617 was used for the mouse imaging. An athymic nude mouse which had grown a prostate tumor derived from the metastases from a patient with lymph node carcinoma of the prostate - androgen receptor (LNCaP-AR) cells [121] was injected with approximately 335 µCi of the radiotracer. The tracer was allowed to biodistribute for 3 hours before acquisition. The mouse was then imaged for 3 hours 30 minutes using the protocol described in Chapter 5. The imaging setup for both the phantom and mouse is shown in Fig. 8.1.

Both the phantom and mouse images were reconstructed using the method developed and tested computationally in Part I, however here it was extended to full 3-dimensional reconstruction. A moment of $n = 8$ was used for both reconstructions, a selected axial slice of which is shown in Fig. 8.2 for the phantom and Fig. 8.3, respectively. In this proof-of-concept study, the optimization of parameters was not conducted. The reconstruction parameters were qualitatively chosen.

8.3 Results

In both acquisitions the measured lifetimes fell within the $\sim$ 1 ns to 5 ns range expected for o-Ps in aqueous and biological media [56]. For the phantom, three regions were segmented from the reconstructed moment image using an intensity threshold. Their mean o-Ps lifetimes

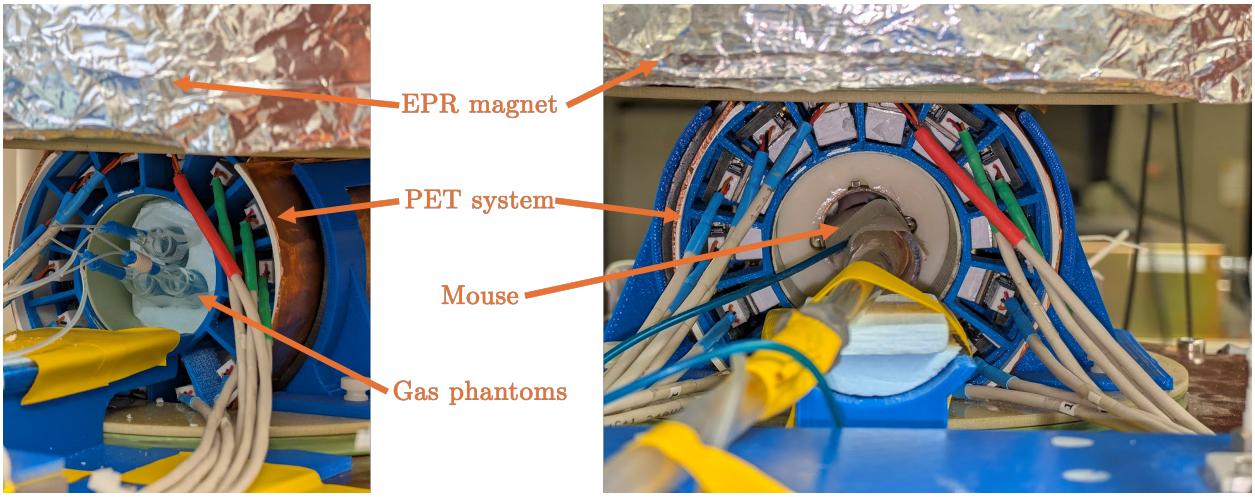

Figure 8.1: Experimental setup for the two PLI acquisitions. The gas phantoms were glass ampoules of deionized water through which air, argon, or carbogen was bubbled continuously during imaging, each loaded with approximately 170 µCi of ^{44g}Sc-PSMA-617. The mouse, bearing an LNCaP-AR prostate tumor, received approximately 335 µCi of the same radiotracer. Both were imaged in the combined EPR-PET scanner following the protocol of Chapter 5.

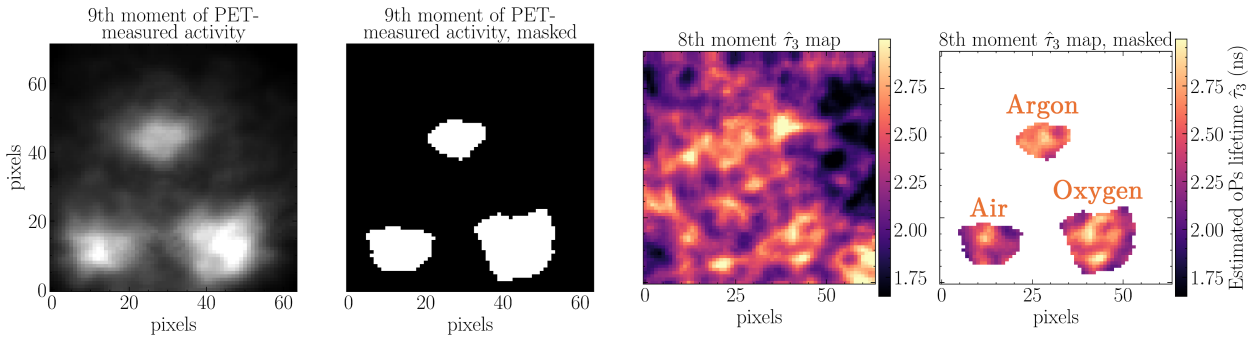

Figure 8.2: Reconstructed o-Ps lifetime map of the gas phantoms for a selected axial slice. The two left panels show the PET moment image used to define the phantom regions. The two right panels show the reconstructed o-Ps lifetime $\hat{\tau}_3$ map and the same map restricted to those regions. The three segmented regions had mean lifetimes of 2.31 ± 0.18 ns for the oxygen phantom, 2.31 ± 0.18 ns for the air phantom, and 2.59 ± 0.18 ns for the argon phantom, where the stated uncertainty is the voxel-level standard deviation within each region.

were 2.31 ns for the oxygen phantom, 2.31 ns for the air phantom, and 2.59 ns for the argon phantom, each with a voxel-level standard deviation of approximately 0.18 ns. For the mouse, the tumor and muscle regions were defined from T2W-MRI anatomical contouring and gave

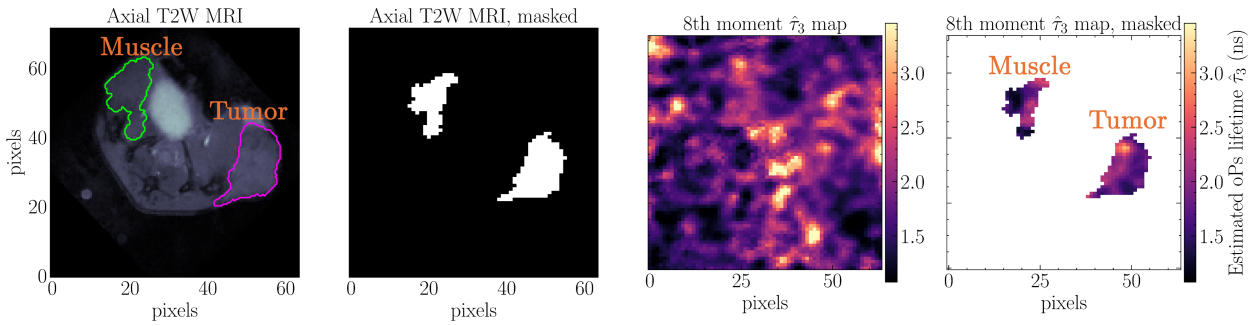

Figure 8.3: Reconstructed o-Ps lifetime map of the tumor-bearing mouse for a selected axial slice. The two left panels show the T2W-MRI used to segment the tumor and muscle, over the full field of view and then restricted to those regions. The two right panels show the reconstructed o-Ps lifetime $\hat{\tau}_3$ map and the same map restricted to the segmented regions. The tumor and muscle had mean lifetimes of 2.12 ± 0.39 ns and 1.97 ± 0.40 ns, where the stated uncertainty is the voxel-level standard deviation within each region.

mean o-Ps lifetimes of 2.12 ns and 1.97 ns, with voxel-level standard deviations of 0.39 ns and 0.40 ns, respectively, however we note that no corrections for random coincidences were applied in this proof-of-concept study, so the standard deviations may be reduced in future work. Also, the small animal PET scanner used for this study was not calibrated for and therefore lacked time-of-flight (TOF) capability.

8.4 Discussion

These results show that the combined EPR-PET scanner and the moment-based reconstruction produce physically reasonable o-Ps lifetime maps for both a phantom and a living animal. The modest oxygen contrast in both acquisitions is consistent with the counting-statistics limitations anticipated for a proof-of-concept in which the reconstruction parameters were not yet optimized.

For the phantom, the large lifetime contrast that was hypothesized between the differently oxygenated ampoules was not observed. The three regions spanned only about 0.3 ns, two of them were essentially identical, and only the longest-lived region separated from the others by more than its voxel-level spread.

For the mouse, even though contrast between regions was difficult to be seen, the tumor lifetime was the longer of the two. This is the expectation when the tumor is the more hypoxic tissue, since a lower oxygen concentration quenches o-Ps less strongly and lengthens its lifetime. The difference of approximately 0.15 ns was nonetheless small compared with the roughly 0.4 ns standard deviation within each region.

8.5 Conclusion

We have taken the first PLI data using ^{44g}Sc with a ground truth oxygen reference. The experimental work for this study was using the protocol developed for previous studies, and the reconstruction method used was developed in-house as well. Future work will involve preparation this data for publication. A simulation of the PET with realistic variables will be done to explore what can be reasonably expected by the system and the reconstruction method. The simulation will be used to guide reconstruction parameters in an effort to increase contrast in the lifetime images. In future experiments, especially proof-of-concepts, higher data rates or longer acquisition times should be considered to increase precision of the measurements and contrast between regions.

CHAPTER 9

CONCLUSIONS AND FUTURE DIRECTIONS

This dissertation was motivated by the goal to map tumor oxygenation *in vivo* with PET accurately enough to guide a radiation boost to hypoxic regions. Two pathways were explored to pursue this goal, a novel reconstruction method for PLI that senses oxygen indirectly through the o-Ps lifetime, and a correction-based refinement of the established hypoxia tracer FMISO, both evaluated against ground truth oxygen imaging.

Part I developed the reconstruction method for PLI. We derived a moment-ratio estimator of the o-Ps rate constant from two successive exponentially weighted, modified moments of the elapsed-time histogram, showed that it isolates the long-lived o-Ps component analytically, and demonstrated that it extends to tomography as two reconstructions followed by a voxel-wise ratio. The estimator was stable and monotonic on one-dimensional PALS data and produced robust two-dimensional o-Ps lifetime/rate constant maps in simulation, with a practical range of moment orders to use. We also demonstrated its capability across a range of PET timing resolutions and event-count numbers.

Part II focused on experimental methods and results using the combined EPR-PET scanner and its ground truth pO_2 maps. A DCE MRI-based correction of FMISO was assessed in a preclinical survival study, the optimal FMISO threshold was re-examined on an enlarged cohort and reproduced the established TMR $\geq$ 1.4 value while a per-image factor threshold agreed more closely with true hypoxia, establishing a new optimal threshold to define hypoxia. Finally, the reconstruction method developed for PLI was applied to the first data acquired with ^{44g}Sc alongside a ground truth oxygen reference, yielding physically reasonable o-Ps lifetime maps of a phantom and a tumor-bearing mouse as a proof of concept.

In conclusion this work strengthens the current FMISO-based route to hypoxia imaging while establishing the analytic and experimental groundwork for PLI. Neither method is yet ready for clinical hypoxia targeting, with the correction-based approach limited by

the fidelity of the FMISO signal and the experimental studies by their sample sizes and counting statistics, especially for PLI. There is a path forward, however. Larger cohorts and improved corrections for the FMISO-based method are needed, and parameter optimization, higher event counts, and improved timing resolution are required for the PLI method. As PET instrumentation and geometries continue to advance, the methods developed here are positioned to help a future introduction into oxygen-guided radiotherapy.

APPENDIX A

SUPPORTING PROOFS AND DETAILS

Below, we write $\mathrm{EXP}(t;\lambda) = \lambda e^{-\lambda t}$ for $t \geq 0$ and 0 otherwise for the exponential density, $\mathcal{N}(t;\sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-t^2/2\sigma^2}$ for the Gaussian of zero-mean, and

$$\psi_{\lambda,\sigma(t)} = \mathrm{EMG}(t;\lambda,\sigma) = \mathrm{EXP}(t;\lambda) * \mathcal{N}(t;\sigma) \tag{A.1}$$

for the exponentially modified Gaussian (EMG), where $*$ indicates convolution. We also make use of the truncated exponential series $E_m(x) = \sum_{k=0}^{m} x^k/k!$, the double factorial notation, i.e., $n!! = \prod_{k=0}^{\lceil \frac{n}{2} \rceil - 1}(n-2k) = n(n-2)(n-4)\cdots$ with conventions $(-1)!! = 0!! = 1$, and note again that the o-Ps component is the longest-lived in the Ps spectrum, so $\lambda_3 < \lambda_1, \lambda_2$.

This Appendix proves the results invoked in the main text: 1) the lower integration limit $t_{d,L} = -2.5\sigma$, 2) the moment-ratio relations Eq. (2.8) and Eq. (2.15), and 3) the closed form of the random-subtraction weight $\alpha_n(s;\mathcal{T})$ of Eq. (2.33).

A.1 Moment properties

For a function $f(t)$ and a shift $a \in \mathbb{R}$, the moment, translated moment, and windowed translated moment are defined as

$$\begin{aligned} \mu_n\{f(t)\} &= \int_{-\infty}^{\infty} t^n f(t)\, dt, \\ \mu_n^{(a)}\{f(t)\} &= \int_{-\infty}^{\infty} (t+a)^n f(t)\, dt, \\ \mu_n^{(a)}\{f(t);\mathcal{T}\} &= \int_{T_1}^{T_2} (t+a)^n f(t)\, dt, \end{aligned} \tag{A.2}$$

respectively, where $\mathcal{T}$ is an interval $[T_1, T_2)$, and $\mu_n^{(0)}\{f(t)\} = \mu_n\{f(t)\}$. We make use of two properties:

Linearity. Because Eq. (A.2) is an integral, for constants a_k and functions $g_k(t)$,

$$\mu_n^{(a)}\left\{\sum_k a_k g_k(t)\right\} = \sum_k a_k \mu_n^{(a)}\left\{g_k(t)\right\}. \tag{A.3}$$

Translation invariance. Translating the argument of f by the moment shift removes the shift. Let $g(t) = f(t+a)$. Then, $\mu_n^{(a0}\left\{g(t)\right\} = \int_{-\infty}^{\infty}(t+a)^n\, g\,(t)\, dt$ and we have

$$\begin{aligned}\mu_n^{(a)}\left\{f(t+a)\right\} &= \int_{-\infty}^{\infty}(t+a)^n\, f\,(t+a)\, dt \\ &= \int_{-\infty}^{\infty} u^n f\,(u)\, du,\ u = t+a \\ &= \mu_n\left\{f\right\},\end{aligned} \tag{A.4}$$

the substitution $u = t + a$ clearing the shift in the weight and in f simultaneously.

A.2 The EMG, its tail, and the lower integration limit

It has been shown [122] that $\psi_{\lambda,\sigma(t)}$ may be written as

$$\psi_{\lambda,\sigma}(t) = \frac{\lambda}{\sqrt{\pi}}\, e^{\sigma^2\lambda^2/2-\lambda t}\int_{-\infty}^{Z(t)} e^{-x^2}\, dx \tag{A.5}$$

where

$$Z(t) = \frac{1}{\sqrt{2}}\left(\frac{t}{\sigma} - \sigma\lambda\right). \tag{A.6}$$

We evaluate the integral in Eq. (A.5),

$$\begin{aligned}\int_{-\infty}^{Z(t)} e^{-x^2}\, dx &= \int_{-\infty}^{0} e^{-x^2}\, dx + \int_{0}^{Z(t)} e^{-x^2}\, dx \\ &= \frac{\sqrt{\pi}}{2} + \frac{\sqrt{\pi}}{2}\mathrm{erf(Z)} \\ &= \frac{\sqrt{\pi}}{2}\left[1 + \mathrm{erf}(Z)\right],\end{aligned} \tag{A.7}$$

and Eq. (A.5) becomes

$$\psi_{\lambda,\sigma}(t) = \lambda\, e^{-\lambda t}\, e^{\sigma^2\lambda^2/2}\, h(t;\lambda,\sigma), \tag{A.8}$$

where

$$h(t;\lambda,\sigma) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{t-\sigma^2\lambda}{\sqrt{2}\,\sigma}\right)\right]. \tag{A.9}$$

As $t \to +\infty$, $h \to 1$, therefore $\lim_{x\to\infty} \psi_{\lambda,\sigma}(t) = 0$.

Left-tail bound. We bound $\psi_{\lambda,\sigma}(t)$ for $t < 0$. Writing Eq. (A.1) in its expanded form gives

$$\begin{aligned}\psi_{\lambda,\sigma}(t) &= \int_0^\infty \lambda e^{-\lambda\tau} \frac{1}{\sqrt{2\pi}\sigma} e^{-(t-\tau)^2/2\sigma^2} d\tau \\ &= \frac{\lambda}{\sqrt{2\pi}\sigma} e^{-t^2/2\sigma^2} \int_0^\infty e^{-(\lambda - t/\sigma^2)\tau} e^{-\tau^2/2\sigma^2} d\tau,\end{aligned} \tag{A.10}$$

where we have used $(t-\tau)^2 = t^2 - 2t\tau + \tau^2$. Let $\lambda^* = \lambda - t/\sigma^2$ for $t < 0$. λ^* becomes $\lambda + |t|/\sigma^2 > 0$. Since $e^{-\tau^2/2\sigma^2} \le 1$,

$$\int_0^\infty e^{-\lambda^*\tau} e^{-\tau^2/2\sigma^2} d\tau < \int_0^\infty e^{-\lambda^*\tau} d\tau = \frac{1}{\lambda^*}. \tag{A.11}$$

Substituting this result into Eq. (A.10) gives

$$\psi_{\lambda,\sigma}(t) < \frac{\lambda}{\lambda^*}\mathcal{N}(t;\sigma) = \frac{\mathcal{N}(t;\sigma)}{1 + |t|/\left(\sigma^2\lambda\right)} < \mathcal{N}(t;\sigma), \quad t < 0. \tag{A.12}$$

Lower integration limit. The measured signal is $\tilde{S}(t_d) = A\sum_k c_k \psi_{\lambda_{k,\sigma}}(t_d)$ with $\sum_k c_k = 1$. From Eq. (A.12), $\mathbb{E}\{\tilde{S}(t_d)\} < A\mathcal{N}(t_d;\sigma)\sum_k c_k = A\mathcal{N}(t_d;\sigma)$. The Gaussian cumulative probability below -2.5σ is

$$\int_{-\infty}^{-2.5\sigma} \mathcal{N}(t;\sigma)\, dt = \Phi(-2.5) \approx 6.2 \times 10^{-3}, \tag{A.13}$$

so for $t_d < t_{d,L}$ with $t_{d,L} = -2.5\sigma$, we approximate $\mathbb{E}\{\tilde{S}(t_d)\} \approx 0$. With Gaussian smoothing for $\sigma \to \sigma^*$, the same argument can be used.

Exponential reweighting property. We now exponentially weight the moment of $\psi_{\lambda,\sigma}$ in an effort to control noise. Evaluating Eq. (A.9) at $t + \sigma^2 s$ with rate constant $\lambda + s$ gives

$$\begin{aligned} h(t + \sigma^2 s; \lambda + s, \sigma) &= \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{(t + \sigma^2 s) - \sigma^2(\lambda + s)}{\sqrt{2}\,\sigma}\right)\right] \\ &= \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{t - \sigma^2\lambda}{\sqrt{2}\,\sigma}\right)\right] \\ &= h(t; \lambda, \sigma). \end{aligned} \tag{A.14}$$

Expanding Eq. (A.8) with the new arguments, and substituting $h(t; \lambda, \sigma) = \psi_{\lambda,\sigma(t)}/(\lambda e^{-\lambda t} e^{\sigma^2\lambda^2/2})$, Eq. (A.8) is simplified to

$$\psi_{s+\lambda,\sigma}\left(t + \sigma^2 s\right) = \frac{s + \lambda}{\lambda} e^{-st} e^{-\sigma^2 s^2/2} \psi_{\lambda,\sigma}(t). \tag{A.15}$$

We solve Eq. (A.15) for $e^{-st}\psi_{\lambda,\sigma}(t)$, taking the translated moment $\mu_n^{(\sigma^2 s)}$, and using the property from Eq. (A.4):

$$\begin{aligned} \mu_n^{(\sigma^2 s)}\left\{e^{-st}\psi_{\lambda,\sigma}(t)\right\} &= \frac{\lambda e^{\sigma^2 s^2/2}}{s + \lambda} \mu_n^{(\sigma^2 s)}\left\{\psi_{s+\lambda,\sigma}\left(t + \sigma^2 s\right)\right\} \\ &= \frac{\lambda e^{\sigma^2 s^2/2}}{s + \lambda} \mu_n\left\{\psi_{s+\lambda,\sigma}\right\}. \end{aligned} \tag{A.16}$$

A.3 Moments of the EMG

Lemma A.1. *For $f_3 = f_1 * f_2$,*

$$\mu_n\{f_3\} = \sum_{m=0}^{n} \binom{n}{m} \mu_{n-m}\{f_1\}\mu_m\{f_2\}. \tag{A.17}$$

Proof. Consider the Fourier transform of $f(t)$ given by $\mathscr{F}\{f(t)\}(\nu) = \hat{f}(\nu) = \int_{-\infty}^{\infty} f(t)e^{-j2\pi\nu}\,dt$.

Differentiating n times under the integral brings down a factor $(-j2\pi\nu)^n$, so that $\hat{f}^{[n]}(\nu = 0) = (-j2\pi)^n \int_{-\infty}^{\infty} t^n f(t)\, dt = (-j2\pi)^n \mu_n\{f(t)\}$. Equivalently,

$$\mu_n\{f(t)\} = (-j2\pi)^{-n} \hat{f}^{[n]}(0). \tag{A.18}$$

By the convolution theorem $\hat{f}_3 = \hat{f}_1 \hat{f}_2$. The general Leibniz rule for the nth derivative of a product, evaluated at $\nu = 0$, gives

$$\hat{f}_3^{[n]}(0) = \sum_{m=0}^{n} \binom{n}{m} \hat{f}_1^{[n-m]}(0) \hat{f}_2^{[m]}(0). \tag{A.19}$$

Multiplying Eq. (A.19) by

$$(-j2\pi)^{-n} = (-j2\pi)^{-(n-m)} (-j2\pi)^{-m}, \tag{A.20}$$

the equation becomes Eq. (A.17) via Eq. (A.18). ■

Lemma A.2. *The moments of Eq. (A.1) are*

$$\mu_n\left\{\psi_{\lambda,\sigma}(t)\right\} = \frac{n!}{\lambda^n} E_{\lfloor n/2 \rfloor}\left(\sigma^2\lambda^2/2\right). \tag{A.21}$$

Proof. The moments of the exponential distribution are $\mu_n\{EXP(t;\lambda)\} = \int_0^{\infty} t^n \lambda e^{-\lambda t}\, dt = n!/\lambda^n$, and the Gaussian central moments are $\mu_m\{\mathcal{N}(t;\sigma)\} = \sigma^m (m-1)!!$ for even m and $\mu_m\{\mathcal{N}(t;\sigma)\} = 0$ for odd m. Setting $m = 2k$ in Eq. (A.17) for $m \in \{0, 2, \ldots, n\}$,

$k \in \{0, 1, \ldots, \lfloor n/2 \rfloor\}$, and using $f_3 = \psi_{\lambda,\sigma}$,

$$\begin{aligned}
\mu_n\{\psi_{\lambda,\sigma}\} &= \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} \mu_{n-2k}\{\mathrm{EXP}\}\mu_{2k}\{\mathcal{N}\} \\
&= \sum_{k=0}^{\lfloor n/2 \rfloor} \frac{n!}{(2k)!(n-2k)!} \cdot \frac{(n-2k)!}{\lambda^{n-2k}} \cdot \sigma^{2k}(2k-1)!! \\
&= \frac{n!}{\lambda^n} \sum_{k=0}^{\lfloor n/2 \rfloor} \frac{(2k-1)!!}{(2k)!}(\sigma\lambda)^{2k}.
\end{aligned} \tag{A.22}$$

Since $n! = n!!(n-1)!!$ and $(2k)!! = 2^k k!$, the summand in Eq. (A.22) becomes

$$\begin{aligned}
\frac{(2k-1)!!}{(2k)!}(\sigma\lambda)^{2k} &= \frac{1}{2^k k!}(\sigma\lambda)^{2k} \\
&= \frac{1}{k!}\left(\frac{\lambda^2\sigma^2}{2}\right)^k.
\end{aligned} \tag{A.23}$$

Substituting this into Eq. (A.22) and recalling that $E_m(x) = \sum_{k=0}^{m} x^k/k!$, we have the desired result. ■

Lemma A.3. *For $x_a > x_b > 0$ and $n \in \mathbb{N}$,*

$$\frac{E_{n+1}(x_a)}{E_{n+1}(x_b)} < \frac{x_a}{x_b}\frac{E_n(x_a)}{E_n(x_b)}. \tag{A.24}$$

Proof. Let $P(x) = \sum_{k=0}^{n} x^k/(k+1)!$. $E_{n+1}(x)$ is then

$$\begin{aligned}
E_{n+1}(x) &= \sum_{k=0}^{n+1} \frac{x^k}{k!} \\
&= 1 + \sum_{k=1}^{n+1} \frac{x^k}{k!} \\
&= 1 + \sum_{j=0}^{n} \frac{x^{j+1}}{(j+1)!} \\
&= 1 + xP(x)
\end{aligned} \tag{A.25}$$

where we reindexed so that $k = j + 1$.

Let $A = x_a P(x_a)$ and $B = x_b P(x_b)$. Since $x_a > x_b > 0$ and P has positive coefficients, $A > B > 0$, and

$$\frac{E_{n+1}(x_a)}{E_{n+1}(x_b)} = \frac{1+A}{1+B} < \frac{A}{B} = \frac{x_a P(x_a)}{x_b P(x_b)}, \tag{A.26}$$

where the inequality follows from

$$\frac{A}{B} - \frac{1+A}{1+B} = \frac{A-B}{B(1+B)} > 0. \tag{A.27}$$

Next we consider if $P(x_a)/P(x_b) < E_n(x_a)/E_n(x_b)$. Let

$$\begin{aligned} S &= P(x_a)E_n(x_b) - P(x_b)E_n(x_a) \\ &= \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} \frac{x_a^{k_1} x_b^{k_2}}{(k_1+1)!\,k_2!} - \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} \frac{x_b^{k_1} x_a^{k_2}}{(k_1+1)!\,k_2!} \\ &= \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} \frac{x_a^{k_1} x_b^{k_2} - x_b^{k_1} x_a^{k_2}}{(k_1+1)!\,k_2!} \\ &= \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} N(k_1,k_2)\, w(k_1,k_2). \end{aligned} \tag{A.28}$$

where $N(k_1, k_2) = x_a^{k_1} x_b^{k_2} - x_b^{k_1} x_a^{k_2}$ and $w(k_1, k_2) = 1/(k_1+1)k_2!$. We factor $x_b^{k_1} x_b^{k_2}$ from $N(k_1, k_2)$:

$$\begin{aligned} N(k_1,k_2) &= x_b^{k_1} x_b^{k_2} \left(\frac{x_a^{k_1}}{x_b^{k_1}} - \frac{x_a^{k_2}}{x_b^{k_2}} \right) \\ &= x_b^{k_1} x_b^{k_2} \left(\rho^{k_1} - \rho^{k_2} \right), \end{aligned} \tag{A.29}$$

where $\rho = x_a/x_b$, $\rho > 1$. The map $k \mapsto \rho^k$ is strictly increasing, so $\rho^{k_1} - \rho^{k_2}$ is positive when $k1 > k2$, is zero when $k_1 = k_2$, and is negative when $k_1 < k_2$. Therefore $N(k_1, k_2) > 0$ if and only if $k_1 > k_2$.

We re-label the dummy indices inside Eq. (A.28) which becomes

$$\begin{aligned} S &= \sum_{k_1=0}^{n} \sum_{k_2=0}^{n} N(k_2, k_1)\, w(k_2, k_1) \\ &= -\sum_{k_1=0}^{n} \sum_{k_2=0}^{n} N(k_1, k_2)\, w(k_2, k_1), \end{aligned} \tag{A.30}$$

since

$$\begin{aligned} N\left(k_2, k_1\right) &= x_a^{k_2} x_b^{k_1} - x_b^{k_2} x_a^{k_1} \\ &= -\left(x_a^{k_1} x_b^{k_2} - x_b^{k_1} x_a^{k_2}\right) \\ &= -N\left(k_1, k_2\right). \end{aligned} \tag{A.31}$$

Adding Eqs. (A.28) and (A.30) yields

$$2S = \sum_{k_1=0}^{n} \sum_{k_2=0}^{n} N\left(k_1, k_2\right) \left[w\left(k_1, k_2\right) - w\left(k_2, k_1\right)\right]. \tag{A.32}$$

Now treating the w terms, we find their common denominator and simplify:

$$w\left(k_1, k_2\right) - w\left(k_2, k_1\right) = \frac{k_2 - k_1}{\left(k_1 + 1\right)!\left(k_2 + 1\right)!}, \tag{A.33}$$

and we see that $w\left(k_1, k_2\right) - w\left(k_2, k_1\right)$ is positive when $k_2 > k_1$, is zero when $k_2 = k_1$, and is negative when $k_2 < k_1$. Substituting this result into Eq. (A.32) we have

$$S = \frac{1}{2} \sum_{k_1=0}^{n} \sum_{k_2=0}^{n} N\left(k_1, k_2\right) \frac{k_2 - k_1}{\left(k_1 + 1\right)!\left(k_2 + 1\right)!}. \tag{A.34}$$

For a single term, if $k_1 > k_2$ then $N(k_1, k_2) > 0$, and since $k_2 - k_1 < 0$, $N\left(k_1, k_2\right)\left(k_2 - k_1\right) < 0$. If $k_1 < k_2$ then $N(k_1, k_2) < 0$, $k_2 - k_1 > 0$, and still $N\left(k_1, k_2\right)\left(k_2 - k_1\right) < 0$. If $k_1 = k_2$ then the term is zero. Therefore, $S \leq 0$ with $S < 0$ for $n \geq 1$.

Since $P(x_b) > 0$ and $E_n(x_b) > 0$, $P(x_b)E_n(x_b) > 0$. $S < 0$ is equivalent to $P(x_a)E_n(x_b) <$

$P(x_b)E_n(x_a)$. Dividing by $P(x_b)E_n(x_b)$ gives

$$\frac{P(x_a)}{P(x_b)} < \frac{E_n(x_a)}{E_n(x_b)}. \tag{A.35}$$

Substituting Eq. (A.35) into Eq. (A.26) gives

$$\frac{E_{n+1}(x_a)}{E_{n+1}(x_b)} < \frac{x_a E_n(x_a)}{x_b E_n(x_b)}. \tag{A.36}$$

■

A.4 The moment-ratio relation

Lemma A.4. *For $\sigma > 0$,*

$$\frac{1}{n+1}\frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda,\sigma}\}}{\mu_n^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda,\sigma}\}} = \frac{1+\delta_n}{s+\lambda} \tag{A.37}$$

where $\delta_n = 0$ for even n, and $0 < \delta_n \ll 1$ for odd n when $\sigma(s+\lambda)$ is small or n is large.

Proof. By Eq. (A.16) both moments carry the order-independent factor $\lambda e^{\sigma^2 s^2/2}/(s+\lambda)$, so

$$\frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda,\sigma}\}}{\mu_n^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda,\sigma}\}} = \frac{\mu_{n+1}\{\psi_{s+\lambda,\sigma}\}}{\mu_n\{\psi_{s+\lambda,\sigma}\}}. \tag{A.38}$$

Using Lemma A.2 with $\lambda \to (s+\lambda)$ we write the ratio

$$\begin{aligned}\frac{\mu_{n+1}\{\psi_{s+\lambda,\sigma}\}}{\mu_n\{\psi_{s+\lambda,\sigma}\}} &= \frac{\frac{(n+1)!}{(s+\lambda)^{n+1}}E_{\lfloor (n+1)/2\rfloor}(x)}{\frac{n!}{(s+\lambda)^n}E_{\lfloor n/2\rfloor}(x)} \\ &= \frac{n+1}{s+\lambda}\frac{E_{\lfloor (n+1)/2\rfloor}(x)}{E_{\lfloor n/2\rfloor}(x)},\end{aligned} \tag{A.39}$$

where $x = \sigma^2(s+\lambda)^2/2$. By inspection we define

$$\delta_n \equiv \frac{E_{\lfloor (n+1)/2 \rfloor}(x)}{E_{\lfloor n/2 \rfloor}(x)} - 1. \tag{A.40}$$

Dividing by $n+1$ in Eq. (A.39) gives

$$\begin{aligned} \frac{1}{n+1}\frac{\mu_{n+1}\{\psi_{s+\lambda,\sigma}\}}{\mu_n\{\psi_{s+\lambda,\sigma}\}} &= \frac{1}{s+\lambda}\frac{E_{\lfloor (n+1)/2 \rfloor}(x)}{E_{\lfloor n/2 \rfloor}(x)} \\ &= \frac{1+\delta_n}{s+\lambda}. \end{aligned} \tag{A.41}$$

For even $n = 2m$, $\lfloor (n+1)/2 \rfloor = \lfloor n/2 \rfloor$ and so the two truncated series are identical and $\delta_n = 0$. When n is odd, $n = 2m+1$, $\lfloor (n+1)/2 \rfloor = m+1$ and $\lfloor n/2 \rfloor = m$, so

$$\begin{aligned} 0 < \delta_{2m+1} &= \frac{E_{m+1}(x)}{E_m(x)} - 1 \\ &= \frac{E_{m+1}(x) - E_m(x)}{E_m(x)} \\ &= \frac{x^{m+1}}{(m+1)!}\frac{1}{\sum_{k=0}^{m} x^k/k!} \leq \frac{x^{m+1}}{(m+1)!}, \end{aligned} \tag{A.42}$$

since $E_m(x) > 1$. When $\sigma(s+\lambda)$ is small or $m = (n-1)/2$ is large, $x^{m+1}/(m+1)! \ll 1$, and is always greater than zero. ■

Lemma A.5. *For* $0 < \lambda_b < \lambda_a$ *and* $x_{a,b} = \sigma^2\lambda_{a,b}^2/2$,

$$\gamma_n(\lambda_a, \lambda_b) \equiv \frac{\mu_n\{\psi_{\lambda_a,\sigma}\}}{\mu_n\{\psi_{\lambda_b,\sigma}\}} = \frac{\lambda_b^n E_{\lfloor n/2 \rfloor}(x_a)}{\lambda_a^n E_{\lfloor n/2 \rfloor}(x_b)} \tag{A.43}$$

decreases over even n *and satisfies*

$$\gamma_n(\lambda_a, \lambda_b) \sim e^{x_a - x_b}\left(\frac{\lambda_b}{\lambda_a}\right)^n \tag{A.44}$$

as $n \to \infty$, *so* $\lim_{n\to\infty} \gamma_n(\lambda_a, \lambda_b) = 0$.

Proof. The formula for Eq. (A.43) is immediately available from Lemma A.2. We compare consecutive terms: $2k$, $2k+1$, and $2k+2$. Since $\lfloor 2k/2 \rfloor = \lfloor (2k+1)/2 \rfloor = k$,

$$\begin{aligned} \frac{\gamma_{2k+1}}{\gamma_{2k}} &= \frac{(\lambda_b/\lambda_a)^{2k+1} E_k(x_a)/E_k(x_b)}{(\lambda_b/\lambda_a)^{2k} E_k(x_a)/E_k(x_b)} \\ &= \frac{\lambda_b}{\lambda_a} < 1. \end{aligned} \tag{A.45}$$

For the next step, we have $\lfloor (2k+2)/2 \rfloor = k+1$, so

$$\frac{\gamma_{2k+2}}{\gamma_{2k}} = \frac{\lambda_b^2}{\lambda_a^2} \frac{E_{k+1}(x_a)/E_{k+1}(x_b)}{E_k(x_a)/E_k(x_b)}. \tag{A.46}$$

Applying Lemma A.3 with $n \to k$, $x_a \to \sigma^2\lambda_a^2/2$, and $x_b \to \sigma^2\lambda_b^2/2$ gives

$$\frac{E_{k+1}(x_a)/E_{k+1}(x_b)}{E_k(x_a)/E_k(x_b)} < \frac{x_a}{x_b} = \frac{\lambda_a^2}{\lambda_b^2}, \tag{A.47}$$

so Eq. (A.46) satisfies the bound

$$\frac{\gamma_{2k+2}}{\gamma_{2k}} < \frac{\lambda_b^2}{\lambda_a^2} \frac{\lambda_a^2}{\lambda_b^2} = 1. \tag{A.48}$$

Thus $\gamma_{2k+1} < \gamma_{2k}$ and $\gamma_{2k+2} < \gamma_{2k}$, $\forall k \geq 1$. For clarity we factor Eq. (A.43) as

$$\gamma_n(\lambda_a, \lambda_b) = \left(\frac{\lambda_b}{\lambda_a}\right)^n \frac{E_{\lfloor n/2 \rfloor}(x_a)}{E_{\lfloor n/2 \rfloor}(x_b)}. \tag{A.49}$$

In the limit $n \to \infty$, $\lfloor n/2 \rfloor \to \infty$ and $E_{\lfloor n/2 \rfloor}(x) \to e^x$, $\forall x > 0$, so

$$\lim_{n\to\infty} \left[\frac{E_{\lfloor n/2 \rfloor}(x_a)}{E_{\lfloor n/2 \rfloor}(x_b)}\right] = e^{x_a - x_b}, \tag{A.50}$$

giving Eq. (A.44). Finally, since $0 < \lambda_b/\lambda_a < 1$,

$$\lim_{n\to\infty} \gamma_n(\lambda_a, \lambda_b) = e^{x_a - x_b} \lim_{n\to\infty} \left(\frac{\lambda_b}{\lambda_a}\right)^n = 0. \tag{A.51}$$

■

Lemma A.6. *Let* $\tilde{S}(t) = A\sum_{k=1}^{3} c_k \psi_{\lambda_k,\sigma}(t)$ *with* $\lambda_3, \lambda_{1,2}$. *For sufficiently large even* n,

$$r^{(n)}(s) = \frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\}}{\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\}} \approx \frac{n+1}{s+\lambda_3} \tag{A.52}$$

which is Eq. (2.8). The per-pixel relation Eq. (2.15) is identical with $\lambda_3 \to \lambda_{3,j}$.

Proof. Let $\lambda'_k = s + \lambda_k$. By Eqs. (A.3) and (A.16),

$$\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}\} = Ae^{\sigma^2 s^2/2} \sum_{k=1}^{3} c_k \frac{\lambda_k}{\lambda'_k} \mu_n\{\psi_{\lambda'_k,\sigma}\}. \tag{A.53}$$

Factoring out the $k = 3$ term of the sum we have

$$\begin{aligned} \mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}\} = A\, e^{\sigma^2 s^2/2} \Bigg[& c_3 \frac{\lambda_3}{\lambda'_3} \mu_n\{\psi_{\lambda'_3,\sigma}\} \\ & + \sum_{k=1}^{2} c_k \frac{\lambda_k}{\lambda'_k} \mu_n\{\psi_{\lambda'_k,\sigma}\} \Bigg], \end{aligned} \tag{A.54}$$

and we factor $c_3(\lambda_3/\lambda'_3)\mu_n\{\psi_{\lambda'_3,\sigma}\}$ giving

$$\begin{aligned} \mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}\} = A\, e^{\sigma^2 s^2/2} c_3 \frac{\lambda_3}{\lambda'_3} \mu_n\{\psi_{\lambda'_3,\sigma}\} \Bigg[& 1 + \\ & + \sum_{k=1}^{2} \frac{c_k}{c_3} \frac{\lambda_k/\lambda'_k}{\lambda_3/\lambda'_3} \frac{\mu_n\{\psi_{\lambda'_k,\sigma}\}}{\mu_n\{\psi_{\lambda'_3,\sigma}\}} \Bigg]. \end{aligned} \tag{A.55}$$

Using

$$\frac{\lambda_k/\lambda_k'}{\lambda_3/\lambda_3'} = \frac{\lambda_k\lambda_3'}{\lambda_3\lambda_k'} \tag{A.56}$$

and

$$\frac{\mu_n\{\psi_{\lambda_k',\sigma}\}}{\mu_n\{\psi_{\lambda_3',\sigma}\}} = \gamma_n(\lambda_k',\lambda_3') \tag{A.57}$$

from Eq. (A.43), and distributing c_3 back through the bracket,

$$\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}\} = A\, e^{\sigma^2 s^2/2}\,\frac{\lambda_3}{\lambda_3'}\,\mu_n\{\psi_{\lambda_3',\sigma}\}\Bigg[c_3 + \sum_{k=1}^{2} c_k \frac{\lambda_k\lambda_3'}{\lambda_3\lambda_k'}\gamma_n(\lambda_k',\lambda_3')\Bigg]. \tag{A.58}$$

Simplifying the $k=3$ term via Eq. (A.16) and expanding the sum we have

$$\begin{aligned}\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\} &= A\,\mu_n^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda_3,\sigma}(t)\} \\ &\times\Big\{c_1\frac{\lambda_1\lambda_3'}{\lambda_3\lambda_1'}\gamma_n(\lambda_1',\lambda_3') + c_2\frac{\lambda_2\lambda_3'}{\lambda_3\lambda_2'}\gamma_n(\lambda_2',\lambda_3') + c_3\Big\}.\end{aligned} \tag{A.59}$$

Since $\lambda_3 < \lambda_1, \lambda_2$ implies $\lambda_3' < \lambda_1', \lambda_2'$, Lemma A.4 gives $\gamma_n(\lambda_1',\lambda_2'), \gamma_n(\lambda_2',\lambda_3') \to 0$ and $n \to 0$, giving $\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\} = c_3 A\,\mu_n^{(\sigma^2 s)}\{e^{-st}\psi_{\lambda_3,\sigma}(t)\}$. The ratio of two consecutive moments is then

$$r^{(n)}(s) \approx \frac{\mu_{n+1}^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\}}{\mu_n^{(\sigma^2 s)}\{e^{-st}\tilde{S}(t)\}}, \tag{A.60}$$

and with Claim 4 using $\lambda \to \lambda_3$ and even n, Eq. (A.60) can be rearranged into the desired result. ■

Lemma A.7. *The random-subtraction weight of Eq. (2.33),*

$$\alpha_n(s;\mathcal{T}) = \mu_n^{(\sigma^2 s)}\{e^{-st};\mathcal{T}\} = \int_{T_1}^{T_2}(t+\sigma^2 s)^n e^{-st}\,dt, \tag{A.61}$$

equals

$$\alpha_n(s;\mathcal{T}) = \begin{cases} -\frac{n!}{s^{n+1}}\left[e^{-st}\sum_{k=0}^{n}\frac{(s(t+\sigma^2 s))^k}{k!}\right]_{T_1}^{T_2}, & s>0 \\ \frac{{T_2}^{n+1}-T_1^{n+1}}{n+1}, & s=0, \end{cases} \tag{A.62}$$

which is Eq. (2.35).

Proof. In the trivial case $s = 0$, the integrand in Eq. (A.61) is t^d and the definite integral evaluates to $({T_2}^{n+1} - T_1^{n+1})/(n+1)$. For $s > 0$, substitute $u = s(t + \sigma^2 s), du = s\,dt$. Then $t = u/s - s\sigma^2$, $e^{-st} = e^{s^2\sigma^2}e^{-u}$, and the integral becomes

$$\alpha_n(s;\mathcal{T}) = \frac{e^{s^2\sigma^2}}{s^{n+1}}\int_{T_1+\sigma^2 s}^{T_2+\sigma^2 s} u^n e^{-u}\,du. \tag{A.63}$$

We claim that, for $n \in \mathbb{N}_0$,

$$F(u) = \int u^n e^{-u}\,du = -e^{-u}\sum_{k=0}^{n}\frac{n!}{k!}u^k + C. \tag{A.64}$$

This is readily verified by taking the derivative:

$$\begin{aligned} F'(u) &= e^{-u}\sum_{k=0}^{n}\frac{n!}{k!}u^k - e^{-u}\sum_{k=1}^{n}\frac{n!}{(k-1)!}u^{k-1} \\ &= e^{-u}\left(\sum_{k=0}^{n}\frac{n!}{k!}u^k - \sum_{k=0}^{n-1}\frac{n!}{k!}u^k\right) \\ &= e^{-u}\left(\sum_{k=0}^{n+1}\frac{n!}{k!}u^k + \frac{n!}{n!}u^n - \sum_{k=0}^{n-1}\frac{n!}{k!}u^k\right) \\ &= u^n e^{-u}. \end{aligned} \tag{A.65}$$

Substituting back $u = s(t + \sigma^2 s)$ into Eq. (A.63),

$$\begin{aligned}\alpha_n(s;\mathcal{T}) &= -\frac{e^{s^2\sigma^2}}{s^{n+1}} e^{-(s(t+\sigma^2 s))} \sum_{k=0}^{n} \frac{n!}{k!} (s(t+\sigma^2 s))^k + C \\ &= -\frac{n!}{s^{n+1}} e^{-st} \sum_{k=0}^{n} \frac{(s(t+\sigma^2 s))^k}{k!} + C.\end{aligned} \tag{A.66}$$

Evaluating at the limits of integration T_1, T_2 gives the desired result. ■

A.5 List-mode evaluation of $\tilde{Y}_i^{(n)}(s)$

Here we show that the first term of Eq. (2.33), the windowed EMM of $y_i(t_d)$, reduces to a sum over the individual detected events in channel i. Consider the approximation

$$\mu_n^{(a)}\{e^{-st_d} y_i(t_d); \mathcal{T}\} \approx \sum_{m=0}^{M-1} (t_m + a)^n e^{-st_m} y_i(t_m) \Delta t_d, \tag{A.67}$$

where $t_m = T_1 + m\Delta t_d$ for $m = 0, \ldots, M-1$ and $\Delta t_d = \frac{T_2 - T_1}{M}$. $y_i(t_m)\Delta t_d$ is then the number of events in the mth bin. Let $S_{i,m}$ be the subset of events in channel i whose elapsed time $t_{d,k}$ is in $[t_m, t_{m+1})$. For small Δt_d, we have $y_i(t_m)\Delta t_d \approx |S_{i,m}| = \sum_{k \in S_{i,m}} 1$ and we may assume $t_{d,k} \approx t_m$ for $k \in S_{i,m}$. Continuing from Eq. (A.67), we have

$$\begin{aligned}\mu_n^{(a)}\{e^{-st_d} y_i(t_d); \mathcal{T}\} &\approx \sum_{m=0}^{M-1} (t_m + a)^n e^{-st_m} \left|S_{i,m}\right| \\ &\approx \sum_{m=0}^{M-1} \sum_{k \in S_{i,m}} (t_{d,k} + a)^n e^{-st_{d,k}} \\ &= \sum_{k \in S_i} (t_{d,k} + a)^n \, e^{-st_{d,k}},\end{aligned} \tag{A.68}$$

where the final step constructs the sum to be over all events S_i since $\left\{S_{i,m}\right\}_{m=0}^{M-1}$ forms a partition of S_i. Setting $a = \sigma^2 s$ and subtracting the random contribution $\hat{r}_i \, \alpha_n(s;\mathcal{T})$, whose

closed form is given in Lemma A.7, yields the accumulation formula Eq. (2.33).

A.6 Linearity of $\tilde{Y}_i^{(n)}(s)$

Let $\mu_n\{f(t_d);\mathcal{T}\}$ denote the nth moment of $f(t_d)$ calculated over $\mathcal{T}$. Beginning from the definition of $\tilde{Y}_i^{(n)}(s)$ in Eq. (2.13), restricting the moment to the measured window $\mathcal{T}$, and using the linearity of $\mu_n^{(\sigma^2 s)}\{e^{-st_d}\cdot\}$ with r_i constant in t_d, we have

$$\begin{aligned}
\tilde{Y}_i^{(n)}(s) &= \mu_n^{(\sigma^2 s)}\{e^{-st_d}\left(y_i(t_d) - r_i\right)\} \\
&\approx \mu_n^{(\sigma^2 s)}\{e^{-st_d}\left(y_i(t_d) - r_i\right);\mathcal{T}\} \\
&= \mu_n^{(\sigma^2 s)}\{e^{-st_d}y_i(t_d);\mathcal{T}\} - r_i\,\mu_n^{(\sigma^2 s)}\{e^{-st};\mathcal{T}\} \\
&\approx \mu_n^{(\sigma^2 s)}\{e^{-st_d}y_i(t_d);\mathcal{T}\} - \hat{r}_i\,\mu_n^{(\sigma^2 s)}\{e^{-st};\mathcal{T}\},
\end{aligned} \tag{A.69}$$

where $\hat{r}_i$ denotes the estimate of r_i.

REFERENCES

[1] M. Hockel and P. Vaupel, "Tumor hypoxia: definitions and current clinical, biologic, and molecular aspects," *Journal of the National Cancer Institute*, vol. 93, no. 4, pp. 266–276, 2001.

[2] J. C. Forster, W. M. Harriss-Phillips, M. J. Douglass, and E. Bezak, "A review of the development of tumor vasculature and its effects on the tumor microenvironment," *Hypoxia*, pp. 21–32, 2017.

[3] B. S. Sørensen and M. R. Horsman, "Tumor hypoxia: Impact on radiation therapy and molecular pathways," *Frontiers in Oncology*, vol. 10, 2020.

[4] M. R. Horsman and P. Vaupel, "Pathophysiological basis for the formation of the tumor microenvironment," *Frontiers in Oncology*, vol. Volume 6 - 2016, 2016.

[5] Z. Chen, F. Han, Y. Du, H. Shi, and W. Zhou, "Hypoxic microenvironment in cancer: molecular mechanisms and therapeutic interventions," *Signal transduction and targeted therapy*, vol. 8, no. 1, p. 70, 2023.

[6] J. C. Walsh, A. Lebedev, E. Aten, K. Madsen, L. Marciano, and H. C. Kolb, "The clinical importance of assessing tumor hypoxia: relationship of tumor hypoxia to prognosis and therapeutic opportunities," *Antioxidants & redox signaling*, vol. 21, no. 10, pp. 1516–1554, 2014.

[7] I. Gertsenshteyn, M. Giurcanu, P. Vaupel, and H. Halpern, "Biological validation of electron paramagnetic resonance (epr) image oxygen thresholds in tissue," *The Journal of physiology*, vol. 599, no. 6, pp. 1759–1767, 2021.

[8] E. Ortiz-Prado, J. F. Dunn, J. Vasconez, D. Castillo, and G. Viscor, "Partial pressure of oxygen in the human body: a general review," *American journal of blood research*, vol. 9, no. 1, p. 1, 2019.

[9] R. Sullivan, G. C. Pare, L. J. Frederiksen, G. L. Semenza, and C. H. Graham, "Hypoxia-induced resistance to anticancer drugs is associated with decreased senescence and requires hypoxia-inducible factor-1 activity," *Molecular cancer therapeutics*, vol. 7, no. 7, pp. 1961–1973, 2008.

[10] G. L. Semenza, "Targeting hif-1 for cancer therapy," *Nature reviews cancer*, vol. 3, no. 10, pp. 721–732, 2003.

[11] S. V. Karakashev and M. J. Reginato, "Progress toward overcoming hypoxia-induced resistance to solid tumor therapy," *Cancer management and research*, pp. 253–264, 2015.

[12] A. M. Shannon, D. J. Bouchier-Hayes, C. M. Condron, and D. Toomey, "Tumour hypoxia, chemotherapeutic resistance and hypoxia-related therapies," *Cancer treatment reviews*, vol. 29, no. 4, pp. 297–307, 2003.

[13] R. M. Hill, S. Rocha, and J. L. Parsons, "Overcoming the impact of hypoxia in driving radiotherapy resistance in head and neck squamous cell carcinoma," *Cancers*, vol. 14, no. 17, p. 4130, 2022.

[14] E. J. Hall, A. J. Giaccia *et al.*, "Radiobiology for the radiologist," *Int J Radiat Oncol Biol Phys*, vol. 66, no. 627, pp. 10–1016, 2006.

[15] INTERNATIONAL ATOMIC ENERGY AGENCY, *Nuclear Medicine Physics*, ser. Non-serial Publications. Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2015. [Online]. Available: http://www.iaea.org/publications/10368/nuclear-medicine-physics

[16] N. Petoussi-Henss, W. Bolch, K. Eckerman, A. Endo, N. Hertel, J. Hunt, M. Pelliccioni, H. Schlattl, and M. Zankl, "Conversion coefficients for radiological protection quantities for external radiation exposures," *Annals of the ICRP*, vol. 40, no. 2-5, pp. 1–257, 2010.

[17] G. Chen, K. Wu, H. Li, D. Xia, and T. He, "Role of hypoxia in the tumor microenvironment and targeted therapy," *Frontiers in oncology*, vol. 12, p. 961637, 2022.

[18] B. Wang, Q. Zhao, Y. Zhang, Z. Liu, Z. Zheng, S. Liu, L. Meng, Y. Xin, and X. Jiang, "Targeting hypoxia in the tumor microenvironment: a potential strategy to improve cancer immunotherapy," *Journal of Experimental & Clinical Cancer Research*, vol. 40, no. 1, p. 24, 2021.

[19] S. Kizaka-Kondoh, M. Inoue, H. Harada, and M. Hiraoka, "Tumor hypoxia: a target for selective cancer therapy," *Cancer science*, vol. 94, no. 12, pp. 1021–1028, 2003.

[20] A. Patel and S. Sant, "Hypoxic tumor microenvironment: Opportunities to develop targeted therapies," *Biotechnology advances*, vol. 34, no. 5, pp. 803–812, 2016.

[21] D. C. Singleton, A. Macann, and W. R. Wilson, "Therapeutic targeting of the hypoxic tumour microenvironment," *Nature reviews Clinical oncology*, vol. 18, no. 12, pp. 751–772, 2021.

[22] X. Yuan, W. Ruan, B. Bobrow, P. Carmeliet, and H. K. Eltzschig, "Targeting hypoxia-inducible factors: therapeutic opportunities and challenges," *Nature Reviews Drug Discovery*, vol. 23, no. 3, pp. 175–200, 2024.

[23] S. Webb, *Intensity-modulated radiation therapy*. CRC Press, 2015.

[24] B. Epel, M. C. Maggio, E. D. Barth, R. C. Miller, C. A. Pelizzari, M. Krzykawska-Serda, S. V. Sundramoorthy, B. Aydogan, R. R. Weichselbaum, V. M. Tormyshev *et al.*, "Oxygen-guided radiation therapy," *International Journal of Radiation Oncology* Biology* Physics*, vol. 103, no. 4, pp. 977–984, 2019.

[25] I. Gertsenshteyn, B. Epel, M. Giurcanu, E. Barth, J. Lukens, K. Hall, J. F. Martinez, M. Grana, M. Maggio, R. C. Miller *et al.*, "Absolute oxygen-guided radiation therapy

improves tumor control in three preclinical tumor models," *Frontiers in Medicine*, vol. 10, p. 1269689, 2023.

[26] M. R. Horsman, L. S. Mortensen, J. B. Petersen, M. Busk, and J. Overgaard, "Imaging hypoxia to improve radiotherapy outcome," *Nature reviews Clinical oncology*, vol. 9, no. 12, pp. 674–687, 2012.

[27] D. B. Cater and I. A. Silver, "Quantitative measurements of oxygen tension in normal tissues and in the tumours of patients before and after radiotherapy," *Acta Radiologica*, no. 3, pp. 233–256, 1960.

[28] J. A. Raleigh, M. W. Dewhirst, and D. E. Thrall, "Measuring tumor hypoxia," in *Seminars in radiation oncology*, vol. 6, no. 1. Elsevier, 1996, pp. 37–45.

[29] H. Lyng, K. Sundfør, and E. K. Rofstad, "Oxygen tension in human tumours measured with polarographic needle electrodes and its relationship to vascular density, necrosis and hypoxia," *Radiotherapy and oncology*, vol. 44, no. 2, pp. 163–169, 1997.

[30] P. Vaupel, M. Höckel, and A. Mayer, "Detection and characterization of tumor hypoxia using po2 histography," *Antioxidants & redox signaling*, vol. 9, no. 8, pp. 1221–1236, 2007.

[31] J. P. O'Connor, J. K. Boult, Y. Jamin, M. Babur, K. G. Finegan, K. J. Williams, R. A. Little, A. Jackson, G. J. Parker, A. R. Reynolds *et al.*, "Oxygen-enhanced mri accurately identifies, quantifies, and maps tumor hypoxia in preclinical cancer models," *Cancer research*, vol. 76, no. 4, pp. 787–795, 2016.

[32] K. A. Krohn, J. M. Link, and R. P. Mason, "Molecular imaging of hypoxia," *Journal of Nuclear Medicine*, vol. 49, no. Suppl 2, pp. 129S–148S, 2008.

[33] B. Epel, G. Redler, and H. J. Halpern, "How in vivo epr measures and images oxygen," *Oxygen Transport to Tissue XXXVI*, pp. 113–119, 2014.

[34] B. Epel, M. K. Bowman, C. Mailer, and H. J. Halpern, "Absolute oxygen r1e imaging in vivo with pulse electron paramagnetic resonance," *Magnetic Resonance in Medicine*, vol. 72, no. 2, pp. 362–368, 2014.

[35] S. R. Cherry, J. A. Sorenson, M. E. Phelps, and S. O. service), *Physics in nuclear medicine*. Saunders, 2003, vol. 3.

[36] N. P. Johnson, S. Sdringola, and K. L. Gould, "Partial volume correction incorporating rb-82 positron range for quantitative myocardial perfusion pet based on systolic-diastolic activity ratios and phantom measurements," *Journal of nuclear cardiology*, vol. 18, no. 2, pp. 247–258, 2011.

[37] S. T. Lee and A. M. Scott, "Hypoxia positron emission tomography imaging with 18f-fluoromisonidazole," in *Seminars in nuclear medicine*, vol. 37, no. 6. Elsevier, 2007, pp. 451–461.

[38] Y. Masaki, Y. Shimizu, T. Yoshioka, Y. Tanaka, K.-i. Nishijima, S. Zhao, K. Higashino, S. Sakamoto, Y. Numata, Y. Yamaguchi *et al.*, "The accumulation mechanism of the hypoxia imaging probe "fmiso" by imaging mass spectrometry: possible involvement of low-molecular metabolites," *Scientific Reports*, vol. 5, no. 1, p. 16802, 2015.

[39] Y. Masaki, Y. Shimizu, T. Yoshioka, K.-i. Nishijima, S. Zhao, K. Higashino, Y. Numata, N. Tamaki, and Y. Kuge, "Fmiso accumulation in tumor is dependent on glutathione conjugation capacity in addition to hypoxic state," *Annals of Nuclear Medicine*, vol. 31, no. 8, pp. 596–604, 2017.

[40] S. S. Foo, D. F. Abbott, N. Lawrentschuk, and A. M. Scott, "Functional imaging of intratumoral hypoxia," *Molecular Imaging & Biology*, vol. 6, no. 5, pp. 291–305, 2004.

[41] S. Zschaeck, S. Löck, F. Hofheinz, D. Zips, L. S. Mortensen, K. Zöphel, E. G. Troost, S. Boeke, M. Saksø, D. Mönnich *et al.*, "Individual patient data meta-analysis of fmiso and faza hypoxia pet scans from head and neck cancer patients undergoing definitive radio-chemotherapy," *Radiotherapy and Oncology*, vol. 149, pp. 189–196, 2020.

[42] K. Sambasivan, S. F. Barrington, S. E. Connor, T. H. Witney, P. J. Blower, and T. G. Urbano, "Is there a role for [18f]-fmiso pet to guide dose adaptive radiotherapy in head and neck cancer? a review of the literature," *Clinical and translational imaging*, vol. 12, no. 2, pp. 137–155, 2024.

[43] S. Welz, F. Paulsen, C. Pfannenberg, M. Reimold, G. Reischl, K. Nikolaou, C. La Fougere, M. Alber, C. Belka, D. Zips *et al.*, "Dose escalation to hypoxic subvolumes in head and neck cancer: A randomized phase ii study using dynamic [18f] fmiso pet/ct," *Radiotherapy and Oncology*, vol. 171, pp. 30–36, 2022.

[44] P. Vera, S. Thureau, P. Chaumet-Riffaud, R. Modzelewski, P. Bohn, M. Vermandel, S. Hapdey, A. Pallardy, M.-A. Mahé, M. Lacombe *et al.*, "Phase ii study of a radiotherapy total dose increase in hypoxic lesions identified by 18f-misonidazole pet/ct in patients with non–small cell lung carcinoma (rtep5 study)," *Journal of Nuclear Medicine*, vol. 58, no. 7, pp. 1045–1053, 2017.

[45] J. R. Brender, Y. Saida, N. Devasahayam, M. C. Krishna, and S. Kishimoto, "Hypoxia imaging as a guide for hypoxia-modulated and hypoxia-activated therapy," *Antioxidants & Redox Signaling*, vol. 36, no. 1-3, pp. 144–159, 2022.

[46] M. Piert, H.-J. Machulla, M. Picchio, G. Reischl, S. Ziegler, P. Kumar, H.-J. Wester, R. Beck, A. J. McEwan, L. I. Wiebe *et al.*, "Hypoxia-specific tumor imaging with 18f-fluoroazomycin arabinoside," *Journal of Nuclear Medicine*, vol. 46, no. 1, pp. 106–113, 2005.

[47] N. Quartuccio, R. Laudicella, P. Mapelli, P. Guglielmo, D. A. Pizzuto, M. Boero, G. Arnone, M. Picchio, and Y. A. W. Group, "Hypoxia pet imaging beyond 18f-fmiso in patients with high-grade glioma: 18f-faza and other hypoxia radiotracers," *Clinical and Translational Imaging*, vol. 8, no. 1, pp. 11–20, 2020.

[48] P. Gouel, P. Decazes, P. Vera, I. Gardin, S. Thureau, and P. Bohn, "Advances in pet and mri imaging of tumor hypoxia," *Frontiers in Medicine*, vol. 10, p. 1055062, 2023.

[49] R. C. Perez, D. Kim, A. W. Maxwell, and J. C. Camacho, "Functional imaging of hypoxia: Pet and mri," *Cancers*, vol. 15, no. 13, p. 3336, 2023.

[50] Y. Huang, J. Fan, Y. Li, S. Fu, Y. Chen, and J. Wu, "Imaging of tumor hypoxia with radionuclide-labeled tracers for pet," *Frontiers in Oncology*, vol. 11, p. 731503, 2021.

[51] I. N. Fleming, R. Manavaki, P. J. Blower, C. West, K. J. Williams, A. L. Harris, J. Domarkas, S. Lord, C. Baldry, and F. J. Gilbert, "Imaging tumour hypoxia with positron emission tomography," *British journal of cancer*, vol. 112, no. 2, pp. 238–250, 2015.

[52] A. E. Ruark, "Positronium," *Physical Review*, vol. 68, no. 11-12, p. 278, 1945.

[53] W. Brandt, S. Berko, and W. W. Walker, "Positronium decay in molecular substances," *Physical Review*, vol. 120, no. 4, p. 1289, 1960.

[54] D. W. Gidley, H.-G. Peng, and R. S. Vallery, "Positron annihilation as a method to characterize porous materials," *Annu. Rev. Mater. Res.*, vol. 36, no. 1, pp. 49–79, 2006.

[55] P. Moskal, "Positronium imaging," in *2019 IEEE nuclear science symposium and medical imaging conference (NSS/MIC)*. IEEE, 2019, pp. 1–3.

[56] P. Moskal, K. Dulski, N. Chug, C. Curceanu, E. Czerwiński, M. Dadgar, J. Gajewski, A. Gajos, G. Grudzień, B. C. Hiesmayr *et al.*, "Positronium imaging with the novel multiphoton pet scanner," *Science Advances*, vol. 7, no. 42, p. eabh4394, 2021.

[57] D. B. Cassidy, "Experimental progress in positronium laser physics," *The European Physical Journal D*, vol. 72, pp. 1–72, 2018.

[58] S. D. Bass, S. Mariazzi, P. Moskal, and E. Stępień, "Colloquium: Positronium physics and biomedical applications," *Reviews of Modern Physics*, vol. 95, no. 2, p. 021002, 2023.

[59] K. Shibuya, H. Saito, F. Nishikido, M. Takahashi, and T. Yamaya, "Oxygen sensing ability of positronium atom for tumor hypoxia imaging," *Communications Physics*, vol. 3, no. 1, p. 173, 2020.

[60] S. V. Stepanov, A. V. Bokov, O. Ilyukhina, and V. M. Byakov, "Dissolved oxygen and positronium atom in liquid media," *Radioelektron. Nanosyst. Inf. Tehnol*, vol. 12, pp. 107–114, 2020.

[61] P. Moskal and E. Ł. Stępień, "Positronium as a biomarker of hypoxia," *Bio-Algorithms and Med-Systems*, vol. 17, no. 4, pp. 311–319, 2021.

[62] B. Jasińska, B. Zgardzińska, G. Chołubek, M. Pietrow, M. Gorgol, K. Wiktor, K. Wysogląd, P. Białas, C. Curceanu, E. Czerwiński *et al.*, "Human tissue investigations using pals technique: free radicals influence," *Acta Physica Polonica A*, 2017.

[63] B. Zgardzińska, G. Chołubek, B. Jarosz, K. Wysogląd, M. Gorgol, M. Goździuk, M. Chołubek, and B. Jasińska, "Studies on healthy and neoplastic tissues using positron annihilation lifetime spectroscopy and focused histopathological imaging," *Scientific Reports*, vol. 10, no. 1, p. 11890, 2020.

[64] A. V. Avachat, K. H. Mahmoud, A. G. Leja, J. J. Xu, M. A. Anastasio, M. Sivaguru, and A. Di Fulvio, "Ortho-positronium lifetime for soft-tissue classification," *Scientific reports*, vol. 14, no. 1, p. 21155, 2024.

[65] P. Moskal, J. Baran, S. Bass, J. Choiński, N. Chug, C. Curceanu, E. Czerwiński, M. Dadgar, M. Das, K. Dulski *et al.*, "Positronium image of the human brain in vivo," *Science Advances*, vol. 10, no. 37, p. eadp2840, 2024.

[66] P. Moskal, D. Kisielewska, C. Curceanu, E. Czerwiński, K. Dulski, A. Gajos, M. Gorgol, B. Hiesmayr, B. Jasińska, K. Kacprzak *et al.*, "Feasibility study of the positronium imaging with the j-pet tomograph," *Physics in Medicine & Biology*, vol. 64, no. 5, p. 055017, 2019.

[67] M. Sitarz, J.-P. Cussonneau, T. Matulewicz, and F. Haddad, "Radionuclide candidates for $\beta+$ γ coincidence pet: an overview," *Applied Radiation and Isotopes*, vol. 155, p. 108898, 2020.

[68] W. M. Steinberger, L. Mercolli, J. Breuer, H. Sari, S. Parzych, S. Niedzwiecki, G. Lapkiewicz, P. Moskal, E. Stepien, A. Rominger *et al.*, "Positronium lifetime validation measurements using a long-axial field-of-view positron emission tomography scanner," *EJNMMI physics*, vol. 11, no. 1, p. 76, 2024.

[69] L. Mercolli, W. M. Steinberger, H. Sari, A. Afshar-Oromieh, F. Caobelli, M. Conti, Â. R. Felgosa Cardoso, C. Mingels, P. Moskal, T. Pyka *et al.*, "In vivo positronium lifetime measurements with a long axial field-of-view pet/ct," *medRxiv*, pp. 2024–10, 2024.

[70] P. Moskal, A. Bilewicz, M. Das, B. Huang, A. Khreptak, S. Parzych, J. Qi, A. Rominger, R. Seifert, S. Sharma *et al.*, "Positronium imaging: history, current status, and future perspectives," *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2025.

[71] B. Huang, T. Li, G. Ariño-Estrada, K. Dulski, R. Y. Shopa, P. Moskal, E. Stępień, and J. Qi, "Split: Statistical positronium lifetime image reconstruction via time-thresholding," *IEEE transactions on medical imaging*, vol. 43, no. 6, pp. 2148–2158, 2024.

[72] B. Huang, Z. Wang, X. Zeng, A. H. Goldan, and J. Qi, "Fast high-resolution lifetime image reconstruction for positron lifetime tomography," *Communications Physics*, vol. 8, no. 1, p. 181, 2025.

[73] B. Huang and J. Qi, "High-resolution positronium lifetime tomography by the method of moments," *Physics in Medicine & Biology*, vol. 69, no. 24, p. 24NT01, 2024.

[74] L. Berens, I. Hsu, C.-T. Chen, H. Halpern, and C.-M. Kao, "An analytic, moment-based method to estimate orthopositronium lifetimes in positron annihilation lifetime (PAL) measurements," *Bio-Algorithms and Med-Systems*, vol. 20(special issue), pp. 40–48, 2024.

[75] K. Wang, "Feasibility of high spatial resolution working modes for clinical pet scanner," *International Journal of Medical Physics, Clinical Engineering and Radiation Oncology*, vol. 7, no. 4, pp. 539–552, 2018.

[76] K. Shibuya, H. Saito, H. Tashima, and T. Yamaya, "Using inverse laplace transform in positronium lifetime imaging," *Physics in Medicine & Biology*, vol. 67, no. 2, p. 025009, 2022.

[77] M. Conti, "Improving time resolution in time-of-flight pet," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 648, pp. S194–S198, 2011.

[78] F. M. Muller, M. E. Daube-Witherspoon, E. J. Li, A. R. Pantel, and J. S. Karp, "Realization of long axial field-of-view pet technology in the clinic and research environment," *British Journal of Radiology*, p. tqag096, 2026.

[79] S. Parzych, S. Niedźwiecki, E. Y. Beyene, N. Chug, M. Conti, C. Curceanu, E. Czerwiński, M. Das, K. Valsan Eliyan, J. Hajduga *et al.*, "Feasibility simulation study of positronium lifetime imaging with the biograph vision quadra and j-pet tomographs," *EJNMMI physics*, 2026.

[80] Z. Chen, C.-M. Kao, H.-H. Huang, and L. An, "Enhanced positronium lifetime imaging through two-component reconstruction in time-of-flight positron emission tomography," *Frontiers in physics*, vol. 12, p. 1429344, 2024.

[81] H.-H. Huang, Z. Zhu, S. Booppasiri, Z. Chen, S. Pang, and C.-M. Kao, "A statistical reconstruction algorithm for positronium lifetime imaging using time-of-flight positron emission tomography," *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2025.

[82] P. Massa and F. Benvenuto, "Predictive risk estimation for the expectation maximization algorithm with poisson data," *Inverse Problems*, vol. 37, no. 4, p. 045013, 2021.

[83] M. Qutbi, "On the convergence of maximum likelihood expectation–maximization algorithm for iterative tomographic reconstruction and the role of spatial frequency,

location, pixel intensity, contrast, and noise," *Journal of Medical Physics*, vol. 50, no. 2, pp. 348–358, 2025.

[84] J. Qi and B. Huang, "Positronium lifetime image reconstruction for tof pet," *IEEE transactions on medical imaging*, vol. 41, no. 10, pp. 2848–2855, 2022.

[85] P. Moskal, D. Kisielewska, R. Y. Shopa, Z. Bura, J. Chhokar, C. Curceanu, E. Czerwiński, M. Dadgar, K. Dulski, J. Gajewski *et al.*, "Performance assessment of the 2 γ positronium imaging with the total-body pet scanners," *EJNMMI physics*, vol. 7, no. 1, p. 44, 2020.

[86] L. Mercolli, W. M. Steinberger, P. V. Grundler, A. Moiseeva, S. Braccini, M. Conti, P. Moskal, N. Rathod, A. Rominger, H. Sari *et al.*, "First positronium lifetime imaging with scandium-44 on a long axial field-of-view pet/ct," *Frontiers in Nuclear Medicine*, vol. 5, p. 1648621, 2025.

[87] B. Huang, B. Dai, S. E. Lapi, G. Liles, J. S. Karp, and J. Qi, "High-resolution positronium lifetime tomography at clinical activity levels on the pennpet explorer," *Journal of Nuclear Medicine*, vol. 66, no. 9, pp. 1464–1470, 2025.

[88] H. H. Barrett, D. W. Wilson, and B. M. Tsui, "Noise properties of the EM algorithm. I. Theory," *Physics in Medicine & Biology*, vol. 39, no. 5, pp. 833–846, 1994.

[89] F. Hashimoto, Y. Onishi, K. Ote, H. Tashima, A. J. Reader, and T. Yamaya, "Deep learning-based pet image denoising and reconstruction: a review," *Radiological physics and technology*, vol. 17, no. 1, pp. 24–46, 2024.

[90] A. Brahme, J.-E. Roos, and I. Lax, "Solution of an integral equation encountered in rotation therapy," *Physics in Medicine & Biology*, vol. 27, no. 10, pp. 1221–1229, 1982.

[91] T. Bortfeld, "Imrt: a review and preview," *Physics in Medicine & Biology*, vol. 51, no. 13, pp. R363–R379, 2006.

[92] B. Epel, M. Krzykawska-Serda, V. Tormyshev, M. C. Maggio, E. D. Barth, C. A. Pelizzari, and H. J. Halpern, "Spin lattice relaxation epr po2 images may direct the location of radiation tumor boosts to enhance tumor cure," *Cell Biochemistry and Biophysics*, vol. 75, no. 3, pp. 295–298, 2017.

[93] B. Epel, M. Maggio, C. Pelizzari, and H. J. Halpern, "Electron paramagnetic resonance po2 image tumor oxygen-guided radiation therapy optimization," in *Oxygen Transport to Tissue XXXIX*. Springer, 2017, pp. 287–296.

[94] J.-L. Lim and M. S. Berridge, "An efficient radiosynthesis of [18f] fluoromisonidazole," *Applied radiation and isotopes*, vol. 44, no. 8, pp. 1085–1091, 1993.

[95] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[96] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[97] J. Ardenkjaer-Larsen, I. Laursen, I. Leunbach, G. Ehnholm, L.-G. Wistrand, J. Petersson, and K. Golman, "Epr and dnp properties of certain novel single electron contrast agents intended for oximetric imaging," *Journal of Magnetic Resonance*, vol. 133, no. 1, pp. 1–12, 1998.

[98] R. Ahmad and P. Kuppusamy, "Theory, instrumentation, and applications of electron paramagnetic resonance oximetry," *Chemical reviews*, vol. 110, no. 5, pp. 3212–3236, 2010.

[99] T. J. Reddy, T. Iwama, H. J. Halpern, and V. H. Rawal, "General synthesis of persistent trityl radicals for epr imaging of biological systems," *The Journal of Organic Chemistry*, vol. 67, no. 14, pp. 4635–4639, 2002.

[100] H. Kim, B. Epel, S. Sundramoorthy, H.-M. Tsai, E. Barth, I. Gertsenshteyn, H. Halpern, Y. Hua, Q. Xie, C.-T. Chen *et al.*, "Development of a pet/epri combined imaging system for assessing tumor hypoxia," *Journal of Instrumentation*, vol. 16, no. 03, p. P03031, 2021.

[101] H. Kim, Y. Hua, C.-T. Chen, B. Epel, S. Sundramoorthy, G.-B. Ko, J. Lee, H. Halpern, and C.-M. Kao, "Development of a pet subsystem for next generation preclinical pet/epr for studying tumor hypoxia," in *2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTSD)*. IEEE, 2023, pp. 1–1.

[102] H. Kim, Y. Hua, C.-T. Chen, B. Epel, S. Sundramoorthy, H. Halpern, and C.-M. Kao, "Validation of the design of a high-sensitivity and high-resolution pet system for a preclinical pet/epr hybrid scanner," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1063, p. 169333, 2024.

[103] G. Redler, E. Pearson, X. Liu, I. Gertsenshteyn, B. Epel, C. Pelizzari, B. Aydogan, R. Weichselbaum, H. J. Halpern, and R. D. Wiersma, "Small animal imrt using 3d-printed compensators," *International Journal of Radiation Oncology* Biology* Physics*, vol. 110, no. 2, pp. 551–565, 2021.

[104] I. Gertsenshteyn, B. Epel, A. Ahluwalia, H. Kim, X. Fan, E. Barth, M. Zamora, E. Markiewicz, H.-M. Tsai, S. Sundramoorthy *et al.*, "The optimal 18f-fluoromisonidazole pet threshold to define tumor hypoxia in preclinical squamous cell carcinomas using po2 electron paramagnetic resonance imaging as reference truth,"

European journal of nuclear medicine and molecular imaging, vol. 49, no. 12, pp. 4014–4024, 2022.

[105] I. H. Gertsenshteyn, "Multi-modal imaging of tumor hypoxia to improve radiation therapy," Ph.D. dissertation, The University of Chicago, 2022.

[106] I. Gertsenshteyn, B. Epel, E. Barth, L. Leoni, E. Markiewicz, H.-M. Tsai, X. Fan, M. Giurcanu, D. Bodero, M. Zamora *et al.*, "improving tumor hypoxia location in 18f-misonidazole pet with dynamic contrast-enhanced mri using quantitative electron paramagnetic resonance partial oxygen pressure images," *Radiology: Imaging Cancer*, vol. 3, no. 2, p. e200104, 2021.

[107] B. Epel, G. Redler, V. Tormyshev, and H. J. Halpern, "Towards human oxygen images with electron paramagnetic resonance imaging," in *Oxygen Transport to Tissue XXXVII*. Springer, 2016, pp. 363–369.

[108] X. Li, W. Huang, and J. H. Holmes, "Dynamic contrast-enhanced (dce) mri," *Magnetic Resonance Imaging Clinics*, vol. 32, no. 1, pp. 47–61, 2024.

[109] U. Simoncic, S. Leibfarth, S. Welz, N. Schwenzer, H. Schmidt, G. Reischl, C. Pfannenberg, C. l. Fougère, K. Nikolaou, D. Zips *et al.*, "Comparison of dce-mri kinetic parameters and fmiso-pet uptake parameters in head and neck cancer patients," *Medical physics*, vol. 44, no. 6, pp. 2358–2368, 2017.

[110] J. S. Rasey, J. J. Casciari, P. D. Hofstrand, M. Muzi, M. M. Graham, and L. K. Chin, "Determining hypoxic fraction in a rat glioma by uptake of radiolabeled fluoromisonidazole," *Radiation research*, vol. 153, no. 1, pp. 84–92, 2000.

[111] M. Elas, K.-H. Ahn, A. Parasca, E. D. Barth, D. Lee, C. Haney, and H. J. Halpern, "Electron paramagnetic resonance oxygen images correlate spatially and quantitatively with oxylite oxygen measurements," *Clinical Cancer Research*, vol. 12, no. 14, pp. 4209–4217, 2006.

[112] D. A. Kovar, M. Lewis, and G. S. Karczmar, "A new method for imaging perfusion and contrast extraction fraction: input functions derived from reference tissues," *Journal of Magnetic Resonance Imaging*, vol. 8, no. 5, pp. 1126–1134, 1998.

[113] R. Polikar, "Bootstrap-inspired techniques in computation intelligence," *IEEE signal processing magazine*, vol. 24, no. 4, pp. 59–72, 2007.

[114] P. Moskal and E. Ł. Stępień, "Perspectives on translation of positronium imaging into clinics," *Frontiers in Physics*, vol. 10, p. 969806, 2022.

[115] P. Moskal and E. Stępień, "Prospects and clinical perspectives of total-body pet imaging using plastic scintillators," *PET clinics*, vol. 15, no. 4, pp. 439–452, 2020.

[116] T. Matulewicz, "Radioactive nuclei for $\beta+$ γ pet and theranostics: selected candidates," *Bio-Algorithms and Med-Systems*, vol. 17, no. 4, pp. 235–239, 2022.

[117] S. Chu, "The lund/lbnl nuclear data search," *http://nucleardata. nuclear. lu. se/nucleardata/toi/index. asp*, 1999.

[118] M. T. Durán, F. Juget, Y. Nedjadi, C. Bailat, P. V. Grundler, Z. Talip, N. P. van der Meulen, P. Casolaro, G. Dellepiane, and S. Braccini, "Half-life measurement of 44sc and 44msc," *Applied Radiation and Isotopes*, vol. 190, p. 110507, 2022.

[119] G. Crișan, N. S. Moldovean-Cioroianu, D.-G. Timaru, G. Andrieș, C. Căinap, and V. Chiș, "Radiopharmaceuticals for pet and spect imaging: a literature review over the last decade," *International journal of molecular sciences*, vol. 23, no. 9, p. 5023, 2022.

[120] J. P. Meier, H. J. Zhang, R. Freifelder, M. Bhuiyan, P. Selman, M. Mendez, P. H. Kankanamalage, T. Brossard, A. Pusateri, H.-M. Tsai *et al.*, "Accelerator-based production of scandium radioisotopes for applications in prostate cancer: Toward building a pipeline for rapid development of novel theranostics," *Molecules*, vol. 28, no. 16, p. 6041, 2023.

[121] C. Abate-Shen and F. Nunes de Almeida, "Establishment of the lncap cell line–the dawn of an era for prostate cancer research," *Cancer research*, vol. 82, no. 9, pp. 1689–1691, 2022.

[122] E. Grushka, "Characterization of exponentially modified gaussian peaks in chromatography," *Analytical chemistry*, vol. 44, no. 11, pp. 1733–1738, 1972.