# Improved Quantification for Respiratory Gated PET/CT: Data-Driven Algorithms for Respiratory Motion Correction in PET/CT

### Alexander Charles Whitehead

Supervisors:

Prof. Kris Thielemans

Dr Jamie R. McClelland

A thesis submitted in partial fulfillment of the requirements for the degree of

**Doctor of Philosophy** 

of

Medical Imaging Physics and Medical Image Computing

University College London.

Department of Medical Physics and Biomedical Engineering
Institute of Nuclear Medicine
University College London

12th June 2025

I, Alexander Charles Whitehead, confirm that the work presented own. Where information has been derived from other sources, I confirmindicated in the work.	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
own. Where information has been derived from other sources, I confirm	
	m that this ha

In dark times, should the stars also go out?

"Steban, the Student Communist"

Disco Elysium

They asked me how well I understood theoretical physics. I said I

had a theoretical degree in physics. They said welcome aboard.

I push buttons. I turn dials. I read numbers. Sometimes I make up

little stories in my head about what the numbers mean.

I know exactly what I'm doing. I just don't know what effect it's

going to have.

"Fantastic"

Fallout: New Vegas

### **Abstract**

Respiratory motion is a significant general problem in Positron Emission Tomography (PET) imaging, affecting both image quality and quantitative accuracy. Respiratory motion not only blurs lesions and other anatomical features in the lungs, but also complicates the application of attenuation correction to the acquired PET data. While respiratory gating can help reduce motion artefacts, it extends acquisition time and reduces effective counts. A further challenge arises from the temporal mismatch between the attenuation map and the emission scan, as patients are typically instructed to hold their breath during the Computed Tomography (CT) acquisition but breathe freely during the PET scan of a combined PET/CT. As a result, the static Attenuation Map ( $\mu$ -Map) often does not correspond to any specific respiratory phase of the PET data, introducing misalignment that degrades both qualitative and quantitative image accuracy.

Motion correction methods exist which incorporate registration in order to attempt to improve upon non-motion corrected results. Often these methods involve separating the PET data into bins, where the respiratory motion is minimal within each bin. However, because of the high level of noise and low spatial resolution of PET data, few bins are often used, leading to a not insignificant amount of respiratory motion still being present in the resultant images. Additionally, this approach does not solve the mismatch of the  $\mu$ -Map. Logically, the bin closest to the  $\mu$ -Map could be used as the reference bin for registration but it is not guaranteed that this bin will be close to the position of the  $\mu$ -Map and as such artefacts will remain in the final image. Furthermore, registration fails when used on dynamic PET data, where the signal from the aorta, at early time points, often leads to mis-registration. More complex motion correction methods exist, however, these methods, in general, tend to be more resource intensive in both the sense of computation time and computational resources.

This thesis focuses on the development of a motion correction method, which seeks to rectify some of the issues above by using respiratory gated data in combination with motion modelling. Firstly, this thesis presents preliminary results, where the bounds of the problem were found. This includes experiments to discover the effectiveness of different motion correction techniques (focusing on motion modelling) in the case of Time of Flight (TOF) vs Non-Time of Flight (Non-TOF) PET data, especially where a high number of bins are used. Then, the thesis explores work related to attempting to solve the mismatch of the  $\mu$ -Map by deforming it to the position during the gates. Furthermore, a comparison between the effects of the reconstruction and Motion Model (MM) fitting process is presented, including using Maximum Likelihood Activity and Attenuation Correction Factors Estimation (MLACF), to approximate attenuation correction, as well as fitting a MM on coarsely binned data and applying it to finely binned data.

Finally, this thesis also presents work on the problems of Data Driven (DD) Surrogate Signal (SS) extraction methods applied to dynamic PET. A SS is imperative to the effectiveness of binning data as well as to MM fitting. By presenting work relevant to SS extraction from dynamic PET this work potentially opens the motion correction methods presented previously to the application of dynamic PET.

The thesis concludes by critically reflecting on the work presented, highlighting both methodological advancements and areas for further refinement. By outlining key future directions, it sets the stage for continued development of clinically viable respiratory motion correction approaches for both static and dynamic PET imaging.

## **Impact Statement**

The work presented in this thesis constitutes a methodological advancement in the domain of respiratory motion correction for PET/CT imaging. The proposed framework was developed to address key limitations of existing techniques, particularly the challenge of aligning attenuation maps with dynamic anatomical motion, and the difficulty of applying motion correction in dynamic PET acquisitions, where conventional gating is often infeasible.

By focusing on DD solutions that operate on data already available in standard clinical protocols, this research aims to enable motion correction that is both technically robust and clinically practical. The integration of diffeomorphic registration, motion modelling, warping to a  $\mu$ -Map reference position, and SS extraction directly from (dynamic) PET data represents a novel contribution at the interface between motion correction and reconstruction.

While the methods developed are currently at the stage of technical and simulation-based validation, they are designed with eventual clinical integration in mind. The SS extraction strategies, in particular, offer a viable path toward extending respiratory motion correction to dynamic PET, a domain where limited signal and variable tracer kinetics have historically impeded progress.

From an academic perspective, the research has been conducted with a strong commitment to transparency, reproducibility, and open scientific practice. The work has been disseminated through international conferences and peer-reviewed publications, and future research directions have been clearly out-

lined, including avenues for clinical validation and integration into joint reconstruction frameworks.

In this context, the impact of the thesis lies in providing a solid methodological foundation for future clinical applications, while contributing to the broader academic discourse on motion correction in medical imaging.

## **Acknowledgements**

This research was supported by General Electric (GE) Healthcare, the National Institute for Health Research (NIHR) University College London Hospitals (UCLH) Biomedical Research Centre and the University College London (UCL) Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training (CDT) in Intelligent, Integrated Imaging in Healthcare (i4health) grant (EP/L016478/1).

The software used was partly produced by the Collaborative Computational Project (CCP) on Synergistic Biomedical Imaging (SyneRBI) United Kingdom (UK) EPSRC grant (EP/T026693/1).

Firstly, I would like to thank my supervisors Kris Thielemans and Jamie McClelland. I am eternally grateful for the opportunity that you have provided to me, both to do a Doctor of Philosophy (PhD) at all (especially at UCL) as well as all of the support I have received throughout the (many) years that it has taken me to get to this point. I am most thankful for how patient you have been with me. It must have been frustrating for you, at times, but I hope you also got something out of it. I hope I didn't disappoint you.

Next, I would like to thank all the people at GE HealthCare for their insightful comments and help in formulating ideas, including Dylan (I hope you continue to enjoy basketball as much as you did last time when your team won), Scott Wollenweber, and Chuck Stearns.

Potentially most importantly (sorry), I would like to write something about my dad. One of the reasons that I wanted to do a PhD at all was because I wanted you to be proud of me. One of the reasons I did anything at all is because I wanted you to be proud of me. You didn't make it to the end of my PhD and I have no idea what you would think now. On one hand I don't think you would have really cared that I did a PhD, not because you don't care at all but because I didn't need to do a PhD for you to be happy with what I am

doing. I almost made the mistake of guessing what you wanted me to do while I was in college, luckily that sorted itself out quite well. I'm going to miss thinking about the fact that on the day I (hopefully) finally pass all of this you would have been so emotional, you were so emotional when I finished my undergraduate, never mind something more. I do not believe at all that you are anywhere other than in a box at the bottom of the stairs, or in my ring (or in Sam's ring, which he used as an extremely expensive method of scattering your ashes). I do not believe that you are watching over me or any other rubbish, like some people like to say. However, I do believe that I wouldn't be the person that I am today if you weren't my dad. I loved all the little things that you found interesting, I know way more about pitot tubes than anyone has any right to know. My interests in electronics and engineering are specifically because of you. I'm disappointed that you never tried to do anything more with what you could have obviously done. I think the best thing I can do for you now is to lead a good (and healthy) life. Thank you, I still miss you.

Additionally, I would like to thank my mum, in your own way, I know, that you have have always wanted to show that you love and care for me. I'm thankful for what you've done for me. I know that things have been difficult since dad died, but I hope someday you can begin to be happy again. I really hope that you will always be proud of me. I would like to thank my brother, (little) Sam, I cherish and appreciate the very special relationship we have, you are one of my best friends. I will always be there for you. Finally, I would like to thank my granddad and my grandma. To my granddad, I appreciate how you always ask about and listened to me talking about my work, even though it must have been very boring at times. I wouldn't have been able to finish anything without you having read and commented on my work, which must have been very time consuming. To my grandma, thank you for showing an interest in my field of work, reading about it in magazines or watching it on the news. I really appreciate the effort you have made, especially because I know you have no interest in technology.

I'd like to thank Nikos because without your supervision during my Master of Science (MSc) I would not even have thought it possible to do a PhD at UCL.

I would like to thank all the people I have met during the time I was doing my PhD. First, I'd like to thank Ludovica, you are one of the most loyal people I know. Even though we have had our differences sometimes (mostly caused by my own stupidity),

you've never held it against me and have been (and still are) always there for me, both as a colleague and as one of my best friends. Next, I would like to thank Ashley for the times we have supported each other during our PhDs, especially during the times we both doubted that we would be able to finish this. I especially appreciated and enjoyed being able to talk about working class things we both can relate to. I want to thank Ander, I appreciate all of your input for my work. I will never forget that you were the one who talked to me on the day my dad died. I hope you are happy now. I would like to thank Ade, you were the first person I met at UCL, you showed me around and made me feel welcome and included. You're an incredibly nice and funny guy and I always enjoyed your company. I would like to thank Elise for all the help you gave me at the beginning of my PhD. Finally, I would like to thank John, Marie, and Cindy for all the good times we spent together before the pandemic, I would have been significantly more lonely when I moved to London without you all.

Not to forget, I would like to thank my friends Bousfield (big Sam), Jack, and Peter for simply being my friends. I have very fond memories of the time we have spent together, so far, and I am looking forward to all the times yet to come. If any of you see this, take a photograph, show me and I will buy you a beer, but only the first time. I would also like to thank Viv for helping me out of a sticky situation with my hand drawn figures.

Abschließend möchte ich mich bei Kathi bedanken. Ich würde gerne scherzen, dass es viel einfacher gewesen wäre, meine Doktorarbeit zu beenden, wenn Du nicht bei mir gewesen wärst, weil ich meine ganze Zeit nur mit Dir verbringen möchte. Aber das stimmt nicht (na ja, das mit der Zeit schon). Ich glaube nicht, dass ich meine Arbeit hätte abschließen können, zumindest nicht so glücklich, wie ich es bin, wenn Du nicht gewesen wärst. Obwohl in letzter Zeit so viel passiert ist, findest Du immer noch etwas in Dir, das Dir hilft, jeden Tag besser zu machen als den letzten. Ich glaube wirklich, dass du meine Arbeit öfter gelesen hast als ich selbst. Dafür möchte ich Dir danken. Ich denke, wir können mit Fug und Recht behaupten, dass wir "in guten wie in schlechten Zeiten" und "in Krankheit und Gesundheit" bereits gemeistert haben. Ich bin gespannt auf "bis dass der Tod uns scheidet".

### **List of Publications**

- Alexander C. Whitehead and Ludovica Brusaferri. 'Mortality Maps: Inherently Interpretable Event-Conditional Survival Analysis'. In: arXiv (2025) (In Preparation)
- Ludovica Brusaferri, Alexander C. Whitehead, Lucia Maccioni, Zaynab Alshelh,
  Mattia Veronese, Federico Turkheimer, Nicola Toschi, Matteo Ferrante, Marianna
  Inglese, Ciprian Catana, Enrico Grisan and Marco L. Loggia. 'Advancing Generalisable Neural Network-Based PET Quantification: A Multicenter [11C]PBR28
  study'. In: 47th Annual International Conference of the IEEE Engineering in
  Medicine and Biology Society (EMBC) (EMBC 2025). 2025 (Accepted)
- Alexander C. Whitehead, Kuan Hao Su, Elise C. Emond, Ander Biguri, Ludovica Brusaferri, Maria Machado, Joanna C. Porter, Helen Garthwaite, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Data driven surrogate signal extraction for dynamic PET using selective PCA: time windows versus the combination of components'. In: *Physics in Medicine & Biology* 69.17 (Aug. 2024), p. 175008. ISSN: 0031-9155. DOI: 10.1088/1361-6560/AD5EF1

### Selected as Editor's Choice article of the month!

- Ahmed H. Shahin, An Zhao, Alexander C. Whitehead, Daniel C. Alexander, Joseph Jacob and David Barber. 'CenTime: Event-conditional modelling of censoring in survival analysis'. In: *Medical Image Analysis* 91 (Jan. 2024), p. 103016. ISSN: 13618423. DOI: 10.1016/J.MEDIA.2023.103016
- Matteo Ferrante, Marianna Inglese, Ludovica Brusaferri, Alexander C. Whitehead,
   Lucia Maccioni, Federico E. Turkheimer, Maria A. Nettis, Valeria Mondelli, Oliver
   Howes, Marco L. Loggia, Mattia Veronese and Nicola Toschi. 'Physically informed

deep neural networks for metabolite-corrected plasma input function estimation in dynamic PET imaging'. In: *Computer Methods and Programs in Biomedicine* 256 (Nov. 2024), p. 108375. ISSN: 0169-2607. DOI: 10.1016/J.CMPB.2024. 108375

- (Oral presentation) A. C. Whitehead, L. Brusaferri, L. Maccioni, M. Ferrante, M. Inglese, Z. Alshelh, M. Veronese, N. Toschi, J. Gilman, K. Thielemans and M. L. Loggia. 'A Bayesian Neural Network-Based Method For The Extraction Of A Metabolite Corrected Arterial Input Function From Dynamic [11 C]PBR28 PET '. in: 2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTSD). Institute of Electrical and Electronics Engineers (IEEE), Dec. 2023, pp. 1–1. DOI: 10.1109/NSSMICRTSD49126.2023.10338687
- (Poster presentation) A. C. Whitehead, A. H. Shahin, A. Zhao, D. C. Alexander, J. Jacob and D. Barber. 'Neural Network Based Methods for the Survival Analysis of Idiopathic Pulmonary Fibrosis Patients from a Baseline CT Acquisition'. In: 2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTSD). Institute of Electrical and Electronics Engineers (IEEE), Dec. 2023, pp. 1–1. DOI: 10.1109/NSSMICRTSD49126.2023.10337922
- Evgueni Ovtchinnikov, Richard Brown, Johannes Mayer, Edoardo Pasca, Casper da Costa-Luis, David Atkinson, Christoph Kolbitsch, Nikos Efthimiou, Matthew Strugari, Ashley Gillman, Sam David Porter, Ander Biguri, Daniel Deidda, Alexander Whitehead, Evangelos Papoutsellis, Gemma Fardell, Ben Thomas, Francesca Leek, Matthias Ehrhardt and Kris Thielemans. 'SIRF Synergistic Image Reconstruction Framework'. In: (Jan. 2023). DOI: 10.5281/ZENODO.7551178. URL: https://zenodo.org/record/7551178
- Ludovica Brusaferri, Elise C. Emond, Alexandre Bousse, Robert Twyman, Alexander C. Whitehead, David Atkinson, Sebastien Ourselin, Brian F. Hutton, Simon Arridge and Kris Thielemans. 'Detection Efficiency Modeling and Joint Activity

- and Attenuation Reconstruction in Non-TOF 3-D PET from Multiple-Energy Window Data'. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 6.1 (2022), pp. 87–97. ISSN: 24697311. DOI: 10.1109/TRPMS.2021.3064239
- (Poster presentation) Alexander C. Whitehead, Kuan Hao Su, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'PET/CT Motion Correction Exploiting Motion Models Fit on Coarsely Gated Data Applied to Finely Gated Data'. In: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845.2022.10399154
- (Oral presentation) Alexander C. Whitehead, Kuan Hao Su, Elise C. Emond, Ander Biguri, Maria MacHado, Joanna C. Porter, Helen Garthwaite, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Data Driven Surrogate Signal Extraction for Dynamic PET Using Selective PCA'. in: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845. 2022.10399196
- (Poster presentation) Alexander C. Whitehead, Kjell Erlandsson, Ander Biguri, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Pseudo-Bayesian DIP Denoising as a Preprocessing Step for Kinetic Modelling in Dynamic PET'. in: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845.2022.10399318
- Matteo Ferrante, Marianna Inglese, Ludovica Brusaferri, Alexander Whitehead, Marco Loggia and Nicola Toschi. 'Physically Informed Neural Network for Non-Invasive Arterial Input Function Estimation In Dynamic PET Imaging'. In: June 2022. DOI: 10.1038/s42254-021-00314-5

- Edoardo Pasca, Ben Thomas, Casper da Costa-Luis, Richard Brown, Ander Biguri, Evgueni Ovtchinnikov, Ashley Gillman, Christoph Kolbitsch, Johannes Mayer, Alexander Whitehead, Georg Schramm, David Atkinson and Kris Thielemans. 'SIRF-SuperBuild'. In: (Aug. 2022). DOI: 10.5281/ZENODO.7022530
- (Poster presentation) Alexander C. Whitehead, Ander Biguri, Kuan Hao Su, Scott D. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'Comparison of Motion Correction Methods Incorporating Motion Modelling for PET/CT Using a Single Breath Hold Attenuation Map'. In: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference Record, NSS/MIC 2021 and 28th International Symposium on Room-Temperature Semiconductor Detectors, RTSD 2022. Institute of Electrical and Electronics Engineers Inc., Sept. 2021, pp. 1–4. ISBN: 9781665421133. DOI: 10.1109/NSS/MIC44867. 2021.9875560
- Ander Biguri, Fotis Kotasidis, Alexander C. Whitehead, Irene Burger, Brian F. Hutton and Kris Thielemans. 'Systematic Evaluation of the Impact of Involuntary Motion in Whole Body Dynamic PET'. in: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference Record, NSS/MIC 2021 and 28th International Symposium on Room-Temperature Semiconductor Detectors, RTSD 2022. Institute of Electrical and Electronics Engineers Inc., 2021. ISBN: 9781665421133. DOI: 10.1109/NSS/MIC44867.2021.9875689
- Evgueni Ovtchinnikov, Richard Brown, Kris Thielemans, Edoardo Pasca, Casper da Costa-Luis, Benjamin A. Thomas, David Atkinson, Johannes Mayer, Ashley Gillman, Christoph Kolbitsch, Matthias J. Ehrhardt and Alexander Whitehead. 'SIRF Virtual Machine'. In: (May 2021). DOI: 10.5281/ZENOD0.7050371
- (Oral presentation) Alexander C. Whitehead, Ander Biguri, Nikos Efthimiou, Kuan Hao Su, Scott W. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'PET/CT Respiratory Motion Correction with a Single Attenuation Map Using NAC Derived Deformation Fields'. In: 2020 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC

2020. IEEE, Aug. 2020, pp. 1–3. ISBN: 9781728176932. DOI: 10.1109/NSS/MIC42677.2020.9507890

- (Poster presentation) Alexander C. Whitehead, Elise C. Emond, Nikos Efthimiou, Adeyemi Akintonde, Scott W. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'Impact of Time-of-Flight on Respiratory Motion Modelling using Non-Attenuation-Corrected PET'. in: 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2019. IEEE, Oct. 2019. ISBN: 9781728141640. DOI: 10.1109/NSS/MIC42101.2019. 9059812
- N. Efthimiou, A. C. Whitehead, M. Stockhoff, C. Thyssen, S. J. Archibald and S. Vandenberghe. 'Preliminary investigation of the impact of Axial Ring Splitting on Image Quality for the Cost Reduction of Total-Body PET'. in: 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2019. Institute of Electrical and Electronics Engineers Inc., Oct. 2019. ISBN: 9781728141640. DOI: 10.1109/NSS/MIC42101.2019.9059650

# **Contents**

1	Intr	roduction 5					
	1.1	Introdu	uction		56		
	1.2	Motiva	ation		56		
		1.2.1	Objective	es of this Work	60		
	1.3	Overvi	iew of this	Thesis	60		
2	Bacl	kground	d		63		
	2.1	Introdu	action		63		
	2.2	PET .			63		
		2.2.1	The Phys	sics of PET	64		
			2.2.1.1	Radiotracers	65		
			2.2.1.2	Decay and Annihilation	66		
			2.2.1.3	Static and Dynamic Acquisition	67		
			2.2.1.4	PET Field of View	68		
			2.2.1.5	Attenuation	69		
		2.2.2	Data acq	uisition	71		
			2.2.2.1	Coincidence Processing	73		
			2.2.2.2	Time of Flight PET	73		
			2.2.2.3	Data Output	75		
			2.2.2.4	PET resolution	77		
		2.2.3	Combine	ed PET/CT	78		
			2.2.3.1	Attenuation Correction	79		
	2.3	Inverse	e Problems	s and Optimisation	83		
		231	Inverse F	Problem Concepts	84		

		2.3.2	Optimisa	tion Concepts	4
			2.3.2.1	Objective Function	5
			2.3.2.2	Optimiser	7
	2.4	PET I	mage Reco	nstruction	9
		2.4.1	Introduct	ion9	0
		2.4.2	Iterative 1	Image Reconstruction	0
		2.4.3	MLAA a	nd MLACF Methods for Attenuation and Activity Re-	
			construct	ion in PET 9	2
			2.4.3.1	MLAA	13
			2.4.3.2	MLACF	4
			2.4.3.3	Incorporation of TOF	4
	2.5	Respir	atory Moti	on in PET	4
		2.5.1	Respirato	ory Motion Artefacts	5
		2.5.2	Respirato	ory Motion Challenges in Combined PET/CT Imaging 9	7
	2.6	Motion	n Correctio	n for PET	8
		2.6.1	Image Re	egistration	0
			2.6.1.1	Rigid Deformation	1
			2.6.1.2	Non-Rigid Deformation	13
		2.6.2	Respirato	ory Gating	16
		2.6.3	Respirato	ory Signal Detection	17
			2.6.3.1	External Devices	19
			2.6.3.2	Data Driven	1
		2.6.4	Applying	Image Registration	8
		2.6.5	Motion N	Modelling	9
			2.6.5.1	Respiratory Correspondence Model	1
		2.6.6	Applying	Motion Correction	4
		2.6.7	Summary	of Motion Correction Methods	:5
	2.7	Summ	ary of Cha	llenges and Motivation for This Work	6
•	Г		7. 1. 6		
3		·	Studies for	<b>Motion Correction Utilising Motion Models on Simple</b>	
	Data	a		12	X

3.1	Overvi	iew					
3.2	Introdu	action					
3.3	Impact of TOF on Respiratory Motion Model Estimation Using Pre-Gated						
	No Int	ra-Cycle Motion NAC PET					
	3.3.1	Introduction					
	3.3.2	Methods					
		3.3.2.1 XCAT Image Generation					
		3.3.2.2 PET Data Simulation					
		3.3.2.3 Image Reconstruction					
		3.3.2.4 Motion Model Estimation					
		3.3.2.5 Evaluation					
	3.3.3	Results					
	3.3.4	Discussion					
3.4	PET/CT Respiratory Motion Correction With a Single Attenuation Map						
	Using	NAC Derived Deformation Fields					
	3.4.1	Introduction					
	3.4.2	Methods					
		3.4.2.1 XCAT Volume Generation					
		3.4.2.2 PET Acquisition Simulation					
		3.4.2.3 Non-Attenuation Corrected Image Reconstruction 139					
		3.4.2.4 Motion Model Estimation					
		3.4.2.5 Attenuation Map Warping					
		3.4.2.6 Motion Corrected Image Reconstruction with Attenu-					
		ation Correction					
		3.4.2.7 Evaluation					
	3.4.3	Results					
	3.4.4	Discussion					
3.5	Compa	arison of Motion Correction Methods Incorporating Motion Mod-					
	elling	for PET/CT Using a Single Breath Hold Attenuation Map 145					
	3.5.1	Introduction					
	352	Methods 146					

			3.5.2.1	XCAT Volume Generation
			3.5.2.2	PET Acquisition Simulation and Non-Attenuation Cor-
				rected Image Reconstruction
			3.5.2.3	Registration
			3.5.2.4	Motion Model Estimation
			3.5.2.5	Attenuation Map Warping
			3.5.2.6	Motion Corrected Image Reconstruction with Attenu-
				ation Correction
			3.5.2.7	Evaluation
		3.5.3	Results	
		3.5.4	Discussion	on
	3.6	Discus	ssion and C	Conclusion
4	A	lesies a N	Istian Can	mestica IItilisias Metica Medelling to Declistic Cincula
4	App tion:		louon Cor	rection Utilising Motion Modelling to Realistic Simula-
	4.1		iow	
	4.1			
	4.2			Correction Exploiting Motion Models Fit on Coarsely
	4.3			ied to Finely Gated Data
		4.3.1		ion
		4.3.2	Methods	
		4.3.2	4.3.2.1	XCAT Volume Generation
			4.3.2.1	PET Acquisition Simulation
			4.3.2.3	MLACF Image Reconstruction
			4.3.2.4	Registration and Motion Model Estimation 165
			4.3.2.5	Motion Corrected Image Reconstruction with Attenu-
			4226	ation Correction
		4.0.0	4.3.2.6	Evaluation
		4.3.3		
		434	Discussion	on 169

	4.4	Evaluation of PET/CT Motion Correction Incorporating Motion Models					
		Using	MLACF a	nd Complex Gating Schemes	. 170		
		4.4.1	Introduct	tion	. 170		
		4.4.2	Methods		. 170		
			4.4.2.1	XCAT Volume Generation and PET Acquisition Simu-			
				lation	. 173		
			4.4.2.2	Image Reconstruction	. 174		
			4.4.2.3	Registration, Motion Model Estimation, and Motion			
				Corrected Image Reconstruction with Attenuation Cor-			
				rection	. 177		
			4.4.2.4	Evaluation	. 178		
		4.4.3	Results		. 178		
		4.4.4	Discussion	on	. 187		
	4.5	Evalua	ation of the	Effect of Lesion Size on the Quality of Motion Correction	on 188		
		4.5.1	Introduct	tion	. 188		
		4.5.2	Methods		. 189		
		4.5.3	Results		. 198		
		4.5.4	Discussion	on	. 203		
	4.6	Discus	ssion and C	Conclusion	. 203		
_	D 4	ъ.	C .		205		
5			Ö	te Signal Extraction for Dynamic PET	207		
	5.1						
	5.2						
	5.3						
		5.3.1	•	quisition and Train/Test Split			
		5.3.2		onal PCA			
		5.3.3	C	Window Method			
		5.3.4		e Interval Method			
		5.3.5	Score, Se	elect, and Combine Method	. 215		
			5.3.5.1		. 215		
			5	.3.5.1.1 Frequency Based	. 215		

			5.	3.5.1.2 Neural Network Based	218
			5.3.5.2		220
	5.4	Evalua	tion		221
		5.4.1	Data Prep	paration	221
		5.4.2	Post-Proc	cessing	224
			5.4.2.1	Parallel Compression	224
			5.4.2.2	Outlier Removal	226
			5.4.2.3	Smoothing	226
		5.4.3	Evaluatio	on Methods	227
	5.5	Results	S		230
	5.6	Discus	sion		235
	5.7	Conclu	ision		237
6	Disc	ussion a	and Conclu	usion	238
	6.1	Overvi	ew		238
	6.2	Main R	Results .		239
		6.2.1	Main Cor	ntributions and Knowledge Advancement	242
	6.3	Outloo	k		242
		6.3.1	Limitation	ns	243
		6.3.2	Future We	ork	244
Ap	pend	ices			247
A	Mot	ion Cor	rection M	ethods Table	247
В	Pseu	do-Bay	esian DIP	Denoising as a Preprocessing Step for Kinetic Mo	odel-
	ling	in Dyna	mic PET		253
	B.1	Abstrac	ct		253
	B.2	Introdu	iction		254
	B.3	Method	ds		255
		B.3.1	Network	Design and Execution	255
		B.3.2	PET Acqu	uisition Simulation and Image Reconstruction	259
		B.3.3	Kinetic M	Modelling	260

		B.3.4	Evaluation	on	260			
	B.4	Result	s		260			
	B.5	Discus	ssion and (	Conclusions	264			
C	PCA	Data 1	Driven Su	arrogate Signal Extraction Methods for Dynamic	PET			
	Full	Overvi	ew		265			
D PCA Data Driven Surrogate Signal Extraction Methods for Dy				urrogate Signal Extraction Methods for Dynamic	PET			
	Resu	ılts Wit	h and Wi	thout Pre- and Post-Processing	267			
E	PCA	Data 1	Driven Su	rrogate Signal Extraction Methods for Dynamic	PET			
	Para	illel Co	mpression	1	271			
F	A Bayesian Neural Network-Based Method for the Extraction of a Metabol-							
	ite C	Correcte	ed Arteria	l Input Function from Dynamic PBR28 PET	275			
	F.1	Abstra	ct		275			
	F.2	F.2 Introduction						
	F.3	3 Methods						
		F.3.1	Data Ac	quisition and Processing	277			
		F.3.2	Neural N	Network Design	278			
			F.3.2.1	Neural Network One Autoencoder	279			
			F.3.2.2	Neural Network Two Signal Extractor	279			
			F.3.2.3	Neural Network Three Metabolite Correction and F	Re-			
				shaping	280			
		F.3.3	Evaluati	on	280			
			F.3.3.1	Candidate Signals	280			
	F.4	Result	s		281			
	F.5	Discus	ssion and (	Conclusions	282			
Bil	bliogr	aphy			284			

# **List of Figures**

2.1	Graphical example of $\beta$ +-decay. Here in the top left of the figure a	
	nucleus can be seen which is unstable as it has an imbalance of protons	
	and neutrons. A positron can be seen exiting the nucleus as a by-product	
	of $\beta$ +-decay converting a proton into a neutron. In the bottom right of	
	the figure a closer example of this can be seen. Here it directly shows a	
	specific proton and the neutron, positron, and neutrino which are produced	
	by $eta+$ -decay	66
2.2	Graphical example of positron range. Here on the right of the figure an	
	atom can be seen travelling with some velocity, a positron is emitted from	
	the nucleus of the atom by $\beta+$ -decay. The path that this positron takes	
	can be seen in the centre of the figure represented by a blue line, this	
	path is the positron range. On the left of the figure the annihilation of the	
	positron with an electron occurs and the $\gamma$ -photons emitted 180.0° apart	
	are shown.	67
2.3	Graphical representation of the difference between a total body PET scan-	
	ner and a standard PET scanner. On the left of the figure a total body PET	
	scanner can be seen, where the rings of detectors completely engulf the	
	patient. However, on the right of the figure a standard PET scanner can be	
	seen where the rings of detectors only cover a portion of the patient. For	
	the case on the right of the figure, in order to take a scan over the entire	
	body, either individual acquisitions will be needed and concatenated, or	
	the bed would have to move while the acquisition was ongoing	68

2.4	Graphical representation of a $\gamma$ -photon scattering off of a particle. The	
	$\gamma$ -photon can be seen entering the figure on the left hand side, before	
	scattering off of a particle in the centre of the figure by an angle $\theta$ and	
	exiting the figure on the top right hand corner. The electron exits the	
	scatter event with some velocity represented by the arrow towards the	
	bottom right.	70
2.5	Graphical representation of the different types of coincidences possible in	
	PET. On the left of the figure a true coincidence can be seen, this is where	
	the $\gamma$ -photons from one annihilation are both detected without scattering.	
	In the middle of the figure a scattered coincidence can be seen, this is	
	where the $\gamma$ -photons from one annihilation are both detected, however in	
	this case one of them has scattered. On the right of the figure a random	
	coincidence can be seen, this is where the $\gamma$ -photons from two unrelated	
	annihilations are detected	71
2.6	Graphical representation of the concept of TOF. The top middle of this	
	figure shows the position where a hypothetical annihilation has occurred	
	(plus the $\gamma$ -photos from this annihilation which have then gone on to be	
	detected by the scanner). The bottom left of this figure shows a traditional	
	Non-TOF acquisition, where the probability of the position of the annihil-	
	ation along the Line of Response (LOR) is constant. The bottom right of	
	this figure shows a TOF acquisition, where the probability of the position	
	of the annihilation along the LOR can be approximated with a Gaussian	
	based upon the difference in arrival time of both photons	74
2.7	Example of some Non-Attenuation Corrected (NAC) Non-TOF data, with	
	noise, with no motion, randoms or scatters, of the thorax with a spherical	
	lesion in the lungs. Coronal view.	74
2.8	Example of some NAC TOF data, with noise, with no motion, randoms	
	or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.	75
2.9	Example of some simulated sinogram data, with no motion, noise, ran-	
	doms or scatters, of the thorax with a spherical lesion in one of the lungs.	
		76

2.10	Example of some simulated viewgram data, with no motion, noise, randoms or scatters, of the thorax with a spherical lesion in one of the lungs.	
		76
2.11	Example of a simulated $\mu$ -Map, with no motion or noise, of the thorax with a spherical lesion in the lungs. Coronal view	80
2.12	Example of simulated NAC Non-TOF reconstructed data, with motion, with no noise randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view.	80
2.13	Example of simulated Attenuation Corrected (AC) Non-TOF reconstructed data, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view	81
2.14	Graphical representation of a Two Dimensional (2D) solution space and an optimiser stepping through this space. In the bottom left of this figure the initial estimate for some optimisation can be seen at $x0$ , the subsequent iterations 1 through 4 can then be seen taking this estimate closer to the centre of the contour plot. Here the contour plot could either be showing a maximisation or a minimisation depending on how it is visualised	87
2.15	Example of a simulated AC Non-TOF Ordered Subset Expectation Maximisation (OSEM) reconstruction, with motion and noise, with no randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view.	91
2.16	Example of a simulated CINE-CT AC Non-TOF OSEM reconstruction, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.	95
2.17	Example of a simulated single static $\mu$ -Map AC Non-TOF OSEM reconstruction, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.	95

2.18	A graphical representation of the concepts of push and pull interpolation.	
	On the left of the figure, under section A, push interpolation can be seen.	
	Here, arrows showing the 'pushing' of the value from the source to the	
	target image can be seen. The blue highlighted pixels in the target image	
	show where holes would be. On the right of the figure, under section ${\bf B}$ ,	
	pull interpolation can be seen. Here, arrows showing the location of the	
	coordinates of the pixels from the target image deformed into the source	
	image space can be seen (Akintonde 2021)	. 99
2.19	Graphical representation of a rigid deformation. On the left of this figure	
	a triangle with vertices $A$ , $B$ and $C$ can be seen. This triangle undergoes	
	a rigid deformation (a translation down and to the right) to a triangle, on	
	the right of this figure, with vertices $A'$ , $B'$ and $C'$	. 101
2.20	Graphical representation of an affine transformation. In the centre of this	
	figure a cyan leaf can be seen which undergoes an affine transformation	
	(a translation down and to the left, a rotation anticlockwise and a scale	
	down) to a red leaf, on the left of this figure. The cyan leaf also undergoes	
	an affine transformation (a translation down and to the right, a rotation	
	clockwise, a scale down and a mirroring) to a blue leaf, on the right of	
	this figure.	. 102
2.21	Graphical representation of a non-rigid deformation. On the left a grid	
	can be seen, this grid undergoes a non-rigid deformation to the grid on	
	the right.	. 104
2.22	Graphical representation of amplitude gating. Here the pseudo sinusoidal	
	signal represents the SS and the horizontal lines, colour coded differently,	
	represent the amplitude gates that the data upon the SS would be binned	
	into	. 106
2.23	Graphical representation of phase gating. Here the pseudo sinusoidal	
	signal represents the SS and the vertical lines, colour coded differently,	
	represent the phase gates that the data upon the SS would be binned into.	106

2.24	Photograph of the Real Time Position Management (RPM). On the bot-
	tom of this figure is the infrared camera and infrared Light Emitting
	Diodes (LEDs) used to locate and track an infrared reflecting marker. On
	the top of this figure is the infrared reflecting marker which is placed onto
	the chest or stomach of the patient in order to track the respiratory amp-
	litude of the patient. Four reflective points are used to track the marker in
	Three Dimensional (3D)
2.25	Example of a point cloud acquired on a Microsoft Kinect camera 109
2.26	Graphical representation of Principal Component Analysis (PCA) applied
	to 2D data (blue circles). PCA finds the 2 eigenvectors indicated in green 111
2.27	Graphical representation of the frequency spectrum for signals correspond
	to three Principal Components (PCs), together with a frequency window
	between 0.1 Hz and 0.4 Hz. In this example, the first signal has the highest
	peak within this window
2.28	Graphical representation of the difference in methodology of pair- and
	group-wise registration. At the top of this figure pair-wise registration
	is shown. Here, the floating image is directly registered to the reference
	image. This process will be repeated independently for the number of
	images which are to be registered. At the bottom of this figure group-wise
	registration is shown. Here, a common reference to all floating images is
	chosen (for instance a weighted sum of the target images). This ensures
	the reference image is located closer on average to all floating images than
	if a single floating image was chosen. Pair-wise registration is complete
	after one iteration of registration. Group-wise registration would involve
	updating the common reference and iteratively registering again for a set
	number of iterations (as seen in figure 2.29)

2.29	Graphical representation of one method of implementing group-wise re-	
	gistration. The sub-images one through four show the initial setup used	
	for this implementation and the sub-images five through 10 show the sub-	
	sequent iterations of the group-wise algorithm. Sub-image one shows	
	an initial pair-wise registration of one image to the rest of the images in	
	the series. Sub-image two shows how to find the mean position of all	
	images by taking the mean of the deformations from the previous step.	
	Sub-image three shows the inverse of the deformation from sub-image	
	two. Sub-image four shows the deformations from sub-image one being	
	composed with the deformation from sub-image three, thus there is now	
	a deformation from each image to the mean position. Sub-image five	
	shows the resultant deformations from the composition in sub-image four.	
	Sub-image six shows the target images resampled using the deformations	
	from sub-image five to form a reference image at the mean position. Sub-	
	images seven through 10 show incremental improvements to the reference	
	image from subsequent iterations of registration and resampling	. 117
2.30	Graphical representation of the process of motion modelling. On the	
	left of the figure the input data of both the image data and the SS can	
	be seen. Then in the centre of the image the image data is taken and	
	registration is applied to it. Then on the right of the image a Respiratory	
	Correspondence Model (RCM) is fit on the Deformation Vector Fields	
	(DVFs) from the registration step and the input SS. In the top right of the	
	figure a naive example of the 2D regression result can be seen. In this	
	case the regression appears to be polynomial	. 119
2.31	Graphical representation of the types of RCM. At the top of the figure is	
	a SS which has been discretised into specific values. On the bottom left	
	of the figure is a linear RCM, fit on a single SS, showing how the points	
	from the discretised SS would be mapped to itself. On the bottom middle	
	of the figure is a RCM fit on SS reflecting the phase of the respiration	
	and on the bottom right of the figure is an RCM where inhalation and	
	exhalation are modelled separately using two RCM (or a 2D SS)	. 122

3.1	Coronal images illustrating results from the estimation of RCMs on NAC	
	images. All volumes correspond to the end inhalation gate. First row	
	from left to right: Four Dimensional Extended Cardiac Torso (XCAT)	
	PET ground truth image, NAC Non-TOF reconstructed data, and NAC	
	TOF reconstructed data. Second row, RCMs applied to mean position	
	XCAT data with RCMs derived from XCAT PET data (left), NAC Non-	
	TOF (middle), and NAC TOF (right) volumes. Colour map ranges are	
	consistent for all images on this row. The third row from left to right,	
	difference between the estimated volumes from the second row with the	
	XCAT end inhalation volume. Colour map ranges are consistent for all	
	images on this row.	. 134
3.2	The path of the Centre of Mass (COM) of the lesion, in voxel indices. Ho-	
	rizontal (respectively vertical) axis corresponds to motion in the Anterior	
	Posterior (AP) (respectively Superior Inferior (SI)) direction over the six	
	gates. Different curves denote COM displacement for ground truth data,	
	the estimated data from the XCAT based RCM, the estimated data from	
	the NAC Non-TOF based RCM, and the estimated data from the NAC	
	TOF based RCM	. 135
3.3	The path of the COM of the lesion, in voxel indices. Horizontal (respect-	
	ively vertical) axis corresponds to motion in the SI (respectively AP).	
	Different curves denote COM displacement for ground truth data, the	
	estimated data from the NAC based MM, and the estimated data from the	
	AC based MM	141
3.4	Ground truth and reconstructions using ungated (CINE-CT), ungated	
	(static CT), NAC MM, and AC MM. Colour map ranges are consistent	
	for all images.	. 142
2.5	A much source the locion for uncertal (CINIE CIT)	
3.5	A profile across the lesion for ungated (CINE-CT), ungated (static CT), NAC MM, and AC MM	143
	NAC IVIVE AND ACTIVIVE	144

3.6	Coronal slices for various methods. First column contains AC motion	
	corrected reconstructions and the second column contains the result of	
	applying the final motion correction on the original XCAT images (for	
	easier assessment of the accuracy of the estimated DVFs). This is for	
	ungated static CT, ungated averaged CINE-CT, pair-wise, pair-wise MM,	
	group-wise, and group-wise MM. Colour map ranges are consistent for	
	all images in each column.	. 151
3.7	A profile across the lesion for ungated static CT, ungated averaged CINE-	
	CT, pair-wise, pair-wise MM, group-wise, and group-wise MM	. 152
4.1	This figure shows an example of the amplitude and gradient 2D SS gat-	
	ing technique as compared to examples of the pseudo-phase based gating	
	technique. On the right of this image is a more traditional gating tech-	
	nique, here using 10 amplitude gates and three gradient gates. The next	
	four images show the pseudo-phase based gating with two, three, four,	
	and eight radial bins, plus an example of how a circular separation can be	
	used for deep and shallow breaths	. 164
4.2	First row contains, AC motion corrected reconstructions (plus Structural	
	Similarity Index Measure (SSIM) to the ground truth). Second row con-	
	tains the results of applying the final MM on the original XCAT volumes.	
	This is using a MM fit on four and 30 gate binned data applied to 30 gate	
	binned data, a MM fit on noiseless four and 30 gate binned data applied	
	to 30 gate binned data, and ungated data AC with a static $\mu$ -Map at end	
	inhalation and all $\mu$ -Maps summed. Colour map ranges are consistent for	
	all images in each row.	. 167
4.3	A profile through the lesion in the SI direction, summed over a window in	
	the AP and Lateral Medial (LM) directions, with median smoothing, for	
	the ground truth XCAT data, a MM fit on four and 30 gate binned data	
	applied to 30 gate binned data, a MM fit on noiseless four and 30 gate	
	binned data applied to 30 gate binned data, and ungated data AC with a	
	static $\mu$ -Map at end inhalation and all $\mu$ -Maps summed	. 168

4.4

A visual analysis which shows the results of performing different reconstruction methods on XCAT data with a 12 mm lesion placed into the base of the right lung (the left of this image). Where noise is present then this data represents a 60 s acquisition on a GE Discovery 710 (roughly equivalent to one gate from a clinical acquisition). At the top of this image is the Ground Truth image (without motion or reconstruction artifacts), second from the top of this image is a NAC OSEM reconstruction, third from the top of this image is a AC OSEM reconstruction using a  $\mu$ -Map at a different respiratory position than the Ground Truth, fourth from the top of this image is a MLACF reconstruction following the literature (Nuyts et al. 2012). From the fifth image from the top to the bottom of this image all images represent adding additional processing steps to attempt to make the reconstructions better for motion correction. Fifth from the top of this image is the result of initialising MLACF with a misaligned AC OSEM reconstruction, using sinogram mashing, and fixing the scale of the image to the scale of the misaligned AC OSEM reconstruction. Sixth from the top of this image is the result of applying a median smoothing (with a kernel size of three) after each activity update (and before every attenuation update). Seventh from the top of this image is the result of interpolating any value outside three times the Inter Quartile Range (IQR) of the image after each activity update using a linear interpolation. The bottom of this image is the result of heavily smoothing the endplanes. All reconstructed images include scatter and random events. All OSEM reconstructions follow standard clinical practise (see chapter 3). For the NAC OSEM and MLACF following the literature (Nuyts et al. 2012) the output was rescaled to the AC OSEM. Colour map ranges are consistent for all images. 

4.5	AC motion corrected reconstructions (plus Perception Based Image Qual-	
	ity Evaluator (PIQE)), where different data was used to fit the MM and the	
	final MM was applied to 30 gate binned data. The data used to fit the MM	
	was as follows. 'Phase' gated and MLACF reconstructed data with three,	
	four, eight and 12 bins respectively. A MM was also fit using a single bin,	
	which is practically the equivalent of doing no motion correction. Colour	
	map ranges are consistent for all images in each row	. 178
4.6	AC motion corrected reconstructions (plus PIQE), where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. Amplitude gated data	
	with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM.	
	A MM was also fit using a single bin, which is practically the equivalent	
	of doing no motion correction. Colour map ranges are consistent for all	
	images in each row.	. 179
4.7	AC motion corrected reconstructions (plus PIQE), where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. Amplitude gated data	
	with five bins reconstructed using MLACF. Here, in the 2D case a 2D SS	
	and in the One Dimensional (1D) case a 1D SS were used. A MM was	
	also fit using a single bin, which is practically the equivalent of doing no	
	motion correction. Colour map ranges are consistent for all images in	
	each row.	. 179
4.8	AC motion corrected reconstructions (plus PIQE), where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. Amplitude gated data	
	with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM.	
	A MM was also fit using a single bin, which is practically the equivalent	
	of doing no motion correction. Colour map ranges are consistent for all	
	images in each row	170

- 4.9 AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The methods presented here are a combination of the best one method from figure 4.5, figure 4.6, and figure 4.8. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row. 179
- 4.11 AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row. . . 180

4.12	AC motion corrected reconstructions (plus PIQE), where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. Amplitude gated data	
	with five bins reconstructed using MLACF. Here, in the 2D case a 2D SS	
	and in the 1D case a 1D SS were used. A MM was also fit using a single	
	bin, which is practically the equivalent of doing no motion correction.	
	Here, results are shown applied to volumes without noise. This is due	
	to it potentially being easier to see smaller differences between motion	
	correction methods when noise is removed. However, noise is still used	
	during the motion correction process. Colour map ranges are consistent	
	for all images in each row.	. 180
4.13	AC motion corrected reconstructions (plus PIQE), where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. Amplitude gated data	
	with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM.	
	A MM was also fit using a single bin, which is practically the equivalent of	
	doing no motion correction. Here, results are shown applied to volumes	
	without noise. This is due to it potentially being easier to see smaller	
	differences between motion correction methods when noise is removed.	
	However, noise is still used during the motion correction process. Colour	
	map ranges are consistent for all images in each row	. 181
111	AC motion corrected reconstructions (plus PIQE), where different data	
4.14		
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The methods presented here are a combination of the best one	
	method from figure 4.10, figure 4.11, and figure 4.13. Here, results are	
	shown applied to volumes without noise. This is due to it potentially be-	
	ing easier to see smaller differences between motion correction methods	
	when noise is removed. However, noise is still used during the motion	
	correction process. Colour map ranges are consistent for all images in	
	each row	. 181

4.15	A profile through the lesion, in the SI direction, summed over a window	
	in the AP and LM directions, with median smoothing. This is for where	
	different data was used to fit the MM and the final MM was applied to	
	30 gate binned data. The data used to fit the MM was as follows. 'Phase'	
	gated and MLACF reconstructed data with three, four, eight and 12 bins	
	respectively. A MM was also fit using a single bin, this is practically the	
	equivalent of doing no motion correction	. 182
4.16	A profile through the lesion, in the SI direction, summed over a window	
	in the AP and LM directions, with median smoothing. This is for where	
	different data was used to fit the MM and the final MM was applied to $30$	
	gate binned data. The data used to fit the MM was as follows. Amplitude	
	gated data with five bins reconstructed using MLACF, AC OSEM, and	
	NAC OSEM. A MM was also fit using a single bin, which is practically	
	the equivalent of doing no motion correction	. 182
4.17	A profile through the lesion, in the SI direction, summed over a window	
	in the AP and LM directions, with median smoothing. This is for where	
	different data was used to fit the MM and the final MM was applied to $30$	
	gate binned data. The data used to fit the MM was as follows. Amplitude	
	gated data with five bins reconstructed using MLACF. Here, in the 2D	
	case a 2D SS was used and in the 1D case a 1D SS was used. A MM was	
	also fit using a single bin, this is practically the equivalent of doing no	
	motion correction.	. 183
4.18	A profile through the lesion, in the SI direction, summed over a window	
	in the AP and LM directions, with median smoothing. This is where	
	different data was used to fit the MM and the final MM was applied to $30$	
	gate binned data. The data used to fit the MM was as follows. Amplitude	
	gated data with 30 bins reconstructed using MLACF, AC OSEM, and	
	NAC OSEM. A MM was also fit using a single bin, which is practically	
	the equivalent of doing no motion correction	. 183

 4.21 A visual analysis which shows the results of performing the method on XCAT data with a 8 mm lesion placed into the base of the right lung (the left of this image). At the top of this image is the Ground Truth image (without motion or reconstruction artifacts), second from the top of this image is where no motion correction has been applied to noiseless data, third from the top of this image is the results of applying a motion correction fit on reconstructed data to noiseless data, fourth from the top of this image is the result of reconstructing an image without motion which has been reconstructed, fifth from the top of this image is the result of reconstructing an image with motion which has been reconstructed, and at the bottom of this image is the result of applying the motion correction to data which has been reconstructed. In this case the input to the motion correction was the 4 Phase MLACF reconstructions and the motion correction has been applied to the 30 bin OSEM reconstructed images.

4.23	A profile through the 8 mm lesion, in the SI direction, summed over a
	window in the AP and LM directions, with median smoothing. The blue
	line shows the profile from the Ground Truth image with a 8 mm lesion,
	the dashed orange line shows the result of not applying motion correction,
	and the dashed green line shows the result of applying a motion correction
	on reconstructed data with a 8 mm lesion. In this case the input to the
	motion correction was the 4 Phase MLACF reconstruction and the motion
	correction has been applied to the 30 bin noiseless images
5.1	This diagram shows the potential workflows for the different methods ex-
	amined here. The path on the left shows only the signal (or signals) being
	extracted without any further processing (until the next steps). The path
	in the middle shows the process for extracting a signal using windows.
	The path on the right shows how one would extract a signal using a PC
	from a point in the data where it could be considered to be static 211
5.2	A plot showing the moving window size optimisation for the PCA method.
	For different fixed window sizes, the correlation of the extracted signal
	to the RPM is shown for the windows sliding over the whole acquisition
	(taken for the first acquisition of patient one). Note that 0.5 s time frames
	were used
5.3	A plot showing the moving window size optimisation for the Spectral
	Analysis Method (SAM) method. For different fixed window sizes, the
	correlation of the extracted signal to the RPM is shown for the windows
	sliding over the whole acquisition (taken for the first acquisition of patient
	one). Note that $0.5\mathrm{s}$ time frames were used
5.4	A diagram showing how PCs can be selected and combined
5.5	An example of the output for the Neural Network (NN) can be seen here
	testing signals extracted from the first acquisition of patient one using
	PCA
5.6	A diagram showing pre-processing performed
5.7	A diagram showing the parallel compression method

5.8	A diagram showing the post-processing performed	. 226
5.9	Output of each method compared to the RPM for the first usable 120 s	
	(between 20 s and 140 s) (taken for the first acquisition of patient one).	
	This is for Conventional PCA, Moving Window PCA, Late Time Interval	
	PC, Score, Select, and Combine using frequency and NN scoring, and the	
	Moving Window SAM method	. 228
5.10	Output of each method compared to the RPM for the first usable 120 s	
	(between 20 s and 140 s) (taken for the first acquisition of patient eight).	
	This is for Conventional PCA, Moving Window PCA, Late Time Interval	
	PC, Score, Select, and Combine using frequency and NN scoring, and the	
	Moving Window SAM method	. 228
5.11	A box plot showing for each method its correlation coefficient to the RPM	
	for the first usable 120 s (between 20 s and 140 s) (taken for the seven test	
	acquisitions). This is for Conventional PCA, Moving Window PCA, Late	
	Time Interval PC, Score, Select, and Combine using frequency and NN	
	scoring, and the Moving Window SAM method	. 229
5.12	A box plot showing for each method its correlation coefficient to the RPM	
	for the entire acquisition (taken for the seven test acquisitions). This is for	
	Conventional PCA, Moving Window PCA, Late Time Interval PC, Score,	
	Select, and Combine using frequency and NN scoring, and the Moving	
	Window SAM method	. 229
5.13	Signals for the first usable 120 s (between 20 s and 140 s) generated using	
	the Select, Score, and Combine NN based scoring method. This is for	
	the five acquisitions which did not have a usable RPM signal and as such	
	could not be used for training or testing	. 230
5.14	RPM signals for the first usable 120 s (between 20 s and 140 s) (for the	
	seven test acquisitions). Notice that only the first acquisition of patient	
	one and patient six shows a steady trace with an average frequency, every	
	other trace shows variable breathing or artefacts	. 231

5.15	The result of applying Fast Fourier Transform (FFT) to the RPM signals	
	for the first usable $120s$ (between $0.05Hz$ and $0.45Hz)$ (for the seven	
	test acquisitions). Notice that four of the seven have peaks very close to	
	or outside the lower boundary of the resiratory window. Also notice that	
	two of the seven have very wide frequency responses, which would be	
	difficult for the automatic selection of a respiratory frequency window.	. 232
5.16	Correlation coefficients to the RPM for the first 120 s in 20 s intervals	
	(between $20\mathrm{s}$ and $140\mathrm{s}$ ) (taken as a mean for all data sets). This is for	
	Conventional PCA, Moving Window PCA, Late Time Interval PC, Score,	
	Select, and Combine using frequency and NN scoring, and the Moving	
	Window SAM method. The stair plots are staggered for the different	
	methods for visual clarity.	. 232
5.17	A plot showing a single 'view' of the original PET data (top) as well	
	as the PCs used to generate the output signal for the Conventional, Late	
	Time Interval and the two Score, Select, and Combine methods (taken for	
	the first acquisition of patient one)	. 233
B.1	Graphical representation of dropout. On the left of this figure a NN	
	which is experiencing a large degree of dropout can be seen. Here, nodes	
	which have been dropped are represented with a cross through it and their	
	outgoing/incoming connections are remove. On the right of this figure	
	the same network can be seen without dropout and will all nodes active	
	and fully connected.	. 255
B.2	Graphical representation of a number of activation functions. In the top	
	left of this figure can be seen the sigmoid function. In the top right of	
	this figure the tanh activation function can be seen. In the middle left of	
	this figure the Rectified Linear Unit (ReLU) activation can be seen. In the	
	bottom right of this figure the Exponential Linear Unit (ELU) activation	
	can be seen. In the bottom right of this figure the leaky ReLU can be seen	. 256

B.3	Graphical representation of the Scaled Exponential Linear Unit (SELU)
	activation function. This figure shows a scaled version of the ELU activa-
	tion function showed in figure B.2. The same advantages and disadvant-
	ages as with ELU are apparent here. However, as an additional advantage
	if the input data to this activation are normalised then the output of the
	activation will also be normalised
B.4	Graphical representation of the process of convolution. This figure shows
	a $3 \times 3$ convolution kernel operating on a padded image at the bottom
	of the figure to create the new image at the top of the figure. For a Con-
	volutional Neural Network (CNN) each value in the convolution kernel
	is determined using weights and biases, like in a fully connected NN.
	However, rather than there being a weight and bias for every element in
	the input data there is only one for the kernel. Therefore, the number of
	parameters is significantly reduced and if the NN is fully convolutional
	then the input size does not have to be fixed
B.5	First column contains, a visual analysis between the ground truth and
	denoised results (taken for the last time point, plus SSIM to the ground
	truth), and the second column contains, the $K_i$ results (all voxels in a
	coronal view) of a Patlak reconstruction of all time points (plus SSIM to
	the ground truth), for the ground truth, and data denoised using, Total Vari-
	ation (TV), the implementation of Deep Image Prior (DIP) from (Gong
	et al. 2019), and our new implementation of DIP, trained sequentially and
	combined (taken for the lung Field of View (FOV)). Last row contains,
	uncertainty volumes, for the data denoised using our new implementation
	of DIP, trained sequentially and combined (taken for the last time point
	of the lung FOV). Colour map ranges are consistent for all images in each
	section
B.6	$K_i$ results (single voxel) of a Patlak reconstruction of all time points, plus
	uncertainty where applicable, for the ground truth, and data denoised
	using, TV, the implementation of DIP from (Gong et al. 2019), and our
	new implementation of DIP, trained sequentially and combined 262

B.7	A Time Activity Curve (TAC) through a lesion, fit as a third order poly-	
	nomial regression, with weighting using uncertainty (where available),	
	for the ground truth, the original noisy data, and this data denoised using,	
	TV, the implementation of DIP from (Gong et al. 2019), and our new	
	implementation of DIP, trained sequentially and combined, both with and	
	without uncertainty (taken for the lung FOV)	. 263
B.8	A profile through a lesion, in the SI direction, for the ground truth, the	
	original noisy data, and this data denoised using, TV, the implementation	
	of DIP from (Gong et al. 2019), and our new implementation of DIP,	
	trained sequentially and combined (taken for the last time point of the	
	lung FOV).	. 263
C.1	A diagram showing an overview of the possible ways in which the method	
	could be executed.	. 266
D.1	A visual comparison between the RPM SS and the SS from the static	
	PCA method using only the first PC for three patients between 20 s and	
	160 s. The three patients shown here are the ones on which parameters	
	for the method were optimised	. 267
D.2	A visual comparison between the RPM SS and the SS from the static PCA	
	method by combining the first 20 PCs for three patients between 20 s and	
	160 s. The three patients shown here are the ones on which parameters	
	for the method were optimised. Here no pre- or post-processing is used.	. 268
D.3	A box plot of the correlation coefficient between static PCA method both	
	when using only the first PC and by combining the first 20 PCs for all	
	patients between 20 s and 160 s and for the entire acquisition. The para-	
	meters used here were optimised for and frozen based on the results shown	
	in other diagrams on three patients. Here no pre- or post-processing is	
	used	268

D.4	A visual comparison between the RPM SS and the SS from the static PCA method by combining the first 20 PCs for three patients between 20 s and 160 s. The three patients shown here are the ones on which parameters for the method were optimised. Here pre- and post-processing are used 2	269
D.5	A box plot of the correlation coefficient between static PCA method both when using only the first PC and by combining the first 20 PCs for all patients between 20 s and 160 s and for the entire acquisition. The parameters used here were optimised for and frozen based on the results shown in other diagrams on three patients. Here pre- and post-processing are used	269
E.1	This figure shows an example of the processed count rate and processed gradient of the processed count rate, between 0 s and 140 s, which are used	
	as part of the parallel compression post-processing step. These have been taken for the first acquisition of patient one. Here the processed count rate is normalised between zero and one by subtracting by the minimum and dividing by the maximum value of the signal	271
F.1	Predicted $V_T \in \mathbb{R}^{r \times s}$ in the test-set subjects, with $r = 69$ and $s = 5$ estimated with the four candidate signals, correlated to the True Volume of Distribution (V <sub>T</sub> ) (obtained with TRUE-Arterial Input Function (AIF)). Please note that for the NN-based methods, the displayed V <sub>T</sub> s were averaged over all realisations.	281

# **List of Tables**

3.1	This table shows a short summary of the highlights of the methods used	
	in section 3.3. This table is useful for quickly referencing the changes	
	between sections. To see the development of the method throughout the	
	thesis please see table A.1 in appendix A	131
3.2	Comparison of the Mean Absolute Percentage Error (MAPE) between	
	the ground truth data and the volumes estimated from the XCAT based	
	RCMs, the volumes estimated from the NAC Non-TOF based RCM, and	
	the volumes estimated from the NAC TOF based RCM	134
3.3	This table shows a short summary of the highlights of the methods used	
	in section 3.3 and section 3.4. This table is useful for quickly referencing	
	the changes between sections. To see the development of the method	
	throughout the thesis please see table A.1 in appendix A	137
3.4	Comparison of Standard Uptake Value (SUV) <sub>max</sub> , SUV <sub>median</sub> and	
	SUV <sub>peak</sub> between ungated (CINE-CT), ungated (static CT), NAC MM,	
	and AC MM	140
3.5	This table shows a short summary of the highlights of the methods used	
	in section 3.4 and section 3.5. This table is useful for quickly referencing	
	the changes between sections. To see the development of the method	
	throughout the thesis please see table A.1 in appendix A	146
3.6	Comparison of $SUV_{max}$ and $SUV_{peak}$ between ungated static CT, ungated	
	averaged CINE-CT, pair-wise, pair-wise MM, group-wise, and group-	
	wise MM	152

4.1	This table shows a short summary of the highlights of the methods used	
	in section 3.5 and section 4.3. This table is useful for quickly referencing	
	the changes between sections. To see the development of the method	
	throughout the thesis please see table A.1 in appendix A	161
4.2	Comparison of $SUV_{max}$ and $SUV_{peak}$ , for the ground truth XCAT data, a	
	MM fit on four and 30 gate binned data applied to 30 gate binned data,	
	a MM fit on noiseless four and 30 gate binned data applied to 30 gate	
	binned data, and ungated data AC with a static $\mu$ -Map at end inhalation	
	and all $\mu$ -Maps summed	168
4.3	This table shows a short summary of the highlights of the methods used	
	in section 4.3 and section 4.4. This table is useful for quickly referencing	
	the changes between sections. To see the development of the method	
	throughout the thesis please see table A.1 in appendix A	171
4.4	Comparison of $SUV_{max}$ , $SUV_{median}$ , and $SUV_{peak}$ , where different data	
	was used to fit the MM and the final MM was applied to 30 gate binned	
	data. The data used to fit the MM was as follows. 'Phase' gated and	
	MLACF reconstructed data with three, four, eight and 12 bins respectively.	
	Amplitude gated data with five bins reconstructed using MLACF, AC	
	OSEM, and NAC OSEM (Here, in the 2D case a 2D SS was used and	
	in the 1D case a 1D SS was used.). Amplitude gated data with 30 bins	
	reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was	
	also fit using a single bin, which is practically the equivalent of doing no	
	motion correction.	184
4.5	This table shows a short summary of the highlights of the methods used	
	in section 4.4 and section 4.5. This table is useful for quickly referencing	
	the changes between sections. To see the development of the method	
	throughout the thesis please see table A.1 in appendix A	190
4.6	Comparison of $SUV_{max}$ and $SUV_{peak}$ over the lesion, where different	
	lesion sizes were used for the XCAT simulation. The different lesion	
	sizes used were $12  \text{mm}$ , $10  \text{mm}$ , $8  \text{mm}$ , and $6  \text{mm}$ . In the case of the $8  \text{mm}$	
	lesion the results without motion correction are also reported	200

5.1	This table shows the correlation coefficient between the result of the	
	Score, Select, and Combine using NN scoring and the RPM for the train	
	dataset where a number of pre- and post-processing methods have been	
	applied to the data. Each new line represents the addition of this method,	
	therefore the last element includes all previous pre- and post-processing.	
	This is for the Freeman-Tukey transform, the Yeo-Johnson transform,	
	the incorporation of a mask to remove low count areas of the sinogram,	
	smoothing and downsampling the sinogram and smoothing the signal,	
	and parallel compression.	221
5.2	This table shows the p-values acquired as part of a statistical analysis of	
	the test dataset using a mixed-effects model. Before applying the mixed-	
	effects model the signals were first transformed to frequency space using	
	FFT. One values between 0.05 Hz and 0.45 Hz were used. This is for	
	Conventional PCA, Moving Window PCA, Late Time Interval PC, Score,	
	Select, and Combine using frequency and NN scoring, and the Moving	
	Window SAM method	227
A.1	This table shows a short summary of the highlights of all methods used	
	in chapter 3 and chapter 4. This table is useful for quickly referencing the	
	changes between sections and the development of the method throughout	
	the thesis	247
B.1	Comparison of $SUV_{max}$ and $SUV_{peak}$ , for the ground truth, the original	
	noisy data, and this data denoised using, TV, the implementation of DIP	
	from (Gong et al. 2019), and our new implementation of DIP, trained	
	sequentially and combined (taken for the last time point of the lung FOV).	262

# **List of Algorithms**

1	Re-interpolate Outliers
2	Smooth Endplanes
3	MLACF for Motion Correction
4	Motion Modelling
5	Conventional Score
6	Moving Window Method
7	Late Time Interval Method
8	Score and Select PCs
9	Frequency Score
10	Combining PCs
11	Parallel Compression
12	Extract Parallel Compression Weighting 27

# Acronyms

[11C]-PBR28 [11C]-Peripheral Benzodiazepine Receptor 62, 275

[<sup>13</sup>N]-NH3 Nitrogen-13 Ammonia 236

[18F]-FDG Fluorine-18 Fludeoxyglucose 65, 259

1**D** One Dimensional 33, 120, 247

2**D** Two Dimensional 26, 57, 247

3**D** Three Dimensional 28, 78, 251

4D Four Dimensional 58, 254

AC Attenuation Corrected 26, 64, 248

ACF Attenuation Correction Factor 82, 250

AI Artificial Intelligence 276

AIF Arterial Input Function 44, 62, 275

AP Anterior Posterior 30, 132

**ATP** Adenosine Triphosphate 65

BFGS Broyden Fletcher Goldfarb Shanno 88

**CCP** Collaborative Computational Project 9

**CDT** Centre for Doctoral Training 9

CNN Convolutional Neural Network 42, 258

**COD** Centroid of Distribution 113

**COM** Centre of Mass 30, 113

**CPG** Control Point Grid 104, 247

**CT** Computed Tomography 5, 56

**DD** Data Driven 6, 58

**DDG** Data Driven Gating 60

**DD-PCA** Data Driven Principal Component Analysis Surrogate Signal Extraction 115

**DIP** Deep Image Prior 42, 61, 253

**DVF** Deformation Vector Field 29, 58, 249

**DWB** Dynamic Whole Body 259

EANM European Association of Nuclear Medicine 140, 260

**ELU** Exponential Linear Unit 41, 256

**EPSRC** Engineering and Physical Sciences Research Council 9

FDG Fluorodeoxyglucose 64, 282

**FFT** Fast Fourier Transform 41, 113

**FOV** Field of View 42, 64, 254

**FWHM** Full Width at Half Maximum 75, 247, 260, 279

**GE** General Electric 9, 73

**GPU** Graphics Processing Unit 255, 278

**HAB** High Affinity Binder 277

HPLC High Performance Liquid Chromatography 277

**HU** Hounsfield Unit 79

i4health Intelligent, Integrated Imaging in Healthcare 9

**IDIF** Image Derived Input Function 275

**IPF** Idiopathic Pulmonary Fibrosis 209

IQR Inter Quartile Range 32, 175

**KBq/mL** Kilo Becquerel per Millilitre 90

KCPS Kilo Counts Per Second 131, 247

**KLD** Kullback Leibler Deviation 223

**KRG** Kinetic Respiratory Gating 208

**L-BFGS** Low memory Broyden Fletcher Goldfarb Shanno 88

L-BFGS-B Low memory Broyden Fletcher Goldfarb Shanno Bounded 89

**LED** Light Emitting Diode 28, 108

LM Lateral Medial 31, 168

**LOR** Line of Response 25, 56

LSTM Long Short Term Memory 45, 282

**MAB** Mixed Affinity Binder 277

MAD Median Absolute Difference 85

MAE Mean Absolute Error 85

 $\mu$ -Map Attenuation Map 5, 57, 249

MAPE Mean Absolute Percentage Error 46, 133

mCi Millicurie 277

MCIR Motion Compensated Image Reconstruction 125

MI Mutual Information 86

MLAA Maximum Likelihood Reconstruction of Activity and Attenuation 82

**MLACF** Maximum Likelihood Activity and Attenuation Correction Factors Estimation 6, 61, 250

MLEM Maximum Likelihood Expectation Maximisation 90, 250

**MM** Motion Model 6, 58, 247

MR Magnetic Resonance 75, 249, 277

**MSc** Master of Science 10

MSE Mean Squared Error 85, 257

NAC Non-Attenuation Corrected 25, 61, 247

**NaN** Not a Number 172, 251, 272

**nCV** Normalised Coefficient of Variance 276

NIHR National Institute for Health Research 9

NMI Normalised Mutual Information 86, 248

**NN** Neural Network 39, 61, 253, 275

NN-AEIF Neural Network Auto-Encoder Input Function 275

**NN-AIF** Neural Network Arterial Input Function 275

**NN-IDIF** Neural Network Image Derived Input Function 275

**Non-TOF** Non-Time of Flight 6, 74, 247, 260

**OSEM** Ordered Subset Expectation Maximisation 26, 91, 247, 260

PC Principal Component 28, 112, 267

PCA Principal Component Analysis 28, 98, 248, 267

**PET** Positron Emission Tomography 5, 56, 253, 275

**PhD** Doctor of Philosophy 9

PIQE Perception Based Image Quality Evaluator 33, 173

**PSD** Power Spectral Density 215

**PSMA** Prostate Specific Membrane Antigen 65

RANSAC Random Sample Consensus 244

**RCM** Respiratory Correspondence Model 29, 58

**RDP** Relative Difference Prior 86

**ReLU** Rectified Linear Unit 41, 256

**RMSE** Root Mean Square Error 85, 278

**ROI** Region of Interest 65, 280

**RPM** Real Time Position Management 28, 98, 248, 267

**SAM** Spectral Analysis Method 39, 113

**SELU** Scaled Exponential Linear Unit 42, 257

SGD Stochastic Gradient Descent 88

SI Superior Inferior 30, 132, 263

SIRF Synergistic Image Reconstruction Framework 132, 259

**SNR** Signal to Noise Ratio 73

**SPECT** Single-photon Emission Computed Tomography 129

**SRF** Sinogram Region Fluctuation 113

**SS** Surrogate Signal 6, 57, 247, 267

SSD Sum of Squared Differences 100, 247

SSIM Structural Similarity Index Measure 31, 163, 254

**STD** Standard Deviation 113

**STFT** Short Time Fourier Transform 208

STIR Software for Tomographic Image Reconstruction 132, 259

SuPReMo Surrogate Parametrised Respiratory Motion Modelling 133

SUV Standard Uptake Value 46, 90, 260

**SVD** Singular Value Decomposition 113

SyneRBI Synergistic Biomedical Imaging 9

**TAC** Time Activity Curve 43, 113, 259

**TOF** Time of Flight 6, 61, 247

TV Total Variation 42, 254

**UCL** University College London 9

**UCLH** University College London Hospitals 9

UK United Kingdom 9, 56

 $V_T$  Volume of Distribution 44, 275

**VOI** Volume of Interest 260

XCAT Four Dimensional Extended Cardiac Torso 30, 61, 247, 253

# Chapter 1

# Introduction

### 1.1 Introduction

An overview of the topics to be covered is given in this chapter, in order to motivate the work performed. This includes introducing the data used, the potential problems anticipated to be encountered, reasons why this problem was chosen to be tackled, and also a small insight into previous work performed in a similar vein (with more detail given in subsequent chapters). This chapter will then move onto covering the intended objectives of this work before outlining the further structure of this report.

### 1.2 Motivation

Positron Emission Tomography (PET)/Computed Tomography (CT) is a hybrid medical imaging modality. It combines PET (a functional imaging technique used to capture data related to internal metabolic processes of a subject) with an anatomical imaging modality called CT. Every year in the United Kingdom (UK) just over 150,000 PET/CT scans are performed, the majority of these scans are used in the diagnosis and treatment of cancer (NHS England 2020).

PET operates by quantifying the distribution of a radioactive tracer that is administered to the patient. This is based on the detection of pairs of  $\gamma$ -photons using rings of detectors. These opposing  $\gamma$ -photons are created by the annihilation of positrons, produced by the (decay of the) radioactive tracer, with electrons. The measurement of these opposing  $\gamma$ -photons gives what is called a Line of Response (LOR) upon which the annihilation may have taken place. Image reconstruction attempts to find from this

measurement space the image space which would reflect the distribution of the radioactive tracer in the patient. This is an inverse problem, meaning that from a number of measurements the underlying cause or system is found or approximated.

The CT and PET acquisitions of a combined PET/CT scan take place temporally apart from one another by a few minutes. Any movement during or between the acquisition of CT and PET data will lead to the occurrence of artefacts and blurring, or the reduction of resolution of the final image. A single bed position PET scan can take upwards of a few minutes to complete (usually between 90 s and 150 s). Clinical PET acquisitions are usually made up of multiple bed positions. A single fast CT scan (lasting approximately between 5 s and 10 s) is usually used to correct for the attenuation of the signal of the radioactive tracer by the matter of the patient. Inter-acquisition artefacts are caused by a misalignment of the CT and PET volumes, leading to attenuation being corrected for where it should not and vice versa. The misalignment of the Attenuation Map  $(\mu$ -Map) can also introduce bias in the scatter correction. For intra-acquisition blurring, these artefacts are formed in a similar way to how blurring may appear during a long exposure photograph. This is where counts from one specific location are spread about amongst multiple voxels (or pixels in the Two Dimensional (2D) case) in the final image or volume. These errors lead, for instance, to difficulty in the detection and location of lesions. Some sources of movement include head motion, cardiac motion, and respiratory motion from patient breathing.

In clinical practice, it is currently still often the case to forego motion correction. This is because of the, usually, high computational expense and a lack of confidence in the motion correction algorithms (and any evaluation techniques used to prove the effectiveness of motion correction algorithms).

Previous motion correction solutions have mainly focused on binning PET data into separate volumes (where the intra-gate motion is low). Then co-registering these gated volumes and summing the result together, this will be discussed further in chapter 2. PET data is usually binned into a histogram based on a Surrogate Signal (SS), which reflects the position that the patient was in during that part of the acquisition. However, if a single CT is used for attenuation correction, the misalignment problem may still exist. It is unlikely and not guaranteed for the position of this CT to match any one PET bin never

mind all of them.

One way to resolve the issue of having no spatially matching CT and PET volumes is to use CINE-CT or Four Dimensional (4D)-CT. These types of CT acquisition are taken continuously with free breathing of the patient, and thus are more likely to have matching data for each position in the respiratory cycle. This type of data can be replayed sequentially like the frames of a video, hence the CINE-CT or cinema and 4D names. However, this requires a higher dose to the patient. In addition, this CT data would not be simultaneously acquired with the PET data so registration and misalignment issues will still be apparent albeit reduced.

More recent research on motion correction has begun to look into methods known as motion modelling. Here, rather than directly co-registering each individual volume, a Respiratory Correspondence Model (RCM) is formed by fitting a model that relates a SS and the data. This means that the RCM, once fit, can be used like a function where values of a SS (potentially other than the ones used to fit the RCM) can be input and the output would be a Deformation Vector Field (DVF). These DVFs could then be used to motion correct the data from which the SS was obtained. An advantage of motion modelling is that it is less susceptible to noise when compared to co-registering. This is because the RCM is fit over all of the data simultaneously and as such any outliers in the data have less of an effect on the overall result. A further advantage would be what was mentioned earlier, which is that a RCM can be fit on less coarsely gated data and then applied to more coarsely gated data. This means that potentially more robust data could be used to fit the model, leading to a more robust model. Motion Models (MMs) offer the option to tailor the coarseness of the data used to fit the model and apply the model independently, which direct co-registering does not.

SSs can be derived from a multitude of sources, including from an external device (which measures the position of the patient) or through a Data Driven (DD) method (where the value of the SS is derived solely from the data of the acquisition). External equipment for estimation of the SS are not desirable. This due to the large impact on clinical workflow, from having to affix the external device to the patient. Furthermore, they cannot be applied retrospectively. PET data can already be gated using DD techniques (including retrospectively) without need for external equipment. However, DD methods

struggle on dynamic acquisitions due to variation caused by tracer kinetics.

This work will specifically focus on correcting for respiratory motion. Some reasons for this include the following.

- Head motion is mostly comprised of rigid deformations. This means that while the object may, for instance, translate or rotate, the space between the points contained within the object do not change. This is in comparison to non-rigid deformations (where the distance between the points contained within the object can change). This concept could be simplified by thinking of a solid and a malleable object. The solid object can be transformed by translating every point of it uniformly but it cannot be deformed, whereas with the malleable object not only can it be translated as a whole but the object itself can be deformed. Respiratory motion is a non-rigid deformation. There is already a large amount of research in the field of rigid deformation motion correction and to a certain degree it could be considered a solved problem (Hill et al. 2001).
- When a PET acquisition is taken of the head, usually measures are taken to immobilise the patient. It is not possible to immobilise the respiratory or cardiac motion of a patient. While it may be true that patients can usually hold their breath during a short CT acquisition, it is not possible for them to hold their breath during a much longer PET acquisition. Thus it is more necessary to correct for these types of inevitable motion.
- Cardiac motion is an autonomous cyclical motion which the patient doesn't have
  conscious control over. Respiratory motion, in contrast, is mostly autonomous.
  However, patients can breath at vastly different rates and depths, and can change
  these parameters over time or completely cease breathing for a period, sometimes
  leading to unpredictable motion patterns. Thus the impact of a method that can
  model non-cyclical unpredictable motion would be greater in the case of respiratory
  motion.

The problems, mentioned above, delay the use of advanced motion management strategies in the clinic. However, further improvements to the method are needed for small lesions at the boundary between the liver, diaphragm and lung. This is because motion in this region is significant and without motion correction it is difficult to quantify or visually assess these lesions. Moreover, a preliminary method to align a single breath hold CT and respiratory gated PET has been developed (Bousse et al. 2016b), (Bousse et al. 2016a). However, this method is likely to be too slow for clinical applications and challenges may arise with larger magnitude or complex motion. This is the case due to this method still relying on co-registration of each gate, which is sensitive to noise. The performance, or susceptibility to noise, and ideally computation time, could be improved by incorporating motion modelling.

### **1.2.1** Objectives of this Work

The aim of this project is to formulate a method which produces PET/CT images that are corrected for respiratory motion and automatically aligned between PET and CT data. This will be achieved through Data Driven Gating (DDG) and motion modelling with minimal impact on the patient and clinical environment, without increased dose, and without increasing scanning time. Ideally the work flow of the method is that of one which is transparent to both the patient and the clinicians, increasing the likelihood of clinical adoption. Evaluation will be performed on simulated and patient data with a comparison to current academic and industry methods.

### 1.3 Overview of this Thesis

The overview of the physics underlying the work is given in **Chapter 2**, of this thesis. Firstly through an introduction to the physics of PET before moving onto describing how the distribution of the radiotracer is quantified by the scanner. After which the different types of scans which can be performed in this manner and the problem of attenuation (why and how it is solved) are expanded upon. Finally, the manner in which the quantified raw data is processed and output is highlighted. This chapter then moves onto discussing the advantages of combined PET/CT scanners, and the method through which the raw data from this are taken and processed in order to give an anatomical or functional image of the patient. Furthermore, the problem of motion is touched upon before introducing the ways in which it can be corrected for, including registration, respiratory gating, and finally covering motion modelling.

Chapter 3 (Whitehead et al. 2019), (Whitehead et al. 2020), (Whitehead et al. 2021) and Chapter 4 (Whitehead et al. 2022c) of this thesis present the development and evaluation of respiratory motion correction methods for PET/CT imaging. Chapter 3 investigates the feasibility of using motion modelling for motion correction, with a particular focus on the impact of Time of Flight (TOF) information on the accuracy of DVF derived from Non-Attenuation Corrected (NAC) data. Simplified simulations are used to explore the effectiveness of various strategies under controlled conditions. Chapter 4 expands upon this by applying the proposed framework to more realistic Four Dimensional Extended Cardiac Torso (XCAT)-based simulations that incorporate variability in respiratory motion, noise, and mismatched  $\mu$ -Maps. This chapter also examines the use of Maximum Likelihood Activity and Attenuation Correction Factors Estimation (MLACF) reconstruction and complex gating schemes to address limitations identified in the earlier feasibility study.

**Chapter 5** (Whitehead et al. 2022b), (Whitehead et al. 2024) extends the motion correction framework developed in the previous chapters to dynamic PET data. Specifically, it focuses on the development and evaluation of DD methods for extracting respiratory SSs directly from dynamic acquisitions, which is essential for enabling motion correction in settings where external hardware or conventional gating is impractical. As with earlier chapters, the structure includes an introduction to the problem, a detailed description of the methods, presentation of the results, and their evaluation.

**Chapter 6** synthesises the findings from the preceding chapters, offering a critical reflection on the work presented. It highlights the methodological contributions of the thesis, discusses current limitations, and outlines key directions for future research aimed at advancing the clinical translation of the proposed motion correction framework.

Additional work conducted during the writing of this thesis is presented in Appendix B (Whitehead et al. 2022a) and Appendix F (Whitehead et al. 2023a). These projects contributed to ongoing collaborations and helped shape the development of related methodologies in PET image analysis and kinetic modelling, broadening the scope of this thesis and its applications. One piece of work explores the use of Deep Image Prior (DIP) denoising as a preprocessing step for improving kinetic modelling in dynamic PET, while the other investigates pseudo-Bayesian Neural Network (NN)-based approaches for ex-

### 1.3. Overview of this Thesis

tracting a metabolite-corrected Arterial Input Function (AIF) in dynamic [<sup>11</sup>C]-Peripheral Benzodiazepine Receptor ([<sup>11</sup>C]-PBR28) PET studies. Although these projects focus on different challenges within the broader field of PET imaging, they helped broaden the methodological foundation that supported the main work of this thesis.

# **Chapter 2**

# **Background**

### 2.1 Introduction

This chapter of the thesis contains the background to the thesis and a literature review of the current state-of-the-art methods, which are to be expanded in later chapters. The first subsection of this chapter introduces a general background to the PET scanner, the physics of its operation, its use, and its combination with other medical imaging modalities.

The second section of this chapter highlights inverse problems in general, before the third section moves on to show how inverse problems are related to the reconstruction of PET acquisition data into a volume representing changes in metabolic processes inside the body of the patient.

The fourth section introduces the problem of respiratory motion in PET and how this can have a significant negative impact on the volumes mentioned in the previous sections. The fifth section proposes methods to address the aforementioned challenge. For instance, the concepts of registration, SS extraction, respiratory gating, and motion modelling (which will be of paramount importance in the following chapters).

The final section brings together the information from the previous sections to highlight a summary of the challenges present in the field and to motivate the rest of the work conducted in the thesis.

### 2.2 **PET**

The physics of PET is introduced in this section. First a general overview of PET is given, including what and how it images (in basic terms), and the use of this clinically.

This general overview is followed by a description of a PET scan, from the compounds which are administrated (to the patient and detected by the scanner) up to how these compounds 'break down' (decay) in the patient and the interaction of the by-products of this annihilation with matter. Along the way, common clinical scanning procedures will be highlighted, as well as issues related to the Field of View (FOV) of the scanner and how these are approached when larger regions of a patient are to imaged.

Secondly, a subsection of the thesis deals with the physical way in which the PET scanner detects individual events. Including, discussing the programmatic ways in which events are determined to be associated. For instance, the timing and energy windows of the scanner, as well as introducing the concept of TOF and how this can affect the data acquired. Next, this subsection covers the expected output data format from the scanner, as well as the effects which determine the maximum resolution of this output.

The final subsection of this section addresses the combined PET/CT scanner. Here, the physics of CT are briefly discussed. Then, methods of Attenuation Corrected (AC) are introduced and CT derived  $\mu$ -Maps are covered. Advantages of AC are supplied before potential pitfalls are highlighted and competing solutions are listed.

### 2.2.1 The Physics of PET

PET is an example of a type of modality known as functional imaging. It is functional because rather than directly capturing images of anatomy (for instance the structure and density of bones, as in CT) it images the metabolic processes. Phantoms can also be imaged which have no metabolic process, this is because it is the radiotracer which is being detected and not the metabolic process itself. This metabolic function is exemplified by how blood flows through and into parts of the body (perfusion) or how glucose is transported to and metabolised by certain cells (using Fluorodeoxyglucose (FDG)). This is useful because, in the case of imaging glucose metabolism, it is possible to quantify the amount of energy that a tissue is using (how much metabolic energy it is demanding). Some cancerous tissues make use of far more energy than non-cancerous tissues. For this reason, PET imaging can be used to diagnose and stage some types of cancers. Moreover, some cancerous tissues are difficult to observe anatomically. For instance, they may have a similar density to the surrounding healthy tissue and therefore may not

be easily detectable on CT. This would be a case where functional imaging could aid with diagnosis. Functional modalities are also useful to highlight increased or decreased metabolic demand or blood flow, which relate to the underlying brain activity.

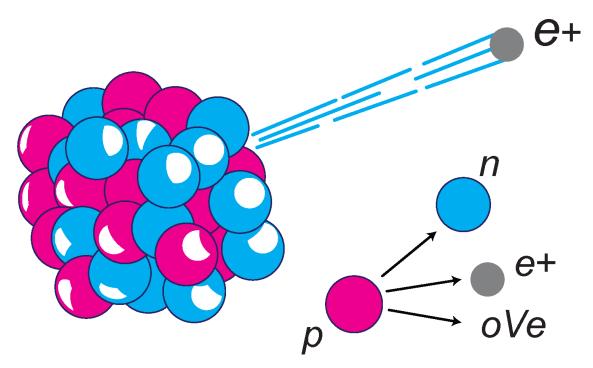
#### 2.2.1.1 Radiotracers

The process through which a PET scan takes place is as follows. Firstly, a patient is injected with a chemical compound called a radiotracer. This is a biologically active molecule which has been labelled with a positron emitting radionuclide. The molecule will have been selected knowing that it has significant uptake in the Region of Interest (ROI), depending on the target tissue.

Some examples of common radionuclides used in radiotracers include fluorine-18, gallium-68, and rubidium-82. The rate at which a radionuclide decays is measured in terms of its half-life. The half-life is defined as the amount of time in which the number of atoms of a radioactive material reduces by half, on average. Half-lives for various radioisotopes can range from a few microseconds to billions of years. For the examples above, these half-lives are approximately 110 min, 66 min, and 66 s respectively (Delbeke et al. 2006).

An example of the use of some of these radionuclides are as follows.

- Glucose molecules (specifically FDG) can be labelled with fluorine-18 and thus called Fluorine-18 Fludeoxyglucose ([<sup>18</sup>F]-FDG). Glucose is used by cells, through glycolosis, in the carbohydrate metabolisation process to produce Adenosine Triphosphate (ATP), which makes energy available to a cell. When a cell requires more energy it also requires more glucose and as such uptake of [<sup>18</sup>F]-FDG is increased in these regions. The concentration of fluorine-18 intra-cellularly increases in certain areas over time because [<sup>18</sup>F]-FDG cannot be fully metabolised. Thus the distribution of fluorine-18 is a good reflection of glucose metabolisation and uptake over time. [<sup>18</sup>F]-FDG is by far the most commonly used radiotracer in PET (Weiss 2016), (Delbeke et al. 2006).
- Gallium-68 is often used to label a radiotracer that targets Prostate Specific Membrane Antigen (PSMA), and can be used in the detection of prostate cancer. PSMA is a protein which is present in prostate cancer cells (Afshar-Oromieh et al. 2013).



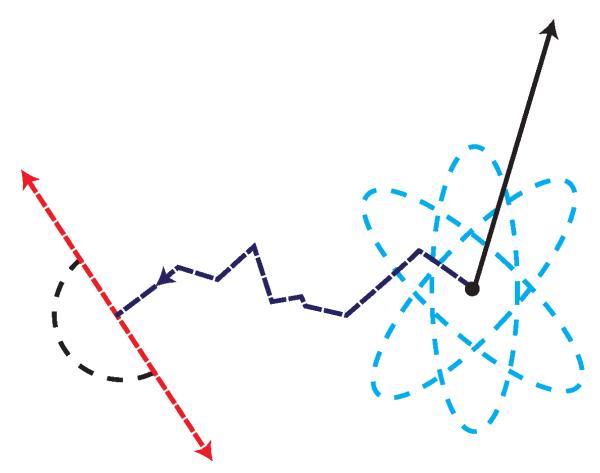
**Figure 2.1:** Graphical example of  $\beta$ +-decay. Here in the top left of the figure a nucleus can be seen which is unstable as it has an imbalance of protons and neutrons. A positron can be seen exiting the nucleus as a by-product of  $\beta$ +-decay converting a proton into a neutron. In the bottom right of the figure a closer example of this can be seen. Here it directly shows a specific proton and the neutron, positron, and neutrino which are produced by  $\beta$ +-decay.

• Rubidium-82 can be used to image the heart in a scan targeting myocardial perfusion (Selwyn et al. 1982).

## 2.2.1.2 Decay and Annihilation

Radionuclides used in PET undergo a type of decay called  $\beta$ +-decay (Conti and Eriksson 2016). This is due to an instability of the radionuclide, because of an imbalance in the number of neutrons to protons in the nucleus. As a consequence, a proton in the atoms nucleus is converted into a neutron, subsequently releasing a positron and a neutrino. This can be seen in figure 2.1. The emitted positron travels with some decreasing velocity, because of sequential collisions, through the body of the patient, for a distance called its positron range. When the positron has dissipated part of its kinetic energy, it will eventually collide with and annihilate its antiparticle the electron. This can be seen in figure 2.2 (Ishkhanov 2012).

The annihilation of positron and electron causes the emission of two 511.0 keV  $\gamma$ photons at an angle of  $\approx 180.0^{\circ}$  apart from one another (thus travelling in approximately



**Figure 2.2:** Graphical example of positron range. Here on the right of the figure an atom can be seen travelling with some velocity, a positron is emitted from the nucleus of the atom by  $\beta$ +-decay. The path that this positron takes can be seen in the centre of the figure represented by a blue line, this path is the positron range. On the left of the figure the annihilation of the positron with an electron occurs and the  $\gamma$ -photons emitted  $180.0^{\circ}$  apart are shown.

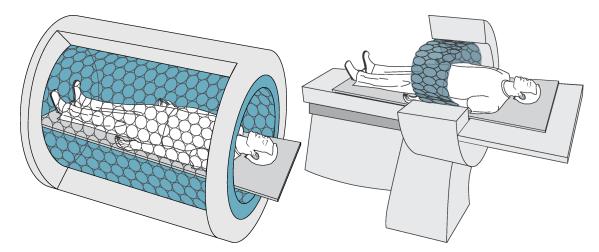
opposite directions). However, because the positron-electron pair may not be at rest at the moment of their annihilation, the two emitted photons can show a certain degree of non-collinearity, according to the laws of conservation of momentum. This means that the  $\gamma$ -photons are almost never exactly 180.0° apart (Khalil 2016).

A PET scanner thus does not image directly the emission of the positron, but in fact more closely images the approximate location of its annihilation.

## 2.2.1.3 Static and Dynamic Acquisition

There are two main types of PET scan useful for determining separate processes. These types of scans and uses are as follows.

• The first and most common type of PET scan is a static PET scan (Muzi et al.



**Figure 2.3:** Graphical representation of the difference between a total body PET scanner and a standard PET scanner. On the left of the figure a total body PET scanner can be seen, where the rings of detectors completely engulf the patient. However, on the right of the figure a standard PET scanner can be seen where the rings of detectors only cover a portion of the patient. For the case on the right of the figure, in order to take a scan over the entire body, either individual acquisitions will be needed and concatenated, or the bed would have to move while the acquisition was ongoing.

2012). The patient is scanned only when the injected radiotracer has distributed throughout their body and eventually approximately stabilised. The time elapsed between injection and acquisition depends on the half-life and metabolisation of the radiotracer. For [18F]-FDG a delay of about 60 min is used.

• The second type of PET scan is a dynamic scan. The acquisition of data for this scan begins before the radiotracer is injected into the patient. The injection of the radiotracer during the acquisition allows for the kinetics of the radiotracer to be observed and quantified with the use of compartmental modelling (Lammertsma 2017). For example, from dynamic PET myocardial perfusion imaging in-vivo studies used in conjunction with radiotracer kinetic modelling enables the quantification of myocardial blood flow, often measured using rubidium-82.

#### 2.2.1.4 PET Field of View

The FOV of the scanner is the area in which  $\gamma$ -photons can be detected. Current clinical PET scanners, usually, have a cylindrical FOV with a length of between 15.0 cm and 25.0 cm and a diameter of between 50.0 cm and 70.0 cm (Pan et al. 2019).

There are multiple ways to acquire data over more than the axial length of the scanner,

three of these methods are as follows.

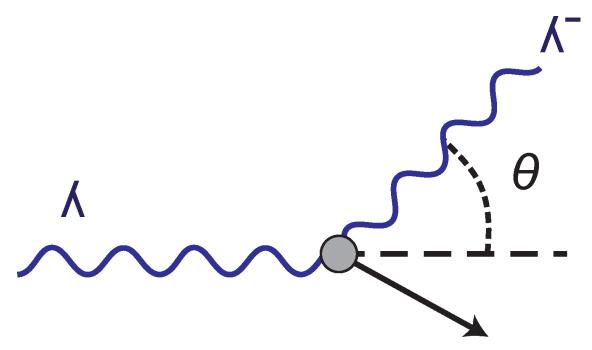
- The most simple and widely used method is to take acquisitions over multiple bed positions and concatenate them.
- A method available on some standard axial length scanners is to continually move the bed through the rings of the scanner while acquiring data. This is advantageous as it is more comfortable for the patient and provides potentially less movement of the patient. A disadvantage of this though, is that it introduces another source of motion to the acquisition from moving the bed. This makes standard motion correction much more difficult.
- Alternatively, total body PET scanners are becoming more viable for research. Total
  body PET scanners have an axial FOV which contains most of the patients body
  (making multiple acquisitions less necessary) while also increasing the sensitivity
  of the scanner, this can be seen in figure 2.3 (Cherry et al. 2018). However, the
  increased price and size constitute a limitation.

#### 2.2.1.5 Attenuation

Attenuation is the process or effect through which counts are 'lost' (from annihilation to detection by the scanner), some become scatter events, while the photons are traversing through the body of the patient. Attenuation can amount to a loss of up to 95.0% of the total initial signal and can cause increased issues in larger bariatric patients (Bailey et al. 2006), (Mettler and Guiberteau 2012).

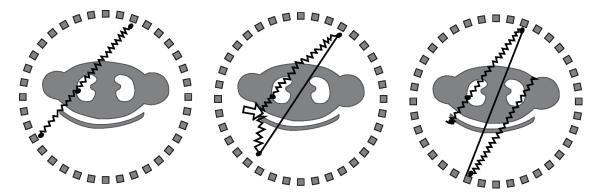
There are three main ways through which the photon signal can be 'lost' (Khalil 2016), these are in ascending order of magnitude as follows.

Rayleigh scattering is the elastic scattering of photons, without loss of significant energy, by particles which are much smaller than the wavelength of the photon.
 A common example of Rayleigh scattering is the scattering of sunlight in the atmosphere, which causes the blue colour of the sky during the day and the red colour of the sky at sunset. Because the wavelength of γ-photons are comparably small, compared to most particles, the probability of Rayleigh scattering occurring is negligible and thus it is normally ignored in PET.



**Figure 2.4:** Graphical representation of a  $\gamma$ -photon scattering off of a particle. The  $\gamma$ -photon can be seen entering the figure on the left hand side, before scattering off of a particle in the centre of the figure by an angle  $\theta$  and exiting the figure on the top right hand corner. The electron exits the scatter event with some velocity represented by the arrow towards the bottom right.

- Absorption through the photoelectric effect is the process through which a high energy γ-photon hits and transfers its energy to a material, causing the emission of a lower energy electron. The likelihood of the photoelectric effect is inversely proportional to the cube of the photon energy. It also increases as the atomic number of the attenuating material increases. In the matter of the patient, the photoelectric effect is most prevalent at photon energies below 100.0 keV and as such the probability of the photoelectric effect occurring for the PET γ-photons is minimal (Bailey et al. 2006). Attenuation through the photoelectric effect occurs mostly in the detectors of the scanner (it's the principle underlying the function of scintillator crystals).
- Compton scattering comprises the majority of interactions between the photon and matter in PET. It occurs where the photon interacts with an electron in a close by atom. The recoiling electron causes the photon to be deflected along another path, transferring energy from photon to electron, this can be seen in figure 2.4. Compton scattering is also known as incoherent scattering. The probability of Compton



**Figure 2.5:** Graphical representation of the different types of coincidences possible in PET. On the left of the figure a true coincidence can be seen, this is where the  $\gamma$ -photons from one annihilation are both detected without scattering. In the middle of the figure a scattered coincidence can be seen, this is where the  $\gamma$ -photons from one annihilation are both detected, however in this case one of them has scattered. On the right of the figure a random coincidence can be seen, this is where the  $\gamma$ -photons from two unrelated annihilations are detected.

scattering is indirectly proportional to the energy of the photon (Bailey et al. 2006).

The relationship between the attenuation of the signal and the material through which it is travelling is given by the Beer-Lambert law. Given  $I_0$  incident photons travelling across a path D, the number of non-scattered photons  $I_D$  is given by

$$I_{D} = I_{0} \cdot \exp\left(\int_{D} -\mu_{E}(r) dr\right)$$
(2.1)

where in equation 2.1  $\mu_E(r)$  is the attenuation coefficient of the media crossed by photons of energy E.

## 2.2.2 Data acquisition

As discussed in section 2.2.1.4, the structure of a PET scanner is that of concentric rings of detectors offset along a central axis. These rings detect each incident photon and attempt to temporally and spatially link opposing photons along a LOR through the scanner. A LOR being a line through the FOV of the scanner linking two detectors. The methods through which the scanner attempts to link related photons together will be discussed in the following section 2.2.2.1.

Because of the photons interaction in matter shown in section 2.2.1.5, there are four different types of event or coincidences that can be detected by the scanner. These are as follows.

- Firstly, the coincidences that originate from the same annihilation event and pass
  through the body of the patient to the detector without being scattered or attenuated.
  These coincidences are called true coincidences as they approximately accurately
  reflect the position of the originating annihilation and thus the underlying activity
  distribution.
- Secondly, there are coincidences which may have originated from the same annihilation event, but from which one or more of the incident photons has undergone Compton scattering before detection. These coincidences are called scattered coincidences. Scattered coincidences can attempt to be corrected for if an accurate μ-Map is given. The density of the matter through which the photons must have travelled indicate the likelihood of scattering.
- Thirdly, there are coincidences where the LOR is determined from two photons from two distinct annihilation events (thus the LOR does not reflect an actual annihilation in reality). This could occur because one or more of the photons, from the original pair of photons, may have been attenuated or scattered, so that it does not arrive at the detector within a reasonable time of its photon pair or that its LOR doesn't go through one of the detectors. These are called random coincidences. Random coincidences can be corrected for if acquisition data of this background level of the scan exists.
- Fourthly, there could be a situation where three or more photons are detected within close temporal proximity to one another. Because of the close time of detection, in this case it is not possible to determine which photons reflect an actual annihilation and which are random coincidences. These coincidences are called multiple coincidences. In normal procedures this is rare.

An example of some of the types of coincidences from above can be seen in figure 2.5.

The total prompts detected during a PET acquisition P can be expressed as

$$P = T + S + R \tag{2.2}$$

where in equation 2.2 T is the number of true coincidences, S is the number of scattered coincidences and R is the number of random coincidences.

## 2.2.2.1 Coincidence Processing

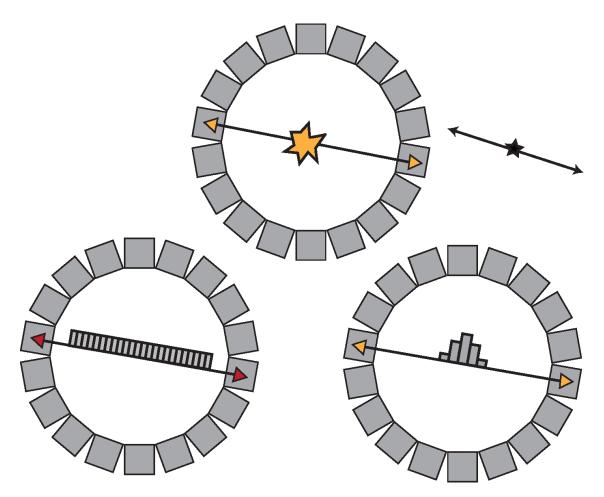
In order to form coincidences, the incident photons must be paired together to an annihilation event. First, before forming coincidences, the photons are filtered by selecting ones which only fall within an energy window of the scanner. For the General Electric (GE) Discovery 690/710 PET/CT this energy window falls approximately between 425.0 keV and 600.0 keV (Bettinardi et al. 2011). Additionally, to attempt to determine temporally if two detected photons belong to the same annihilation event a timing coincidence window is used. If the events arrive more than the time of the coincidence window apart then they are determined to be unrelated. A standard coincidence window size would be about 5.0 ns.

## 2.2.2.2 Time of Flight PET

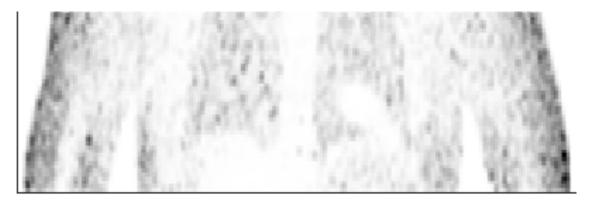
As stated above in section 2.2.2.1, in order for the coincidence window based method to determine which specific detected photons represent LORs, the scanner must be able to asses the difference in arrival time of each photon. It had been hypothesised for some time (since the 1960s) that given the speed of light and the difference in arrival time at each detector, for each photon that makes a specific LOR, then it should be possible to approximately calculate the position upon the LOR at which a given annihilation occurred (Surti 2015), (Spanoudaki and Levin 2010). This can be seen in figure 2.6.

The reason for the uncertainty of the position along the LOR is because of the relatively course timing resolution of each given scanner. Generally modern PET/CT scanners have a timing resolution ranging between 200.0 ps and 600.0 ps. This represents an approximate spatial uncertainty of between 30.0 mm and 90.0 mm. The uncertainty within these TOF bins is usually modelled using a Gaussian distribution centred around the estimated position of annihilation by the scanner, with standard deviation dependent on the time resolution.

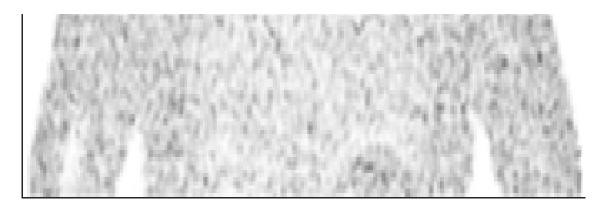
Currently, TOF is a focus for research because of the drastic improvements that it can have on the Signal to Noise Ratio (SNR) (Lecoq 2017), (Cates and Levin 2018). An example of some NAC Non-TOF and NAC TOF data can be seen in figure 2.7 and



**Figure 2.6:** Graphical representation of the concept of TOF. The top middle of this figure shows the position where a hypothetical annihilation has occurred (plus the  $\gamma$ -photos from this annihilation which have then gone on to be detected by the scanner). The bottom left of this figure shows a traditional Non-Time of Flight (Non-TOF) acquisition, where the probability of the position of the annihilation along the LOR is constant. The bottom right of this figure shows a TOF acquisition, where the probability of the position of the annihilation along the LOR can be approximated with a Gaussian based upon the difference in arrival time of both photons.



**Figure 2.7:** Example of some NAC Non-TOF data, with noise, with no motion, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.



**Figure 2.8:** Example of some NAC TOF data, with noise, with no motion, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.

figure 2.8 respectively. Notice the difference in distribution of counts in the centre of the thorax and lungs. In the Non-TOF case, counts have been disproportionately placed at the periphery of the patient. However, in the TOF case, counts can be seen more evenly distributed throughout the patient.

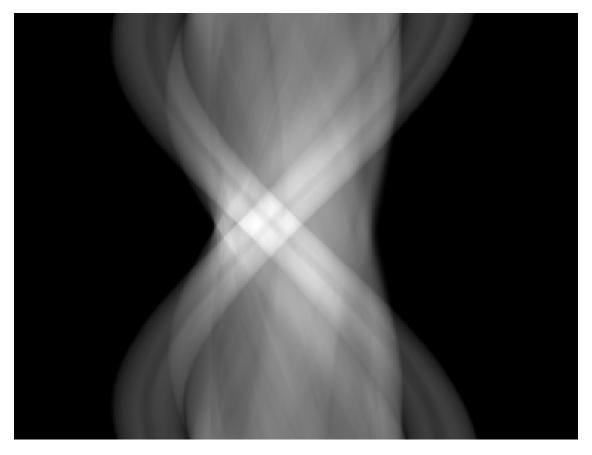
The current (as of 2021) PET/CT scanner with the highest TOF resolution, which is commercially available, is the Siemens Vision with an approximate Full Width at Half Maximum (FWHM) of 210.0 ps or 31.5 mm (Van Sluis et al. 2019). The PET/Magnetic Resonance (MR) scanner with the highest TOF resolution, which is commercially available, is the GE Signa with an approximate FWHM that is sub 400.0 ps or 60.0 mm (Grant et al. 2016), (Hsu et al. 2017), (Caribé et al. 2019).

## 2.2.2.3 Data Output

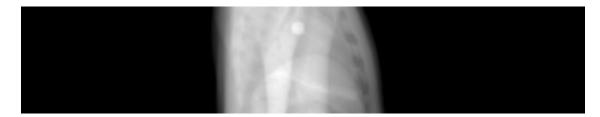
The output from a PET scanner must be stored in a file format in order to be of any use. This file format will usually contain information related to the prompts from the acquisition, discussed above in section 2.2.2. Each prompt stored represents a LOR connecting the centre of two detectors. Where TOF information is available it is stored as an extra dimension in this file. This file can then be taken and reconstructed in order to estimate the original distribution of the radiotracer, this will be discussed later in section 2.4.

There are two main formats in which this information is stored from the scanner.

• The most common way, that is used in clinical practise (as of 2023), is a format called a sinogram. During acquisition, if a sinogram output is specified, then the



**Figure 2.9:** Example of some simulated sinogram data, with no motion, noise, randoms or scatters, of the thorax with a spherical lesion in one of the lungs.



**Figure 2.10:** Example of some simulated viewgram data, with no motion, noise, randoms or scatters, of the thorax with a spherical lesion in one of the lungs.

coincidences detected by the scanner are binned into a histogram which represents their plane orthogonal to the scanner, their orientation angle, their average axial location, and their distance from the centre of the gantry. TOF can be added as an additional dimension if it is used. If a single point source were imaged it would produce a sinusoid when binned into a sinogram, hence the name. An example of sinogram (where the LOR are in a given transaxial plane) and viewgram (where the LOR are along a given view, parallel to each other) data, with no noise, can be seen in figure 2.9 and figure 2.10 respectively. Because data is being binned into a histogram with this method, information is lost, it could be considered a lossy compression method.

A less common method, but one which is becoming more prevalent, is a format
called listmode data. Here each coincidence is recorded sequentially in a file. The
information stored for each coincidence includes its arrival time, the coordinates
of the detector and its detected energy. TOF information can also be stored if it is
used. A listmode file can be directly reconstructed or first unlisted into a sinogram
post acquisition.

#### 2.2.2.4 PET resolution

There are five main effects which impact the resolution of a PET acquisition, these are as follows.

- Firstly, there is, as has been discussed above in section 2.2.1.2, the effect of positron range. Because the positron travels a small distance before undergoing annihilation the PET scanner will always, at best, be measuring the position of the annihilation rather than the position of the decay and as such not directly measuring the position of the radiotracer (Levin and Hoffman 1999).
- Secondly, again as discussed above in section 2.2.1.5, acollinearity of the  $\gamma$ -photons introduces errors which affect the resolution of PET. This is because the positron will almost always enter the annihilation event with some velocity then the  $\gamma$ -photons produced will exit with the same additional velocity. This effect is exacerbated by the amount of time that the photons are allowed to travel, thus the larger

the bore of the PET scanner the larger this effect will have on the resolution. The effect of acollinearity on [<sup>18</sup>F]-FDG gives an error of approximately 0.54° (Shibuya et al. 2007)

- Thirdly, the size and number of detectors in each ring and the thickness of each ring dictates the maximum resolution of the reconstruction before voxels begin to potentially have zero LORs pass through them.
- Fourthly, the block construction of each detector negatively impacts the resolution. This is because a number of scintillation crystals is paired with, usually, fewer photodetectors This means if a photon interacts with one crystal it may incorrectly be attributed to another crystal (Nieman et al. 2015). So called digital PET scanners are beginning to be seen which use a one to one ratio of photodetectors to scintillation crystal, effectively removing issues related to block construction (Schillaci and Urbano 2019).
- Finallya scintillation crystal has a stopping power. This power describes the approximate depth at which photons will undergo attenuation by the photoelectric effect. In some instances, depending upon the position and angle at which the incident photon hits the scintillation crystal, it is possible for the photon to travel through the crystal and into an adjacent crystal before being detected. This means that the photon is incorrectly positioned and will result in blurring of the reconstructed volume (Nieman et al. 2015).

#### 2.2.3 Combined PET/CT

A CT scanner consists of two devices which sit one on either side of the bore of the scanner. One device is an X-ray emitter and the other is an X-ray detector. If the X-ray emitter were to operate in one fixed position the result would be similar to a standard diagnostic X-ray. The difference here comes from the fact that for CT, during a continuous acquisition, the device moves around the axis of the scanner taking continuous measurements. This allows for an X-ray image at every angular position. While the CT is acquiring, the bed of the scanner travels along the axis of the scanner, this allows for the collection of data over a Three Dimensional (3D) volume (similar as to what PET collects over). This method

of acquiring 3D data by taking measurements at a number of angles, or section, around a subject using a penetrating wave is known as tomography and is a similar concept to the tomography of PET.

When the X-ray beam intersects with the body of a patient, it is possible for the beam to be attenuated by the photoelectric effect or scattered, similarly to discussed in section 2.2.1.5. The intensity of the detected X-ray correlates with the density of the material between the emitter and the detector. Where the intensity of the beam detected is less then expected, it can be assumed that there is a higher density object attenuating more of the beam between it and the emitter. If this information is collected over a 3D volume it allows, for instance, for the generation of a 3D volume that reflects the attenuation of the body of the patient. This attenuation is normally expressed in Hounsfield Units (HUs).

The energy of the X-ray used in CT consists of many different wavelengths, hence it is polychromatic. The wavelength range usually used in a PET/CT acquisition is between 40.0 keV and 140.0 keV (Alvarez and MacOvski 1976). The wavelength range is determined by the settings of the scanner, these mainly consist of the peak kV used and the electric current applied, in mA.

In a standard PET/CT acquisition the CT component comes first before the PET. In modern PET/CT scanners the CT and PET are inline with the same bed. However, the first PET/CT acquisitions were taken on different machines entirely and as such the position of the patient differed more drastically between acquisitions. A standard CT scan (as part of a combined PET/CT acquisition) over the thoracic region specifically will last approximately between 2.0 s and 3.0 s (Goerres et al. 2003).

#### 2.2.3.1 Attenuation Correction

As discussed previously in section 2.2.1.5 attenuation represents the loss of coincidences by photon interactions in matter. Attenuation is an issue in PET as it causes the loss of signal and a degradation in image quality. The opposite is true for CT where the modality itself relies on attenuation in order to differentiate anatomical structure. In order to find reasonable quantitative results the attenuation of the patient must be taken into account in PET. An example of a  $\mu$ -Map can be seen in figure 2.11. An example of some NAC Non-TOF and AC Non-TOF reconstructed data can be seen in figure 2.12 and figure 2.13



**Figure 2.11:** Example of a simulated  $\mu$ -Map, with no motion or noise, of the thorax with a spherical lesion in the lungs. Coronal view.



**Figure 2.12:** Example of simulated NAC Non-TOF reconstructed data, with motion, with no noise randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view.



**Figure 2.13:** Example of simulated AC Non-TOF reconstructed data, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view.

respectively. Similarly to as seen previously in figure 2.7, in figure 2.12 without AC the counts present in the data are placed by reconstruction at the periphery of the patient (the counts that remain within the patient are still misquantified). However, in figure 2.13, similarly to as seen previously in figure 2.8, with AC the counts are better distribution throughout the patient. However, in contrast to in figure 2.8, in figure 2.13 this time rather than the change being the addition of TOF it is instead the addition of AC.

There are multiple methods to acquire a  $\mu$ -Map for AC. One method would be to take a transmission scan of a known point source rotated around the body of the patient prior to the injection of the radiotracer. This allows for the estimation of the attenuation for each angle (Bailey 1998). Another method involves the use of the known attenuation from the CT scan.

In order to apply the CT based  $\mu$ -Map to AC in PET, first it must undergo either bilinear or trilinear conversion to asses the Attenuation Correction Factors (ACFs). This is because of the relative energy difference of the two modalities (Carney et al. 2006). As discussed above in section 2.2.2.1 and section 2.2.3 PET and CT operate at two different energy levels, of between 425.0 keV and 600.0 keV and between 40.0 keV and 140.0 keV respectively (Bettinardi et al. 2011), (Alvarez and MacOvski 1976).

Issues with CT based AC include the following: Firstly, as mentioned above in section 2.2.3, CT is acquired sequentially to PET rather than simultaneously (like MR and PET in a combined PET/MR) meaning that there can be mismatches in anatomy between the scans. Secondly, the propagation of any artefacts from the CT volume into the PET volume. Regardless of these issues CT is currently considered to be the gold standard for  $\mu$ -Map estimation for AC. Transmission scans are now very rarely used for many reasons, including, because of the necessity of an additional external source, their significant increase in scan time, and low image quality compared to CT.

Recently a method to jointly estimate both the activity and attenuation distributions from PET data only has been proposed called Maximum Likelihood Reconstruction of Activity and Attenuation (MLAA) (Nuyts et al. 2013). This method attempts to reconstruct both distributions by iteratively estimating one distribution while keeping the other one fixed. This can be considered as at each step performing either emission or transmission tomography depending on the distribution being optimised for (Fuin et al. 2017),

(Brusaferri et al. 2020). A disadvantage of this solution is that without TOF information it is a highly ill-posed problem and even with TOF data the solution can only be found up to an arbitrary scaling factor. It is also possible for there to be significant cross-talk artefacts between the activity and attenuation estimate (Salomon et al. 2011), (Defrise et al. 2012). Additionally, optimising for the attenuation distribution increases the complexity and computational effort required of a reconstruction algorithm.

A variation of MLAA called MLACF has also been developed. Here, rather than optimising for the  $\mu$ -Map instead during the transmission optimisation step the ACFs are updated (Nuyts et al. 2012), (Rezaei et al. 2012). This offers a significant speed improvement over MLAA, without impacting the quality of reconstructed output. However, if the ACFs are found rather than the  $\mu$ -Map, then if a  $\mu$ -Map is required for subsequent computation this method is inappropriate. Furthermore, this method still suffers from the other disadvantages of MLAA.

More details about MLAA and MLACF are given in the image reconstruction section below in section 2.4.3

## 2.3 Inverse Problems and Optimisation

This section of the thesis introduces the concept of inverse problems. First, a definition of an inverse problem is given, including where these may be applicable to the field of PET. Next, the definition of an inverse problem is expanded upon by highlighting the difficulty of solving them analytically and how this is usually overcome or addressed. The general form of an iterative approach to solving an inverse problem is presented.

The second subsection expands upon the approach given previously to solving inverse problems. This Includes defining what optimisation is, applications of optimisation (for instance, PET reconstruction) and listing the requirements of a simple optimisation problem. Next, this subsection moves onto addressing the components of an optimisation, including the objective function. The purpose of an objective function in optimisation is presented before a number of common and robust similarity measures are introduced, the merits of different functions are discussed and common applications of specific instances are given. Regularisation is then briefly mentioned before moving onto the optimiser. As with the previous subsection, here, the purpose of an optimiser is initially introduced be-

fore families of optimisation algorithms are compared and variations of these families, as well as their applications, are addressed. Finally bounds or constraints upon optimisation are mentioned, as well as providing examples of optimisers which can incorporate such functionality.

## 2.3.1 Inverse Problem Concepts

An inverse problem is one where the original conditions of a system are estimated from its effects. For instance, the data from a PET scanner represents the observations of the distribution of the radiotracer, reconstruction is an attempt to find the distribution from these observations.

However, direct inversion is often not feasible due to the complexity, noise, or ill-posed nature of the problem. As a result, solutions are typically obtained through iterative optimisation methods that gradually approximate the best-fit reconstruction. In order to attempt to find the solution to an inverse problem there are two things which are required. First, the forward operator, second, ideally, a model of the noise present in the system (Brusaferri 2020), (Emond 2020).

## 2.3.2 Optimisation Concepts

Optimisation means to find values that best parametrise a given function based on some criteria or objective. For instance, in image reconstruction an optimisation could be to find the image that when the forward operator is applied to it, the result best matches the measured data. Another example of an optimisation would be to find the motion parameters that when applied to a given image most closely deform that image to match another image. Optimisation is also used in fields such as deep learning to train NNs, here the optimisation is used to find parameters for a model that maps one set of values to another.

In order to perform a basic optimisation four components are required. Firstly, an objective function (which returns the goodness of the current estimate) and a method to update the estimate based on this function are paramount. Additionally, an initial estimate (the closer this is to the ideal estimate then the less computation time is required overall) and a method to determine when optimisation should cease (for instance when the current objective function value or the number of iterations exceeds a threshold) are

also necessary. These will be discussed in the following sections in section 2.3.2.1 and section 2.3.2.2 respectively.

#### 2.3.2.1 Objective Function

Optimisation of some values requires a function that represents, for instance, the similarity of two measures or the likelihood of a measure. This function is necessary as its output reflects the accuracy of the current estimate. The gradient of the objective function describes the direction in which the optimiser should update the estimate. The optimiser attempts to find a solution by either maximising or minimising the result of applying the objective function and updating the estimate iteratively.

An example of an objective function would be Mean Absolute Error (MAE). For a vector of some values, MAE subtracts the estimated value from the true value, finds the absolute value of this and takes the means of all values in the vector. The absolute value is taken because the error should be the distance to the true value regardless of if the estimated value is greater or less than the true value. If the estimated values approach the true values then the value of the MAE will approach zero. MAE suffers from the disadvantage that its gradient is undefined at zero, hence it is not differentiable at all points, which is a problem for some optimisation strategies. If Mean Squared Error (MSE) is used then the square of the error is taken rather than the absolute value, taking the square causes the error to increase quadratically as it becomes larger. This can be advantageous as it penalises large errors more than smaller ones, which can lead to a better result in some circumstances. Additionally, the square is differentiable at all points.

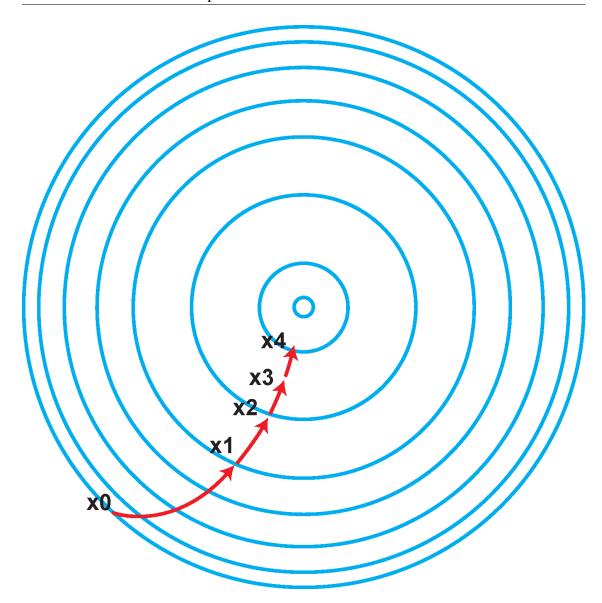
More complex objective functions include Root Mean Square Error (RMSE). Here, the square root of MSE is taken, which scales the value of the error back to the units of the original estimate. Median Absolute Difference (MAD) is similar to MAE, however, rather than taking the mean of all error values the median is taken instead. This would give an objective function which is less sensitive to noise or more robust as outliers are diminished or ignored by the median. An objective function which differs more significantly would be correlation coefficient. Here, the correlation of the values in the estimate are compared to the measured data, thus if correlation coefficient was used in an optimisation then it is not guaranteed that the estimate will have the same scale as the measured data (but its

shape should be similar). Correlation coefficient can be used on very noisy data as it is less sensitive to this noise. MAE, MSE, RMSE, and correlation coefficient are usually used in regression-like problems where a line is fit though a data set. However, correlation coefficient is more often used as an accuracy measure rather than an objective function.

Normalised Mutual Information (NMI) is an example of another type of objective function. NMI is the normalised version of Mutual Information (MI). This is a clustering objective function. Here, values in both the estimate as well as the target are binned into a given number of bins and the clusters and locations of clusters are compared between the estimate and the target. This objective function is useful when the estimate and target are not necessarily in the same scale or units. For instance, this objective function is useful when performing registration of two different imaging modalities. This is because their units are different but it could be expected that they would represent similar underling structure. This will be expanded upon in section 2.6.1.

Another example of an objective function would be likelihood. This function for any given sample of data computes the goodness of fit to a statistical model. Strictly speaking, the likelihood is related to the inverse of the goodness of fit. The likelihood function describes a planes whose peak, if there is one distinct peak, represents the combination of model parameter values that maximise the probability of drawing the sample obtained (Myung 2003). Poisson log likelihood is often used as the objective function in PET reconstruction, this can be seen in section 2.4.

As an additional term to the objective function a regularisation term is often added, summed to the objective function value (after being scaled by an  $\varepsilon$ ). Regularisation terms are used in order to decrease the optimisers sensitivity to noise. Inverse problems are often ill posed, this means they are not well posed. A well posed problem is one for which the following properties hold. Firstly, the problem has a solution. Secondly, the solution is unique. Finally, the solution's behaviour changes continuously with the initial conditions. An ill posed problem breaks at least one of these properties and as such often the problem is highly sensitive to changes in its conditions. An example of a regularisation term is Relative Difference Prior (RDP) used by GE in Q.Clear (Ross 2014). Another example is bending or linear energy which are used in registration to penalise against rapid changes in the DVF (Modersitzki 2009). L1 and L2 are common regularisation methods which can



**Figure 2.14:** Graphical representation of a 2D solution space and an optimiser stepping through this space. In the bottom left of this figure the initial estimate for some optimisation can be seen at x0, the subsequent iterations 1 through 4 can then be seen taking this estimate closer to the centre of the contour plot. Here the contour plot could either be showing a maximisation or a minimisation depending on how it is visualised.

be used in almost any optimisation problem. Here, the L1 is the mean absolute value and L2 is the mean squared value of the model parameters. These methods try to encourage small model parameters. Large model parameters are often associated with over-fitting.

## 2.3.2.2 Optimiser

An optimiser takes an estimate of the solution and an objective function and returns an update of the estimate based on the current objective function value, as seen in section 2.3.2.1. The direction in which the optimiser updates the estimate is based on the gradient of the objective function. The gradient can be determined through three methods. Firstly, if the estimate is permuted slightly the gradient can be approximated, this method is extremely slow due to having to calculate the objective function multiple times (this is the finite differences method). Secondly, a hand crafted gradient function can be provided for each objective function. This method is almost always the fastest to execute but takes the most work to implement. Finally, an automatic differentiation algorithm can be used. Automatic differentiation, as the name suggests, automatically differentiates the objective function to find the gradient function. Automatic differentiation has a slight overhead when compared to a hand crafted gradient function. However, the time saved in implementation is often worth the overhead. NNs are often trained with the help of robust auto differentiation packages, otherwise every time a model was changed at all a new hand crafted gradient function would be needed.

The method by which the optimiser updates the estimate is what differentiates optimisation algorithms. Gradient descent is a commonly used family of optimisation algorithms. Gradient descent itself takes the gradient of the objective function at the current estimate and takes steps, of a given size, in the direction calculated from the gradient. The step size of gradient descent can be set to a fixed value or found using a line search which optimises the step size. Momentum can be used as an improvement of gradient descent where the current update direction is a linear combination, with a predefined weighting, of the current gradient and the previous update direction.

Stochastic Gradient Descent (SGD) is an extension of gradient descent where the current update is calculated using the gradients of a randomly selected subset of the data, this is significantly more computationally efficient than gradient descent as it reduces the number of calculations needed for each update.

Conjugate gradient is another extension of Gradient descent. Here the direction of subsequent updates are confined so that they are orthogonal to the previous update direction, this can decrease convergence time (Tustison et al. 2009).

Broyden Fletcher Goldfarb Shanno (BFGS) and Low memory Broyden Fletcher Goldfarb Shanno (L-BFGS) are second order optimisers in that they take the second order partial derivatives of the objective function. BFGS and its derivatives determine the

descent direction by preconditioning the gradient with curvature information. L-BFGS is differentiated from BFGS in that BFGS stores a dense approximation of the inverse Hessian, whereas L-BFGS stores a history of a past window of updates. Thus L-BFGS uses less memory than BFGS and can therefore converge faster (Fletcher 2000).

An optimiser can also be provided with bounds or constraints, a simple bound would be a box bound where the values of the estimate cannot exceed a threshold. Low memory Broyden Fletcher Goldfarb Shanno Bounded (L-BFGS-B) is an implementation of L-BFGS which accepts box bounds. This can be useful, for instance, in PET reconstruction, where it is not expected that negative values should exist, so they could be constrained.

A different family of optimisers would be exemplified by the Adam optimiser (Kingma and Ba 2014). These optimisers are commonly used to train NNs and are an alternative to SGD. BFGS and its derivatives, which are second order optimisers, are often too memory intensive to train large NNs. Adam uses estimations of the first and second moments of the gradient to adapt the learning rate for each parameter of the model. AdamW is an extension of Adam which incorporates weight decay. Weight decay has a similar effect to L2 regularisation (Loshchilov and Hutter 2019).

## 2.4 PET Image Reconstruction

This section of the thesis follows on from the previous section in that it focuses on the inverse problem of PET reconstructions specifically. First it shows how PET reconstruction is an inverse problem, by likening aspects of the reconstruction problem to those presented previously, before elaborating on the two general families of reconstruction algorithm. Analytical and numerical approaches to PET reconstruction are discussed and finally a common output scale from this process is highlighted.

The first subsection expands upon the numerical approach (or optimisation) to PET reconstruction given previously. Initially the advantages and disadvantages of iterative PET reconstruction in general are introduced. Then common algorithms are described and their operation, advantages and disadvantages are compared.

The second subsection

#### 2.4.1 Introduction

PET image reconstruction is an inverse problem, as stated in section 2.3.1. This means that a PET reconstruction algorithm takes as an argument the effects of the PET acquisition system and attempts to determine its initial conditions, for instance the distribution of the radiotracer.

There are two main types of methods through which a PET reconstruction can be performed. First there are analytical PET reconstruction algorithms, an analytical solution attempts to calculate the exact solution. Secondly there are numerical solutions, for instance, iterative PET reconstruction. Numerical approaches to solving inverse problems in general are discussed in section 2.3.1 and iterative PET reconstruction algorithms are discussed in section 2.4.2.

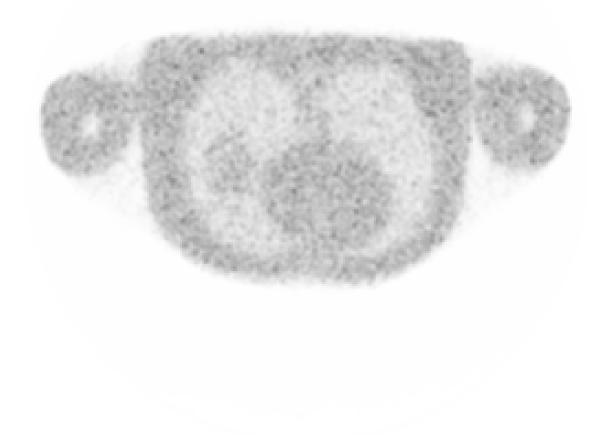
The data output from a PET acquisition are usually expressed in Kilo Becquerel per Millilitre (KBq/mL). However, for pseudo quantitative analysis the values are usually normalised to Standard Uptake Value (SUV) by dividing the activity by, for instance, the mass of the patient and the injected activity.

## 2.4.2 Iterative Image Reconstruction

An iterative method has the advantage that the model which is used can take into account the noise properties of the data and also the physical properties of the scanner (The noise associated with PET is often assumed to be Poisson distributed). However, iterative methods have the disadvantage that they require substantial computational effort to execute when compared to analytical reconstruction.

Maximum likelihood is often combined with the expectation maximisation algorithm for the optimisation of the model parameters and in this case is called Maximum Likelihood Expectation Maximisation (MLEM) (Dempster et al. 1977), (Shepp and Vardi 1982), (Lange and Carson 1984). When maximising the likelihood the natural logarithm of the likelihood, also known as the log-likelihood, is taken for computational efficiency. Here the goal is to reconstruct the emission distribution  $\lambda$ . The image estimate is updated as follows.

$$\lambda_j^{t+1} = \frac{\lambda_j^t}{\sum_i a_i c_{ij}} \sum_i \frac{c_{ij} g_i}{b_i^t}$$
 (2.3)



**Figure 2.15:** Example of a simulated AC Non-TOF Ordered Subset Expectation Maximisation (OSEM) reconstruction, with motion and noise, with no randoms or scatters, of the thorax with a spherical lesion in the lungs. Transverse view.

where in equation 2.3 t is the iteration number,  $b_i^t = \sum_k c_{ik} \lambda_k^t$  is the total activity along the i-th LOR,  $c_{ij}$  is the detector sensitivity, and  $a_i = \exp(-\sum_k l_{ik}\mu_k)$  (Brusaferri 2020). The output from MLEM commonly has a Gaussian blur applied in order to smooth noise (Llacer et al. 1993). Disadvantages associated with MLEM include, for noisy data iterating for too long can cause the output to accentuate the noise present in the data. One way to avoid this is to purposefully cease iterating early before the noise can take over the image (Johnson 1994). Another solution to this problem is to incorporate a regularisation term, into the objective function, which penalises rapid local changes in magnitude. Another issue is that MLEM is exceptionally slow even when compared to other iterative algorithms.

To combat the slow execution speed of MLEM OSEM was developed. In OSEM the LOR or detector pairs of the scanner are binned into a number of subsets. This is in contrast to MLEM where there is logically one subset, MLEM would be applied to all LOR or detector pairs simultaneously. However, in OSEM the LOR or detector pairs could be binned into at least two subsets. MLEM is then applied to each subset, in a specific order, sequentially (Hudson and Larkin 1994). Because the image is updated after each iteration of MLEM on each subset of OSEM then the execution speed of OSEM is increased by the number of subsets used. However, if subsets are used then OSEM will converge to a limit cycle around true convergence (Mettivier et al. 2011). OSEM is to MLEM as SGD is to gradient descent (except for how the subsets are sampled, ordered or stochastic). If a reasonable number (commonly between two to 32 or potentially 64, 24 for the GE Discovery 710) of subsets are used it has been found that OSEM will accelerate MLEM without affecting the accuracy of quantification too drastically (Morey and Kadrmas 2013). An example of an OSEM reconstructed image can be seen in figure 2.15.

# 2.4.3 MLAA and MLACF Methods for Attenuation and Activity Reconstruction in PET

In PET imaging, accurate attenuation correction is essential for both qualitative and quantitative analysis. Traditionally, this relies on external anatomical imaging, such as CT or MR, to provide a  $\mu$ -Map. However, several methods have been proposed to derive

attenuation information directly from the PET emission data, most notably MLAA and its variant, MLACF (Nuyts et al. 2013), (Nuyts et al. 2012), (Rezaei et al. 2012).

#### 2.4.3.1 MLAA

MLAA aims to reconstruct both the activity distribution  $\lambda$  and the attenuation map  $\mu$  directly from the PET emission data by alternately updating each component (Nuyts et al. 2013).

The algorithm maximises the joint log-likelihood function:

$$\mathcal{L}(y; \lambda, \mu) = \sum_{i} [y_i \ln \bar{y}_i(\lambda, \mu) - \bar{y}_i(\lambda, \mu)], \qquad (2.4)$$

where in equation 2.4  $y_i$  denotes the measured sinogram data, and  $\bar{y}_i(\lambda, \mu)$  represents the predicted data, computed as a function of both the activity  $\lambda$  and the attenuation  $\mu$ .

Typically, the predicted data is calculated via:

$$\bar{y}_i(\lambda, \mu) = n_i \exp\left(-\sum_k l_{ik} \mu_k\right) \sum_j c_{ij} \lambda_j$$
 (2.5)

where in equation 2.5  $n_i$  is the detector-pair sensitivity,  $l_{ij}$  is the intersection length of LOR i with voxel j,  $\mu_j$  is the linear attenuation coefficient map in voxel j,  $c_{ij}$  is the probability that an event in voxel j is detected along LOR i, and  $\lambda_j$  is the activity in voxel j.

In practice, MLAA alternates between:

- Updating  $\lambda$  using fixed  $\mu$  via an expectation-maximisation step.
- Updating  $\mu$  by holding  $\lambda$  constant, treating the problem similarly to transmission tomography.

A key advantage of MLAA is its potential to recover both activity and attenuation information from emission data alone, which is particularly attractive in PET/MR where CT-based  $\mu$ -Maps are unavailable. However, since  $\mu$  is not explicitly reconstructed, MLACF is unsuitable in workflows that require  $\mu$ -Maps for further processing (for instance, in MR-based attenuation correction pipelines or hybrid imaging involving multiple modalities).

#### 2.4.3.2 MLACF

MLACF offers an alternative strategy by reconstructing ACFs along each LOR rather than a full  $\mu$ -Map. In this formulation, the attenuation component of the system matrix  $a_{ij}$  is treated as a set of parameters to be estimated directly. The benefit of this approach is that it avoids reconstructing a full image of  $\mu$ , thereby significantly reducing computational load while maintaining good image quality (Nuyts et al. 2012), (Rezaei et al. 2012).

The forward model in this case becomes:

$$\bar{y}_i(\lambda, a) = n_i a_i \sum_j c_{ij} \lambda_j \tag{2.6}$$

where in equation 2.6  $n_i$  is the detector pair i sensitivity,  $a_i$  is the attenuation correction factor for LOR i,  $c_{ij}$  is the probability that an activity in voxel j is detected along LOR i, and  $\lambda_j$  is the activity in voxel j.

The update equations alternate between estimating the activity  $\lambda$  and the ACFs, typically using maximum-likelihood estimation steps adapted to this decoupled formulation.

Although this formulation provides significant computational savings, a key limitation is that since  $\mu$  is not explicitly reconstructed, the method is not suitable when  $\mu$ -Maps are needed for other purposes (for instance, PET/MR co-registration or quantitative multimodal analysis).

Additionally, although MLACF is less sensitive to cross-talk artefacts, it still inherits some of the ill-posedness of MLAA in Non-TOF scenarios.

## 2.4.3.3 Incorporation of TOF

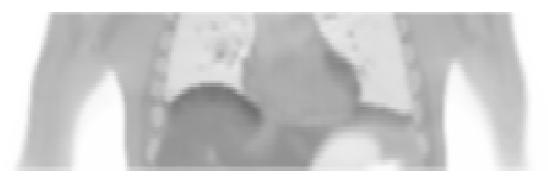
TOF information can be incorporated into both MLAA and MLACF frameworks. TOF narrows the spatial uncertainty along each LOR, thereby improving identifiability and reducing cross-talk. This results in more stable and accurate estimation of both activity and attenuation or ACFs.

# 2.5 Respiratory Motion in PET

This section of the thesis introduces the problem of respiratory motion in PET/CT. First it addresses how respiratory motion presents, specifically in reconstructed PET volumes, and what this can mean for clinical diagnosis. Next the issue of respiratory motion in



**Figure 2.16:** Example of a simulated CINE-CT AC Non-TOF OSEM reconstruction, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.



**Figure 2.17:** Example of a simulated single static  $\mu$ -Map AC Non-TOF OSEM reconstruction, with motion, with no noise, randoms or scatters, of the thorax with a spherical lesion in the lungs. Coronal view.

combined PET/CT is highlighted and the challenges associated with the misalignment of the data between modalities is addressed. Including how this could negatively impact the AC reconstructed PET volume (issues introduced by the PET/CT workflow in the clinic are also presented).

The second subsection expands upon the challenges introduced by respiratory motion in the combined PET/CT workflow and lists some methods from the clinic and the literature which have been developed to combat this.

## 2.5.1 Respiratory Motion Artefacts

A static single bed position acquisition on a conventional PET scanner takes, on average, approximately 120 s. This means that, because the patient is in respiratory motion throughout the acquisition, then the result of the scan will contain data from different respiratory states. During the different respiratory states the position and volume of the lungs, diaphragm, and any lesion will change. If the data is reconstructed without account-

ing for this, there will be the presence of blurring artefacts (especially prevalent around the anatomy that is moving the most, such as the diaphragm). An example of a PET reconstruction with motion artefacts can be seen in figure 2.16, notice the blurring above the diaphragm on the right side of the figure. Artefacts originating from the moving anatomy pose the largest challenge in imaging of the thorax (Beyer et al. 2003), (Sureshbabu and Mawlawi 2005).

Artifacts caused by respiratory motion lead to clinical issues with cancer staging and follow-up. This is because the size of the lesion is often overestimated and the activity underestimated (the activity in the lesion is spread over more voxels). Thus, respiratory motion has the capacity to cause lesions to potentially be missed due to reduced detectability (Erdi et al. 2004).

If AC is used then the position of the  $\mu$ -Map in relation to the PET data also poses an issue. Where the  $\mu$ -Map does not match the position of the anatomy then it will cause either under- or over-correction of the attenuation. This can cause a type of artefact often referred to as a 'banana' artefact due to the shape of the shadow that it causes to appear above the diaphragm (Bai and Brady 2011). An example of this can be seen in figure 2.17. Notice the black arc shaped artefact over the diaphragm and on the heart. The mismatch of the  $\mu$ -Map and PET data does not just cause this artefact but it can also change the expectation and thus the quantification of the reconstructed image. To combat intra- $\mu$ -Map motion the patient will often be asked to hold their breath, if they can, as the CT acquisition can last for only between 2.0 s and 3.0 s (as part of a combined PET/CT acquisition) (Nyflot et al. 2015). An issue with this is, that often the breath hold CT will be taken at full inspiration. If this  $\mu$ -Map is then used to correct for attenuation in data that is in a respiratory phase other than full inspiration then a lot of this anatomy will have been moved into the FOV (it is not present in the  $\mu$ -Map). Furthermore, when a patient is asked to hold their breath they often inhale deeper than they otherwise would, meaning that this  $\mu$ -Map can be in an unrealistic respiratory state that will not be represented in the PET data taken during free breathing.

# 2.5.2 Respiratory Motion Challenges in Combined PET/CT Imaging

To overcome the issues mentioned in section 2.5.1, specifically related to the mismatch between  $\mu$ -Map and PET data a number of solutions have been proposed.

- Firstly, a method which acquires PET data over a prolonged acquisition and discards any data where the patient is in a respiratory position other than the one that corresponds most closely to the breath hold CT  $\mu$ -Map. For this all data which is not at full inspiration would be removed (Liu et al. 2010), (Grootjans et al. 2014). The breath hold CT  $\mu$ -Map could then be deformed to this data (Nehmeh et al. 2007). An advantage of this approach is that it would not only correct for the misalignment of  $\mu$ -Map and PET data but it could also eradicate most blurring associated with averaging over respiratory phases. A disadvantage is that, either there would be substantially more noise in the reconstructed data (if the acquisition was the same length as a standard one) or the acquisition could take significantly longer (to acquire an equivalent number of counts), as so much data is being removed (Nehmeh and Erdi 2008). Additionally, dynamic scans would not be possible with this correction method. This is because the radiotracer kinetics could be shorter than one respiratory cycle and thus would be lost when those parts of the acquisition are removed.
- Secondly, a variation of the previous method has been proposed, where the PET data is separated into individual images representing the separate respiratory phases and warped to a common respiratory phase (seen in section 2.6.1). The breath hold CT μ-Map could then be warped to the same common respiratory phase. This method provides the advantage over the first in that it uses all of the data from the PET acquisition and provides a more robust reconstruction (Bai and Brady 2009). A disadvantage is, that the reconstruction of each respiratory phase is likely to contain more noise than if all of the PET data was reconstructed simultaneously. This is because the iterative reconstruction algorithm (seen in section 2.4.2) is non-linear (analytical reconstructions are linear but are not used clinically any more) and summing reconstructed volumes is not equivalent to summing projection data,

reconstructing and then summing again. In addition, the higher levels of noise in the reconstructed data can pose a problem when attempting to deform the  $\mu$ -Map to them.

• Finally, a method where the reconstruction and motion parameters can be estimated simultaneously, directly from the PET data, for one breath hold CT  $\mu$ -Map has been recently proposed (Jacobson and Fessler 2003), (Rezaei et al. 2018), (Bousse et al. 2016b). Here, the PET data is split into the respiratory phases, as above. Then the method iterates between a reconstruction step (seen in section 2.4.2) and a motion parameter estimation step (seen in section 2.6.1) where the same parameters are used to deform both the PET data and the  $\mu$ -Map for each respiratory position. Thus the  $\mu$ -Map does not have to correspond to any one respiratory phase as each set of PET data will be reconstructed at the position of the  $\mu$ -Map. This method works especially well when TOF data is available (Bousse et al. 2016a). A disadvantage of this method is that it takes more computation than the above methods and that it has not been as extensively evaluated.

#### 2.6 Motion Correction for PET

This section of the thesis discusses motion correction (specifically when applied to PET). The first subsection introduces the concept of registration. This section initially highlights how motion correction can have some similarity to PET reconstruction (in the sense that they are both optimisation problems), before moving on to discussing the classification of motion types and how they can be corrected. Types of deformation are then introduced, including rigid deformation (also affine transformation) and non-rigid deformation. Classic approaches to motion correction are next highlighted. Including parametric and non-parametric registration and regularisation terms used with these methods.

The second subsection moves on to introduce the concept of respiratory gating, including both amplitude and phase gating. The third subsection explains how the SSs, which are used in respiratory gating, are acquired. Firstly, methods incorporating external devices to extract SSs are highlighted (including using the Real Time Position Management (RPM)) before DD methods are introduced. Principal Component Analysis (PCA) is explained generally (as a dimensionality reduction technique) and then its application

#### 

**Figure 2.18:** A graphical representation of the concepts of push and pull interpolation. On the left of the figure, under section **A**, push interpolation can be seen. Here, arrows showing the 'pushing' of the value from the source to the target image can be seen. The blue highlighted pixels in the target image show where holes would be. On the right of the figure, under section **B**, pull interpolation can be seen. Here, arrows showing the location of the coordinates of the pixels from the target image deformed into the source image space can be seen (Akintonde 2021).

specifically for acquiring SSs directly from PET acquisition data is provided.

The fourth subsection gives examples of how to combine registration and respiratory gating. This subsection specifically contrasts pair- and group-wise registration, before the next subsection introduces an extension to these methods by incorporating motion modelling. Here, both iterative and simultaneously with registration MMs are explained. Types of MM including simple motion modelling (for instance, a linear regression of DVFs and SS) as well as more complex motion modelling approaches are highlighted. Furthermore, different types of RCM and their formation, advantages, and disadvantages are addressed briefly.

The fifth subsection is related to the actual application of motion correction. For instance, how and where it is applied. This includes highlighting the benefits of both a post reconstruction and iterative with reconstruction schema.

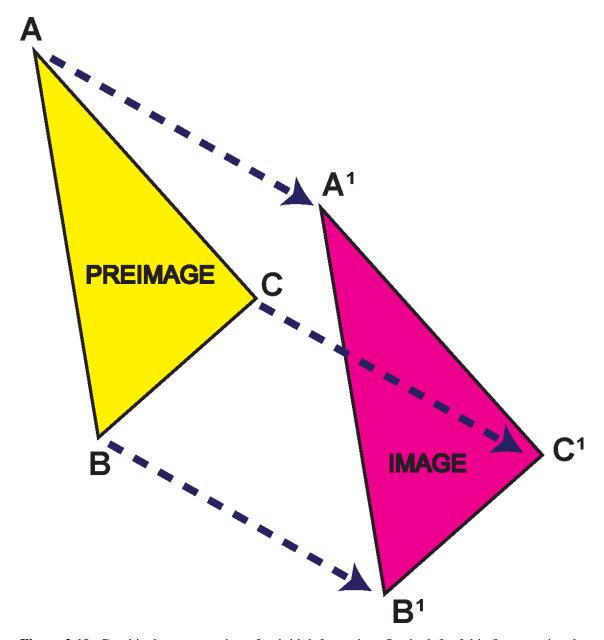
The final subsection provides a summary highlight of some recent motion correction literature in order to, in one place, demonstrait the research landscape that the rest of this thesis fits into.

## 2.6.1 Image Registration

There are different classes of deformation. These include rigid deformations and non-rigid deformations, which will be discussed in the following sections, in section 2.6.1.1 and section 2.6.1.2 respectively. The images used for registration do not necessarily need to be from the same modality and as such CT data can be registered to PET data or vice versa. A common use for registration in medical imaging is to aid in the process of motion correction.

Registration is an optimisation problem which attempts to deform one image (often called the floating image) to another image (often called the reference image). The process through which one image is deformed to another image is by the use of a DVF transformation (or a transformation parametrised by a DVF) and a re-sampler. The resampler takes the floating image and the DVF (a transformation can be parametrised by a DVF) and uses interpolation to resample the floating image with regard to the DVF. There are two types of interpolation which can be used. The first is push interpolation. Here, the values from the source image are 'pushed' into the target image space and then new values for each regular pixel or voxel are interpolated. However, if no values are 'pushed' into the space of a given pixel or voxel then it is undefined what value these spaces should be. As such, with push interpolation there can be holes in the target image. The second is pull interpolation. Here, the coordinates of the pixels or voxels in the target image are transformed to the source image space, new values are interpolated from the source image and they are 'pulled' back into the target image space. This method has the advantage over push interpolation because it is guaranteed that there are no holes in the target image. A graphical representation of the above concepts can be seen in figure 2.18.

As discussed in section 2.3.2 an optimisation requires an objective function. In the case of registration the most common objective functions are MSE (or Sum of Squared Differences (SSD)) and NMI. MSE is discussed in section 2.3.2.1 but simply assumes that once the floating image has been deformed to the reference image then the images should be close to identical. Thus MSE is best used when the only difference between the two images is from, for instance, motion and not from a change in modality. NMI is less reliant on the specific intensity values of an image and instead look for relationships between

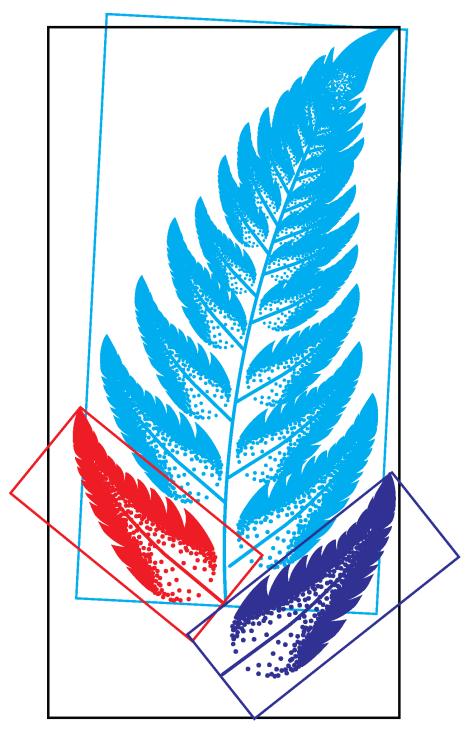


**Figure 2.19:** Graphical representation of a rigid deformation. On the left of this figure a triangle with vertices A, B and C can be seen. This triangle undergoes a rigid deformation (a translation down and to the right) to a triangle, on the right of this figure, with vertices A', B' and C'.

intensities. Thus they are more suitable to registering between different modalities than MSE (Hill et al. 2001), (Oliveira and Tavares 2014).

## 2.6.1.1 Rigid Deformation

A rigid deformation could be a rotation or a translation of the entire contents of an image where the same rotation or translation is applied at every point. A rigid deformation is one where the euclidean distance between every pair of points in the image is consistent before



**Figure 2.20:** Graphical representation of an affine transformation. In the centre of this figure a cyan leaf can be seen which undergoes an affine transformation (a translation down and to the left, a rotation anticlockwise and a scale down) to a red leaf, on the left of this figure. The cyan leaf also undergoes an affine transformation (a translation down and to the right, a rotation clockwise, a scale down and a mirroring) to a blue leaf, on the right of this figure.

and after the deformation is applied, this can be seen in figure 2.19. Rigid deformations are a subset of a type of deformation called an affine transformation. A 3D rigid deformation has six degrees of freedom, being rotation and translation in every axis, whereas a 3D affine transformation has 12 degrees of freedom, rotation, translation, scaling and sheering in every axis. An affine transformation does not guarantee that the euclidean distance between pairs of points are maintained but it does guarantee that sets of parallel lines remain parallel after they are transformed. This can be seen in figure 2.20.

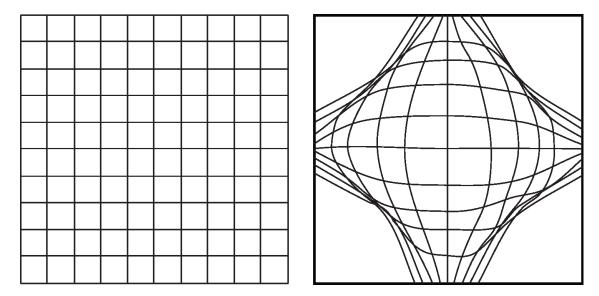
Rigid deformations are often seen (or used as part of image registration) in medical imaging where the anatomy which is being registered is not expected to undergo individual internal motion. For instance, rigid deformations are often used in the registration of patient head motion (Hill et al. 2001). Affine transformations are not often used in medical imaging as anatomy does not usually deform in ways that affine transformations can capture but rigid deformations cannot. Although, an affine transformation can be used as an initial estimate for image registration when subsequently fitting a more complex non-rigid deformation.

The output from a rigid deformation or affine transformation is usually the values from the approximately six, or 12, value transformation matrix respectively. As such they do not take up much computational memory or storage. However, there is no reason that a rigid deformation or affine transformation could not be parametrised differently, for instance as a DVF, it is just not necessary. However, a non-rigid deformation could not be parametrised completely using a transformation matrix.

#### 2.6.1.2 Non-Rigid Deformation

A non-rigid deformation is one where the euclidean distance between pairs of points is not maintained, this can be seen in figure 2.21. Notice that the euclidean distance between where the lines of the grids intersect changes between the grid on the left of the figure and the grid on the right of the figure. Non-rigid deformations (as an motion artifact and as a class of image registration) are commonly seen in medical imaging in respiratory motion. This is due to the diaphragm and lungs being displaced and displacing parts of the anatomy by different amounts.

A non-rigid deformation is often parametrised by a DVF. A DVF could be represen-



**Figure 2.21:** Graphical representation of a non-rigid deformation. On the left a grid can be seen, this grid undergoes a non-rigid deformation to the grid on the right.

ted by deformation vectors for all voxels in a volume. A deformation vector could specify the position of a voxel in reference to a source position. In the same way that there are push and pull interpolation re-samplers, a DVF can also be defined in both directions.

However, for reasons of computational efficiency (and usually to regularise the registration optimisation) a DVF can be parametrised using fewer parameters. One parametrisation could be to use control points on a Control Point Grid (CPG). Here, there are usually fewer control points than pixels or voxels in the original DVF, often these control points are spaced evenly on a regular grid called the CPG. To get a DVF back from this CPG, values are interpolated using, for instance, linear or B-spline interpolation to find the vector to be applied at each voxel (Bardinet et al. 1996), (Rueckert 1999), (Mattes et al. 2003), (Jacobson and Fessler 2003). There are also registration methods which forego further parametrisation and instead directly fit the DVF. These registration methods often use another form of regularisation, for instance a regularisation which penalises rapid local changes in the DVF. An example of such a registration method would be Daemon registration (also known as non-parametric registration) (Vercauteren et al. 2009). For completeness, it is possible to treat a parametrised registration method as a non-parametric registration method by setting the number of control points equal to the number of voxels in the DVF.

Regularisation terms are often employed for non-rigid deformation registration. This

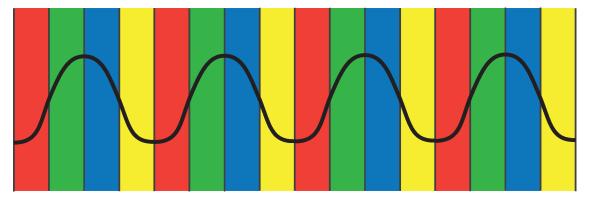
is because otherwise, with a high enough resolution CPG, it is possible for the optimisation to fit the noise present in the data rather than fitting the motion, as discussed in section 2.3.2. One common form of regularisation is a smoothness penalty, for instance bending or linear energy. Here, simply, the term is calculated as the integral of the square of the second derivative of the DVF. This term is multiplied by some value  $\varepsilon$  (representing the weighting of this penalty term) and then the scaled term is summed to the current value of the objective function. This regularisation term attempts to enforce that adjacent control points should not rapidly change, with regards to one another, as this type of motion is unlikely physically (Duchon 1977). Or in other words this regularisation term attempts to encourage a smooth DVF. In the case of non-parametric registration, a Gaussian smoothing of the input data (or DVF) is sometimes used for regularisation (Vercauteren et al. 2009).

A further parametrisation of a DVF would be a velocity field parametrisation. This can be applied on top of an already parametrised method or a non-parametric method. Here, the aim is to generate a diffeomorphic DVF, or in other words a DVF free from folding. Folding would be where the deformation vectors of a DVF cross. Not only is this non-anatomical but it also prevents the DVF from being inverted. In simple terms, for a velocity field parametrised DVF it is expected that the DVF will be integrated over multiple times. The number of times the DVF is integrated over is determined by the minimum number of integration steps necessary to ensure that none of the deformation vectors in the DVF cross. If during optimisation two or more of the deformation vectors cross then the optimisation can revert an iteration and increase the number of integration steps (Modat et al. 2012).

Another method to help to ensure that is registration is diffeomorphic would be to treat a registration as a symmetric registration, rather than a asymmetric registration. A asymmetric registration is as described above, the floating image is registered to the reference image. However, in a symmetric registration the floating image is registered to the reference image and the reference image is registered to the floating image. This helps to ensure that the DVFs generated are invertible (Modat et al. 2014).



**Figure 2.22:** Graphical representation of amplitude gating. Here the pseudo sinusoidal signal represents the SS and the horizontal lines, colour coded differently, represent the amplitude gates that the data upon the SS would be binned into.



**Figure 2.23:** Graphical representation of phase gating. Here the pseudo sinusoidal signal represents the SS and the vertical lines, colour coded differently, represent the phase gates that the data upon the SS would be binned into.

## 2.6.2 Respiratory Gating

Methods of respiratory gating were briefly addressed on previously in section 2.5.2. Here, the process of specifically how respiratory gating works will be addressed. In order to separate PET acquisition data (so that counts only in a certain window or respiratory phase are combined into specific gates or bins), a SS which reflects the respiratory state of the patient over time must be acquired. This SS can either directly reflect the amplitude of the patient's breathing, or can be a percentage of the phase through which the patient is in the respiratory cycle at any one time (Kitamura et al. 2017). These two types of SSs directly influence the type of gating that will be performed. These two types are as follows.

• Firstly, amplitude gating takes the maximum and minimum value of the SS and splits the values between them into a number of gates. This can be seen in figure 2.22. The gates can be chosen so that they are either equally spaced apart or so that each gate has a similar number of counts binned into them. The acquisition data is gated by taking its relevant SS value and summing the acquisition data into

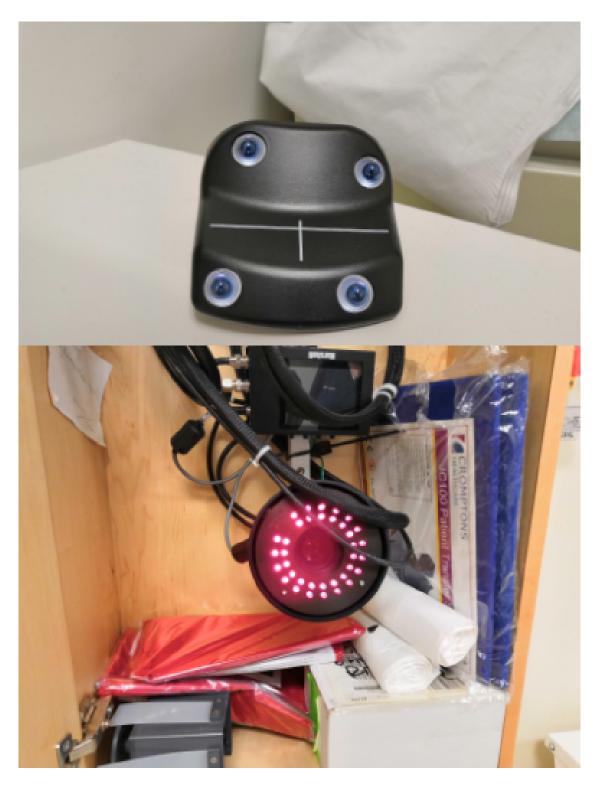
the bin where its SS value falls between the maximum and minimum threshold of the bin.

• Secondly, phase gating works exactly the same as amplitude gating but rather than splitting the data up along the SS it can be conceptualised as splitting the data by the phase of the respiratory cycle. This can be seen in figure 2.23.

Both types of respiratory gating can be augmented by incorporating additional respiratory signals. For instance, amplitude gating can be refined by also sorting events based on the gradient of the respiratory SS. This effectively bins data separately for inspiration and expiration portions of the cycle (Low et al. 2005). Similarly, gating schemes that combine both the phase of the respiratory cycle and the amplitude of the motion have been proposed in the last decade (Manber et al. 2015), (Manber et al. 2016), (Manber et al. 2018). In such phase/amplitude gating, the respiratory signal is partitioned, not only by amplitude, but also by its phase within each cycle. The primary motivation for these hybrid strategies is to avoid mixing different motion states (expanding lungs during inhale vs. contracting lungs during exhale) in the same gate. By separating data in this manner, the breathing hysteresis can be accounted for, rather than blurring the two directions together. A practical limitation of any 2D (or higher dimensionality) gating is that it increases the number of gates, which reduces the counts per gate and can lead to high noise in each gated image (Lamare et al. 2022).

## 2.6.3 Respiratory Signal Detection

As mentioned in section 2.6.2, it is necessary to acquire a signal which represents the position in the respiratory cycle that the patient is in, over the acquisition. This is the case not only for respiratory gating but also for motion modelling (as will be discussed in section 2.6.5). There are two types of methods through which a SSs can be obtained. These are from external mechanical or electrical devices (which directly physically measure the patient) or through DD algorithms (which attempt to extract the SS from the data of the acquisition itself).



**Figure 2.24:** Photograph of the RPM. On the bottom of this figure is the infrared camera and infrared Light Emitting Diodes (LEDs) used to locate and track an infrared reflecting marker. On the top of this figure is the infrared reflecting marker which is placed onto the chest or stomach of the patient in order to track the respiratory amplitude of the patient. Four reflective points are used to track the marker in 3D.

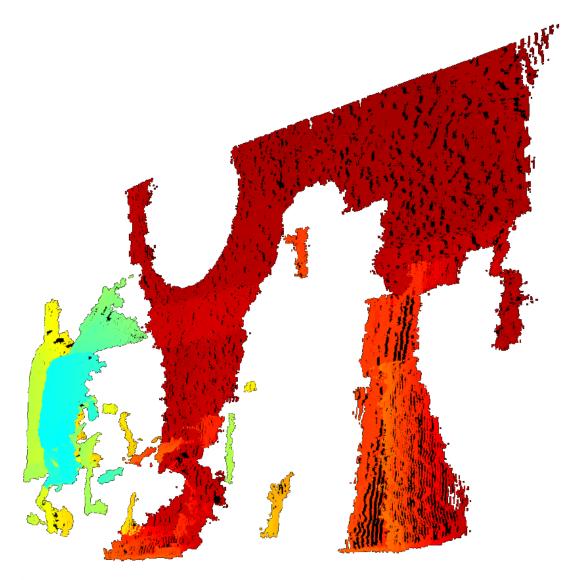


Figure 2.25: Example of a point cloud acquired on a Microsoft Kinect camera.

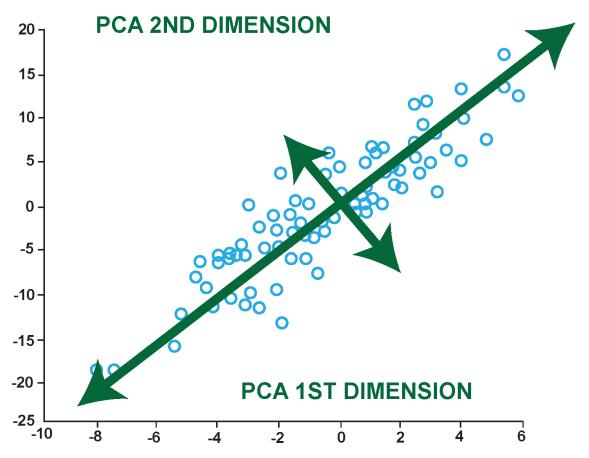
#### 2.6.3.1 External Devices

There are numerous external device methods used to track the patient and acquire a SS. Four example devices are as follows.

• Firstly, the oldest method of SS tracking (presented here) is the use of a spirometer (Voscopoulos et al. 2013). A spirometer is a device with a tube which is inserted into the mouth of the patient, through which they breath. The spirometer measures the volume of air and the velocity with which the patient inhales or exhales (Guivarc'h et al. 2004). Some disadvantages of this method include because spirometers are not designed for highly accurate measurement, over time, they are

susceptible to value drift. This is where the mean position of the respiratory cycle is not consistent between cycles (Hoisak et al. 2004).

- Secondly, a methods borrowed from radiotherapy for SS tracking, is through the use of the Varian RPM (Bettinardi et al. 2013). The RPM was designed to be used in radiotherapy to turn on and off the beam of a linear accelerator (when the beam is estimated to target and not target the appropriate part of the patient). This would work by using the infrared camera of the RPM to track a reflective marker placed on the stomach or chest of the patient (tracking the displacement of the chest wall over time). This can be seen in figure 2.24. The Anzai AZ-733V system attempts to acquire the displacement of the abdomen, similarly to the RPM, but uses a pressure belt wrapped around the patient (Yu et al. 2016). In radiotherapy, as the patient moves, the target of the linear accelerator also moves so the beam is only turned on when the target is within a set range. In PET the RPM has been modified to output a clock tick to the computer acquiring the PET data. This computer will record timing information into the PET data in order to align the RPM SS post acquisition. This method has the advantage over the spirometer in that it is significantly less susceptible to drift. However, the use of the RPM increases scan time (and often has problems with failing) and as such experiences resistance from radiographers.
- Finally, depth sensing, optical, or laser based cameras, such as the Microsoft Kinect (Silverstein and Snyder 2018), (Xia and Siochi 2012). These depth sensing cameras (using the TOF of lasers, for instance, to determine depth) can be used to acquire a point cloud of the patient at each time point. An example of a point cloud acquired on a Microsoft Kinect camera can be seen in figure 2.25. A point cloud is a collection of coordinates measured as points on the surface of the object being scanned at some displacement. A SS can be acquired by finding the difference in these point clouds over time (Miranda et al. 2017). Advantages of this solution include that it doesn't make use of a reflective marker placed on the patient, like the RPM, and as such shouldn't increase scan time. Additionally, an optical or laser based camera can track motion over a larger FOV than the RPM. For instance, an optical or laser based camera could conceivably simultaneously track both respirat-



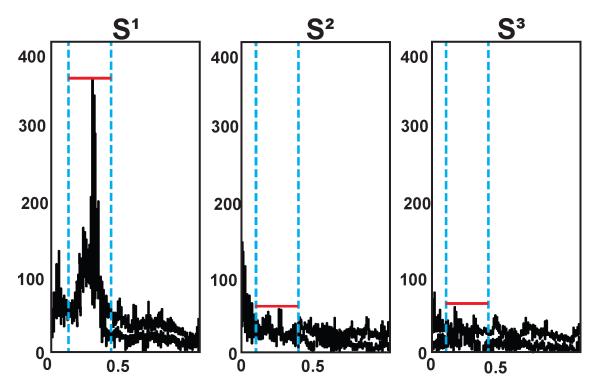
**Figure 2.26:** Graphical representation of PCA applied to 2D data (blue circles). PCA finds the 2 eigenvectors indicated in green

ory and head motion, producing directly a DVF or a separate SSs. A disadvantage is that it is much more difficult to spatially and temporally align the acquisition of both a PET scanner and a stand alone optical or laser based camera without direct manufacturer support. Furthermore, these methods requiring a constant line of sight (Noonan et al. 2012), (Noonan et al. 2015), (Whitehead 2018).

Further limitations of the methods presented above include, most external device methods can only track the surface deformations of the patient, as opposed to the internal ones and such methods require the use of additional equipment, a change to clinical practise, and must be acquired alongside any other data, not retrospectively.

#### 2.6.3.2 Data Driven

There are numerous DD methods used to extract a SS directly from static PET data (Kesner et al. 2014). Some methods require that the PET data be reconstructed first and then markers, possibly inserted into the patient, are tracked over time. Methods requiring



**Figure 2.27:** Graphical representation of the frequency spectrum for signals correspond to three Principal Components (PCs), together with a frequency window between 0.1 Hz and 0.4 Hz. In this example, the first signal has the highest peak within this window.

reconstruction are mostly inferior to methods which work on PET acquisition data. This is because requiring that reconstruction be performed takes significant time and the SS, itself, is usually required for most motion correction methods so the initial reconstructions will be poor. DD methods include, but are not limited to, those which attempt to spatially track aspects of the acquisition data, be that reconstructed or not, or ones which use methods such as dimensionality reduction (Lamare et al. 2022).

Some DD solutions make use of aspects of the image acquisition itself, by reconstructing short time frame images and tracking aspects of them over time. Such as, an external or inserted radioactive fiducial marker (Büther et al. 2013), (Zimmermann et al. 2003), a tumour (Bundschuh et al. 2007), or a combination of patterns from different voxels (Kesner et al. 2009). Some methods make use of MR information (tracking the position of the diaphragm using a pencil shaped navigator) (Taylor et al. 1997), (Fürst et al. 2015). A disadvantage of image space based methods is computation time and potential poor quality due to low count statistics. Obviously methods which require the insertion of objects into a patient have the unnecessary side effect of causing harm to the

patient. Furthermore, to make use of MR information requires a combined PET/MR.

Alternatively, aspects of the data in sinogram space can be individually tracked directly from the list mode data, such as the Centroid of Distribution (COD) or Centre of Mass (COM) (Klein et al. 2001), (Bruyant et al. 2002), (Ren et al. 2017), (Feng et al. 2018). A potential disadvantage, is that these methods require structures with high contrast in sinogram space. A final class of methods uses short time frame sinograms (often at low spatial resolution) and detects motion patterns in the whole sinogram. Such methods rely on the fact that, in static PET, the main cause of (non-stochastic) change in the data is motion. The main sinogram-based methods are the Spectral Analysis Method (SAM) method (Schleyer et al. 2009), (Schleyer et al. 2011), (Schleyer et al. 2018), the Sinogram Region Fluctuation (SRF) method (Kesner and Kuntner 2010) and a method based on PCA (Thielemans et al. 2011), (Bertolli 2018), briefly described below.

- SAM identifies regions in sinograms which are likely to be experiencing respiratory motion. This is achieved by analysing the frequency spectrum of the result of applying a Fast Fourier Transform (FFT) to each bin in the sinogram. A bin which is experiencing respiratory motion will have a peak in the frequency spectrum at the frequency of the respiratory motion. Through this, areas in the sinogram which are experiencing respiratory motion are determined and the total number of counts in these regions, over time, is used to estimate a SS (Schleyer et al. 2009), (Schleyer et al. 2011), (Schleyer et al. 2018).
- SRF proposes to recursively combine signals from sinogram bin Time Activity Curves (TACs). This is performed using a score based on the ratio between respiratory and non-respiratory content (with a positive and negative sign) in order to maximise its Standard Deviation (STD) (Kesner and Kuntner 2010). However, a disadvantage of the use of STD, as the objective, would be that there are many ways to increase STD, which are not acquiring a better respiratory trace. For instance, noise may increase the STD of a signal.
- PCA works similarly to Singular Value Decomposition (SVD), in fact most implementations of PCA use SVD. The goal of this method is to find linear transforms of the data, such that it is projected to a space along which its axis point in the

direction of greatest variance (and then second greatest variance and so on) (Pearson 1901). PCA produces, for the data, a series of eigenvectors and weights. Here, the eigenvectors (called PCs) are the orthogonal vectors of descending variance through the data (usually normalised) and the weights are the magnitude of the contribution of the components to the data. Thus, the first eigenvector from PCA will represent the vector of greatest variance through the data. For SS extraction PCA is applied across a time series of sinograms. The weighting of each PC for each time point would be the signal. Generally multiple PCs are extracted, and the one which contains the most respiratory information (determined using FFT) is selected, as will be discussed further below (Thielemans et al. 2011), (Bertolli 2018).

Focusing on PCA, when applied specifically to PET acquisition data (following (Thielemans et al. 2011)) PCA is often used on either sinograms (or unlisted listmode data). The input sinograms are usually spatially downsampled. This is beneficial for a number of reasons including that the noise present in the full data would obscure the motion. Additionally, the large size of the PET sinograms introduce issues when it comes to storing the number needed in memory (and the computational expense necessary) to apply PCA. Furthermore, the non-downsampled sinograms contain more than enough information than is required for PCA to be able to extract the relevant variation. Thus if all the data was used, then time would be wasted processing all of this data. Generally, additional filtering (or smoothing) can be applied and in most cases will, in fact, improve results, see section 5.3. Usually when used to extract respiratory variation the sampling rate of the PET sinograms is chosen as 0.5 s. This sampling rate is chosen so as to attempt to mitigate cardiac motion (by averaging most of it in each frame) while still allowing for respiratory motion between frames (Bertolli 2018).

The PC which contains the variation present in the data caused by respiratory motion must be identified, as discussed briefly above. One method to do this is to compute a 'score' for every PC and select the one with the highest score. Here, we will describe the original method developed in (Thielemans et al. 2011), but also see additional methods described in Chapter 5.

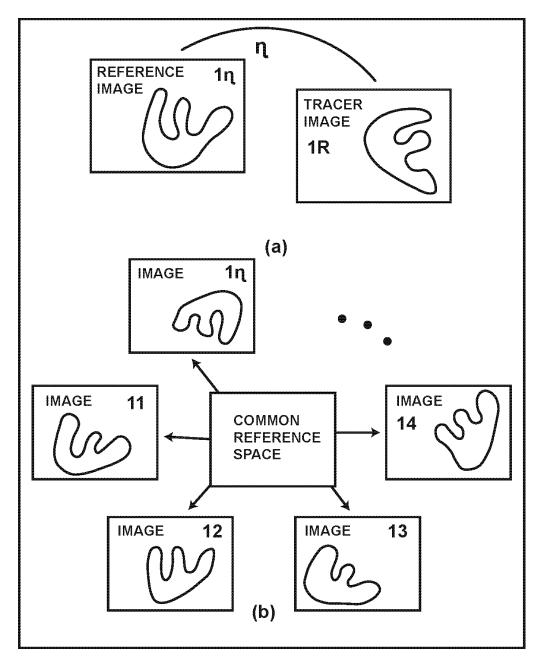
First, the frequency spectrum of the weight of each PC is computed. Then a fre-

quency window is determined, this is usually between 0.1 Hz and 0.4 Hz (so as to coincide with the approximate frequency of respiration). The max value (or peak) in the frequency spectrum is then found for each PC in this window. The PC which has the greatest peak in the window is determined to be the one best representing variance caused by respiratory motion. This can be seen in figure 2.27 (Thielemans et al. 2011).

To-date, evaluations of these methods have been almost exclusively performed on static PET data, mostly using [<sup>18</sup>F]-FDG. They include comparisons with external devices (such as the RPM), MR navigator based SSs (Manber et al. 2015), as well as image quality (Büther et al. 2020), (Walker et al. 2019). Preliminary investigations indicated that many sinogram-based methods all perform similarly. Specifically Data Driven Principal Component Analysis Surrogate Signal Extraction (DD-PCA) showed a correlation of 0.89 (over nine patients) when compared to the MR navigator based SSs (Thielemans et al. 2013), (Manber et al. 2015), (Walker et al. 2018), (Walker et al. 2019), (Büther et al. 2020), (Sigfridsson et al. 2021).

An advantage of DD SS extraction methods over an external device based SS tracking method is that they allow for retrospective extraction of the SS. DD methods can be applied automatically without the intervention of a clinician, nor radiographer. Thus not affecting acquisition time or inserting additional complex steps into clinical practise. DD methods are also more comfortable for the patient as there is no need to be either strapped into or otherwise have to interact with an additional device. Additionally, DD-PCA does not require any other modality, like the MR navigator does. This is all the while, as discussed above, providing accurate results when compared to the external devices. As such, DD methods are universally appropriate wherever PET acquisition data is available.

However, an issue with DD-PCA, and most DD methods, when applied to dynamic PET is that it is sensitive to dynamic radiotracer kinetics from dynamic PET scans, this is similar to as discussed in section 2.2.1.3. This means that wherever radiotracer kinetics are apparent in the data, they can mask the respiratory SS derived via DD-PCA. This will be investigated further in chapter 5.



**Figure 2.28:** Graphical representation of the difference in methodology of pair- and group-wise registration. At the top of this figure pair-wise registration is shown. Here, the floating image is directly registered to the reference image. This process will be repeated independently for the number of images which are to be registered. At the bottom of this figure group-wise registration is shown. Here, a common reference to all floating images is chosen (for instance a weighted sum of the target images). This ensures the reference image is located closer on average to all floating images than if a single floating image was chosen. Pair-wise registration is complete after one iteration of registration. Group-wise registration would involve updating the common reference and iteratively registering again for a set number of iterations (as seen in figure 2.29).

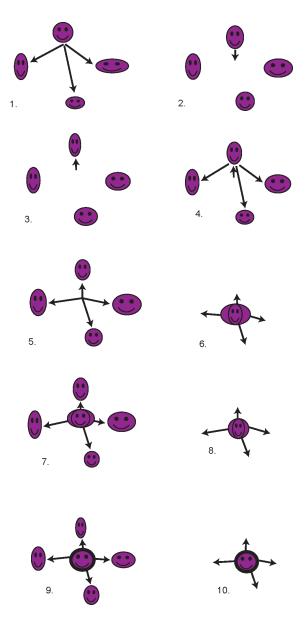


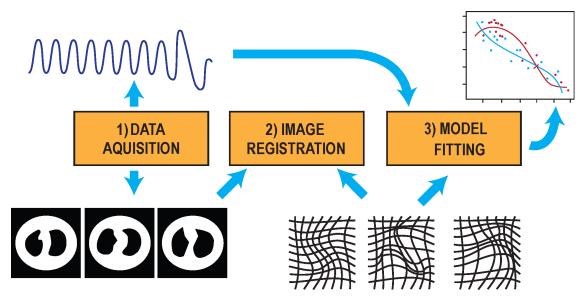
Figure 2.29: Graphical representation of one method of implementing group-wise registration. The sub-images one through four show the initial setup used for this implementation and the sub-images five through 10 show the subsequent iterations of the group-wise algorithm. Sub-image one shows an initial pair-wise registration of one image to the rest of the images in the series. Sub-image two shows how to find the mean position of all images by taking the mean of the deformations from the previous step. Sub-image three shows the inverse of the deformation from sub-image two. Sub-image four shows the deformations from sub-image one being composed with the deformation from sub-image three, thus there is now a deformation from each image to the mean position. Sub-image five shows the resultant deformations from the composition in sub-image four. Sub-image six shows the target images resampled using the deformations from sub-image five to form a reference image at the mean position. Sub-images seven through 10 show incremental improvements to the reference image from subsequent iterations of registration and resampling.

#### 2.6.4 Applying Image Registration

While there are multiple ways to motion correct a series of respiratory gated PET volumes using registration the two most common approaches are reported below:

- First, pair-wise registration. Here, one of the gates is selected and all other gates are registered to it. Usually the gate with the highest number of counts is selected, due to it having the highest SNR. Once registration is complete, each gate is resampled using the corresponding DVF and summed to form the final motion compensated image. An advantage of this algorithm is that it is relatively simple, usually only one pass of registration is required and thus it could be considered relatively quick. However, the accuracy of the motion correction entirely hinges on the quality of the volume selected as the reference, if it is low quality then it would be difficult to accurately register the gates. Additionally, the reference gate may be far from a number of the other gates if the reference is close to max inhalation or exhalation (compounded with a large amount of the anatomy possibly being outside the FOV if the reference is close to max inhalation) again hindering the ease of registration.
- Second, group-wise registration. Here, all gates are summed to form a reference volume. Gates are registered to the reference, resampled and summed once more to form a new reference volume. This process is repeated a set number of iterations, with the quality of the reference, ideally, increasing each time. As an additional step the inverse mean of all DVFs can be composed with each DVF before resampling the gates to ensure that the reference gate is at the mean respiratory position. Composing the inverse mean DVF offers a significant advantage in the sense that it ensures that the reference image is 'centred' with respect to the floating gates. This means that the reference image is as 'close' as it can be, on average, to all floating gates. Therefore the deformation, on average, from each gate is as small as it could possibly be. A graphical representation of group-wise registration can be seen in figure 2.29. A further advantage of this algorithm is, it has the capacity to provide better registration results due to the higher SNR of the reference volume. However, it takes longer to perform due to the multiple full iterations of registration.

A graphical comparison of pair- and group-wise registration can be seen in fig-



**Figure 2.30:** Graphical representation of the process of motion modelling. On the left of the figure the input data of both the image data and the SS can be seen. Then in the centre of the image the image data is taken and registration is applied to it. Then on the right of the image a RCM is fit on the DVFs from the registration step and the input SS. In the top right of the figure a naive example of the 2D regression result can be seen. In this case the regression appears to be polynomial.

ure 2.28.

#### 2.6.5 Motion Modelling

MMs offer an alternative solution to motion correction when compared to direct coregistration of respiratory gated PET. MMs, similarly to the RPM discussed in section 2.6.3.1, is a technique borrowed from radiotherapy. In radiotherapy the MM would be used to help predict the respiratory position of the patient in real time. MMs themselves attempt to relate the DVFs or B-spline parametrised CPGs from, for instance, registration to the SS acquired for the data used to generate the DVFs. In this case a RCM would be generated which is fit on the DVFs and SS, this can be seen in figure 2.30. In its simplest form a post registration MM could be a linear regression of both the DVFs and the SS (Kim et al. 1997). In the case described previously the DVFs are found first before fitting the RCM. However, in more recent research a method to simultaneously fit both the B-spline parametrised CPGs and the RCM has been presented (McClelland et al. 2006), (McClelland et al. 2013), (McClelland et al. 2014), (McClelland et al. 2017). RCMs are discussed in more detail in section 2.6.5.1.

In the most common case, there are four components required to fit a MM. The

data itself which is to be motion corrected, a SS with a value representing the respiratory position at each data point (the SS can be *n*-dimensional), a model (such as B-spline) for the deformations, and the RCM linking the DVFs or B-spline parametrised CPGs and SS. In some cases the SS can be replaced with other features of the acquisition or data.

An advantage of MMs over direct co-registration is that it allows for prediction of data not used to fit the RCM. This advantage is particularly useful for radiotherapy. Here, some data is acquired on the machine used to perform radiotherapy (usually on the patient on the day of therapy). This allows for a RCM to be fit (pre-therapy) on the data acquired with a SS from, for instance, the RPM. Then, as a dose is administered to the patient, from a linear accelerator, SSs values can continue to be acquired and DVFs calculated using the RCM (fit previously). This is done in order to, in real time, change the trajectory of the linear accelerator to apply a dose taking the movement into account. The new SS values need not match any of the SS values used to generate the RCM as the RCM is a continuous model. This advantage could provide the ability in PET to fit a RCM while acquisition is still ongoing. This could be used for motion correction during reconstruction (or indeed for a RCM to be fit on a subset of the data to improve motion correction speed). An additional advantage this could provide would be that the MM could be fit on more coarsely gated data and then applied to less coarsely gated data. If the RCM is One Dimensional (1D) the minimum number of gates required is two. This is because the 1D RCM would be fitting a line. If the RCM is 2D the minimum number of gates requires is three. This is because the 2D RCM would be fitting a plane. If the minimum number of gates is used there may be a higher degree of motion artefacts in the image but there would also be a lot less noise. Thus there is a potential to tune the number of gates used for RCM fitting specifically to find the cross over between where motion artefacts and noise impact the registration more. This MM could then be used to motion correct any number of gates.

An additional advantage of MMs is that compared to direct co-registration it is more robust to noise. This is because direct co-registration will register each individual volume together (regardless of how well that registration fits the underlying motion of the patient). In comparison, a MM would then be fit on top of (or be applied simultaneously with) registration using the SS, this could regularise the effect of outlying data points and noise

in general.

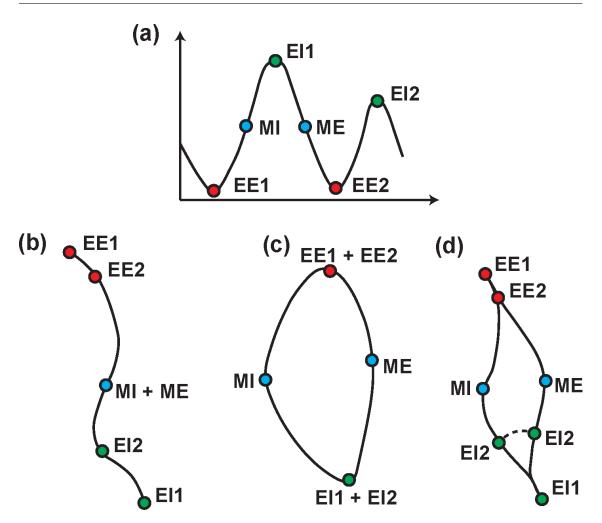
In summary, registration literature contains many examples of spatial regularisers, however there are few temporal regularisers. MMs can act as such a temporal regulariser, as well as allowing for the interpolation of unseen motion correction results. MMs parametrise DVFs in terms of a SS. They could be considered as an addition to a standard spatially regularised motion correction technique, as they impose a degree of temporal or gate-wise regularisation. In addition, they also allow obtaining DVFs for time points/gates, not used to fit the model (as long as a relevant SS exists (McClelland et al. 2013)).

#### 2.6.5.1 Respiratory Correspondence Model

The type of RCM fit depends mostly upon the number and types of SS given to the MM algorithm and the number of models fit. Some of the different types of RCM can be seen in figure 2.31.

Three examples of types of RCM are as follows.

- The most simple form of RCM would be to fit a linear model using a single SS (initially a SS directly representing the amplitude of respiration, over time, could be given). A disadvantage of this approach would be that inhalation and exhalation are treated alike when in reality they differ. For instance the gradient of the SS is not taken into account. This can be seen in figure 2.31, in the bottom right of the figure where mid-inhalation and mid-exhalation fall on the same point.
- To combat the disadvantage of the previous model a RCM could be fit instead on the phase of respiration from the initial displacement SS. A disadvantage of this type of RCM is that it treats every breath as if it is exactly the same regardless of the max displacement of the breath. This can be seen in figure 2.31, in the bottom middle of the figure where the two end exhalation and end inhalation points fall on the same point.
- A final, more complex, RCM is one where a separate model is fit for both the inhalation and exhalation portion of the data. This would mean that combined the RCMs can model both inhalation and exhalation separately, meaning that both inter- and intra-respiratory variation can be maintained. Alternatively, a 2D SS



**Figure 2.31:** Graphical representation of the types of RCM. At the top of the figure is a SS which has been discretised into specific values. On the bottom left of the figure is a linear RCM, fit on a single SS, showing how the points from the discretised SS would be mapped to itself. On the bottom middle of the figure is a RCM fit on SS reflecting the phase of the respiration and on the bottom right of the figure is an RCM where inhalation and exhalation are modelled separately using two RCM (or a 2D SS).

could be given to one RCM where the additional dimension is the gradient of the displacement (or using amplitude and phase simultaniously, as discussed in 2.6.3). This would allow the RCM to fit a plane which can account for the difference in inhalation and exhalation. However, moving from inhalation to exhalation may incur a 'jump' in the output from the RCM (McClelland et al. 2013). Although the RCM fit with a 2D SS can forego the 'jump' associated with fitting two RCMs, it does suffer with an additional disadvantage. The additional disadvantage of this method is that if gating of the input data is employed then each gate would have far less data within it when compared to where a single SS and single RCM is used.

For more detail on this final type of 2D RCM see chapter 3 and chapter 4.

Furthermore, although it has not been explored here, the RCM does not have to be linear. As an example, if a post-registration RCM is fit as a linear regression, then there is no reason that a polynomial regression or a type of non-linear regression could not be used. As an extreme example, a NN could be used as a RCM.

The reference position of the SS with regards to the RCM can be defined in two ways.

- One, the reference position is defined as the mean respiratory position. This is similar to how the reference position is defined for group-wise registration, as discussed in section 2.6.4. A value of zero for the SS would correspond to the respiratory position which is in the 'centre' of all the input data or is at the mean respiratory position. In other words, it is equally 'close' to all input data or time weighted average position.
- Two, if an additional dimension is added to the SS that does not vary, it allows the intercept of the RCM to be fit. Thus the reference position does not need to sit at the mean respiratory position, which may in some applications be useful. For instance, if the reference position was set as the position of the μ-Map.

As an example, fitting a RCM using two SSs (for instance, the displacement and gradient, as mentioned previously) would be as follows. 3D B-spline can be used to model spatial deformations with the corresponding warping operation denoted as  $W(\alpha_t)$ , with  $\alpha_t$  a vector with B-spline coefficients at time t. The respiratory SSs s contained two components. Following (McClelland et al. 2013) a direct correspondence MM could be used where the B-spline coefficients at time t are expressed as a linear combination of the two SSs,  $s_{1,t}$  and  $s_{2,t}$ :

$$\forall t \in [[1, n_t]], \alpha_{k,t} = R_{1,k} s_{1,t} + R_{2,k} s_{2,t} + R_{3,k}$$
(2.7)

where in equation 2.7  $\alpha_{k,t}$  is the 3D B-spline coefficient for motion parameter k at time point t, and  $R_{i,k}$  are the model parameters.

#### 2.6.6 Applying Motion Correction

As discussed throughout this chapter (in section 2.5.2 and section 2.6.2) there are a number of methods through which the DVFs (found from direct co-registration or MMs) can be applied to the task of motion correction. These methods mostly fit into three different classes, these are as follows.

- Post-reconstruction motion correction: DVFs are used to register and fuse
  already reconstructed gated images into a single corrected volume Polycarpou et al.
  2012. Motion correction is performed independently after reconstruction. However,
  it suffers from noise amplification due to the lower counts in individual gates and
  does not modify the forward model, limiting its effectiveness.
- Motion-compensated iterative reconstruction: DVFs are incorporated directly into the forward and/or backward projection steps during iterative reconstruction. This allows the entire gated dataset to contribute to a single image reconstruction, improving both spatial resolution and quantitative accuracy Kyme and Fulton 2021. In some implementations, such as those by Bousse et al. Bousse et al. 2016b; Bousse et al. 2016a, DVFs are iteratively updated alongside the reconstruction in an alternating optimisation framework.
- Joint motion-estimation and reconstruction: Rather than alternating between
  reconstruction and motion estimation, this approach jointly optimises both components via a unified objective function Kalantari et al. 2016. While this has shown
  improved performance, it introduces challenges in balancing terms in the cost function and increases computational complexity.

Early strategies for motion correction in PET/CT were after reconstruction methods. In such methods, one acquires the PET data in multiple gates corresponding to different motion states, this was introduced in section 2.6.2. Each gate is reconstructed separately, yielding a series of images. The motion is then estimated by applying non-rigid registration on all floating images to a reference image, this was explained in section 2.6.4, and fusing them into a single motion corrected image by summing (Fayad et al. 2015). This post-reconstruction registration approach can reduce blurring, but it has drawbacks. The

gated images have lower counts (higher noise) and any error in registration (due to noise) will propagate into the final image (Kyme and Fulton 2021).

More advanced methods therefore move the motion compensation inside the reconstruction algorithm itself. In a Motion Compensated Image Reconstruction (MCIR) approach, all gated data contribute directly to forming one final image during reconstruction. Incorporating motion into the iterative reconstruction is conceptually achieved by modifying the forward-projection operation, effectively deforming the projectors according to the DVF. This means that for each gate the system model accounts for where the event originated in the reference image space. In practical terms, the estimated DVFs are used to map each gate to the reference image coordinates during reconstruction. By accounting for the DVF at each forward-projection, the algorithm uses all the PET data and aligns them in the image update step. In summary, the forward-projection operation is dynamically deformed according to the DVF, so that the reconstruction process itself produces a motion corrected image from all the raw data (Kyme and Fulton 2021).

#### 2.6.7 Summary of Motion Correction Methods

Respiratory gated CT has been historically been regarded as the definitive method for accurately estimating respiratory DVFs. This is achieved by using CINE-CT synchronized with the PET acquisition (Pan et al. 2004), (Pan et al. 2006), (Li et al. 2006), (Manjeshwar et al. 2006), (Qiao et al. 2006), (Lamare et al. 2007). However, CINE-CT incurs significantly higher radiation doses, rendering it impractical for clinical PET. Additionally, variations between the PET and CINE-CT acquisitions introduce discrepancies in DVFs, as discussed in 2.2.3.1 (Lamare et al. 2022).

To circumvent the challenges mentioned above, alternative approaches were proposed which suggest estimating DVFs directly from gated, NAC PET reconstructions (Jin et al. 2005), (Dawood et al. 2008b), (Liu et al. 2011), (Fayad et al. 2013), (Chan et al. 2013). These methods simplify acquisition by requiring only gated PET. Regardless, motion estimation from NAC PET suffers from low image contrast (Lamare et al. 2022).

It has been demonstrated that mismatched  $\mu$ -Maps introduce errors into DVFs estimation, subsequently compromising the quality of MCIR (Lu et al. 2018). Furthermore, it has been noted that DVFs derived from NAC PET are unreliable in central anatomical

regions such as the spine. Therefore, optimal motion correction requires matched AC volumes (for each respiratory gate) to ensure accurate motion estimation (Lu et al. 2018), (Lamare et al. 2022).

Methods like MLAA and MLACF, as presented in section 2.2.1.5 and section 2.4.3, can now jointly estimate both the activity volume and the  $\mu$ -Map directly from the PET data alone, with the integration of TOF. This can bypass CT potentially eliminating mismatch issues (Defrise et al. 2012), (Rezaei et al. 2012), (Panin et al. 2013). However, MLAA derived  $\mu$ -Maps (and MLACF derived ACFs) tend to exhibit high noise levels, which is particularly problematic when applied to gated data with fewer counts per gate (Lamare et al. 2022).

In (Rezaei et al. 2016) a method which simultaneously reconstructs the activity volumes and computes DVFs that deform the  $\mu$ -Map to the position of the reconstructed activity volumes is presented. A similar concept was also applied in (Bousse et al. 2016a), (Lamare et al. 2022).

In (Chan et al. 2018) a method is presented to correct for both inter- and intrarespiratory cycle motion. This approach involves two main components: Firstly, continuous DVFs are estimated using the internal-external motion correlation method (Chan et al. 2013). Secondly, these DVFs are then integrated into a MCIR framework, deforming the PET system matrix to align to a reference phase matching the  $\mu$ -Map.

In (Lu et al. 2018) a method is proposed where phase gated PET images are reconstructed using MLACF (producing AC images at each respiratory gate) on which DVFs are obtained by registration of these volumes (Joshi et al. 2011). A linear MM is then fit on the DVFs and a SS acquired using a belt.

### 2.7 Summary of Challenges and Motivation for This Work

In summary, while numerous techniques have been developed to mitigate respiratory motion in PET/CT imaging, major challenges remain unresolved. In the context of static PET, methods such as respiratory gating and image registration, have been employed to reduce motion artefacts. However, these approaches often rely on coarse binning of

data, which limits temporal resolution and allows residual motion to persist. Additionally, they typically use a static  $\mu$ -Map acquired separately from the emission data, which may not correspond to any specific respiratory phase in the PET acquisition. This mismatch introduces significant quantification errors, further degrading image accuracy.

Extending motion correction to dynamic PET is even more challenging. Conventional gating methods are difficult to apply due to the low SNR in early time frames, and because tracer kinetics confound signal extraction. Moreover, many existing approaches depend on external hardware to provide respiratory SSs, which limits their clinical practicality. As a result, motion correction is often poorly implemented or entirely absent in dynamic imaging protocols.

To address these limitations, this thesis proposes a motion correction framework initially developed for static PET, combining respiratory gating with motion modelling to estimate DVF across the respiratory cycle. These DVF are used to correct both the emission data and the misaligned  $\mu$ -Maps, thereby improving image quality and quantification.

Building on this foundation, the framework is then extended to the more complex setting of dynamic PET, where conventional gating is often impractical due to low signal levels and the confounding effects of tracer kinetics. To enable motion correction in this context, a novel DD approach is introduced for extracting respiratory SSs directly from the PET data itself. This method removes the need for external monitoring devices and enables robust motion modelling even in challenging dynamic acquisitions. Overall, this thesis presents a clinically feasible solution for respiratory motion correction across both static and dynamic PET/CT workflows.

#### **Chapter 3**

# Feasibility Studies for Motion Correction Utilising Motion Models on Simple Data

#### 3.1 Overview

This chapter investigates the feasibility of performing respiratory motion correction using simulated PET data, with a particular focus on estimating deformation fields directly from NAC images. Simplified simulation data is used in this early-stage study to allow precise control over motion and anatomy, enabling systematic evaluation of key factors such as TOF information and reconstruction accuracy. The central question addressed here is whether motion models derived solely from NAC PET, particularly TOF PET, are sufficient to enable effective motion correction. This work serves as a stepping stone for more complex and clinically realistic studies in the following chapter.

The first section of this chapter introduces preliminary work on the feasibility of motion correction when attenuation correction is not applied, both with and without TOF (Whitehead et al. 2019).

The second section of this chapter underlines the work needed to incorporate attenuation correction into the methods from the first section. Here, methods to deform a single CT derived  $\mu$ -Map to position at each gate using a MM are introduced. In addition, the advantage of fitting another MM on the newly AC data is evaluated (Whitehead et al. 2020).

The third section of this chapter introduces a comparison between different motion correction methods, both with and without MMs. This section introduces a modular framework for the comparison of different aspects of the motion correction pipeline, for instance, the use of pair- and group-wise registration. Here, the data used is of lower count rate than in the previous two sections and uses a  $\mu$ -Map at end inhalation (close to a breath hold  $\mu$ -Map), rather than one at the mean respiratory position (Whitehead et al. 2021).

The final section chapter discusses the problems or limitations of the work performed in the previous three sections, leading to the next chapter.

#### 3.2 Introduction

As discussed in section 2.6, many methods have been proposed to correct for respiratory motion in PET. These methods usually involve direct co-registration between a reference volume and a set of volumes in different positions in the respiratory cycle obtained by gating (floating volumes) (Oliveira and Tavares 2014). However, such pair-wise registration is sensitive to noise and intra-gate motion. Another problem specific to PET (and Single-photon Emission Computed Tomography (SPECT)) is that motion affects both the location of the emission and the attenuation of the photons. The impact of attenuation is particularly important as the  $\mu$ -Map is often static. This chapter investigates the feasibility of various motion correcting techniques. For instance, is it necessary to use AC for the reconstruction of the gated images used for motion estimation and will TOF alone provide an adequate NAC reconstruction to forego AC? If AC is required then any misalignment of the  $\mu$ -Map will severely hinder the motion estimation and another approach will have to be taken, for instance MLACF or MLAA could utilise the TOF information to provide higher quality reconstructions. However, these methods are more complex and time consuming.

This chapter will investigate this on XCAT simulations, where the realism of the simulations is gradually increased. The work in this chapter was conducted under the hypothesis that access to a PET scanner with higher TOF resolution than the GE Discovery 710 would be available. This assumption was made to explore the potential benefits of improved TOF performance for motion correction, and to better understand whether

future scanner technologies with enhanced TOF resolution might be necessary to fully realise the advantages of the proposed methods.

A linear model was selected in this feasibility study. While real respiratory motion is often non-linear and more complex, the linear model provides a baseline for evaluating the core feasibility of PET-based MM estimation.

# 3.3 Impact of TOF on Respiratory Motion Model Estimation Using Pre-Gated No Intra-Cycle Motion NAC PET

This section investigates the possibility of estimating MMs for respiratory motion from NAC reconstructed images. In particular this section evaluates the impact of incorporating TOF information on the accuracy of the MMs derived from NAC reconstructed images (Whitehead et al. 2019).

#### 3.3.1 Introduction

To avoid misregistration due to attenuation mismatches in the gated images used for motion estimation, most existing methods rely on pair-wise registration of NAC reconstructions (Bai and Brady 2011), (Kalantari and Wang 2017), (Dawood et al. 2008a), (Dawood et al. 2006). However, for PET/CT of the lung, this is a challenging problem due to the low contrast and high noise of these volumes. The use of MMs as opposed to pair-wise registration could overcome these challenges.

In the absence of TOF, there is no information on the activity position along the LOR. NAC Non-TOF reconstructions therefore have high intensity near the surface and low contrast in the internal part of the body. In TOF, the time information constrains the activity position along the LOR. This changes the nature and extent of the artefacts associated with NAC PET (as well as changing noise properties) (Ter-Pogossian et al. 1981).

The aim of this section is to investigate whether accurate MM can be estimated from NAC images, and if TOF is needed to sufficiently increase the contrast and lower the noise of NAC images to facilitate this.

#### 3.3.2 Methods

**Table 3.1:** This table shows a short summary of the highlights of the methods used in section 3.3. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.3	XCAT:     40.0 mm lesion     1D XCAT respiratory traces     No hysteresis     6 volumes     No intra-gate motion      Simulation:     2D Simulation     120 s simulation     American clinical count rate (True 60 Kilo Counts Per Second (KCPS))     Non-TOF and TOF     375.0 ps TOF resolution     No scatter or random events      Gating:     XCAT SS     6 respiratory bins	NAC OSEM (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm FWHM)	<ul> <li>Jointly estimating registration and MM</li> <li>Non-diffeomorphic</li> <li>1D SS</li> <li>SSD objective function</li> <li>Default parameters (CPG spacing 5 times voxel size, 5e-3 bending energy weight)</li> </ul>

Ground truth PET images of one bed position, with a FOV including the base of the lungs and the diaphragm, are obtained using the XCAT software, section 3.3.2.1. These are then used to simulate PET data, section 3.3.2.2, and TOF and Non-TOF volumes are reconstructed using OSEM, section 3.3.2.3. These volumes are used as input for MM estimation, section 3.3.2.4.

The TOF output is evaluated against the Non-TOF output from the method both visually, using a profile, as well as with SUV analysis, section 3.3.2.5.

#### 3.3.2.1 XCAT Image Generation

XCAT (Segars et al. 2010) was used to generate six volumes over a 5.0 s second breathing cycle. The volumes were generated with one at full expiration (beginning of the cycle) and one at full inspiration (end of the cycle), this corresponds to the time in the respiratory

cycle (which is roughly equivalent to phase gating). Settings used for generation were the default XCAT settings for the extent of Anterior Posterior (AP) and Superior Inferior (SI) motion (default XCAT settings being a single generic breathing trace) of 2.0 cm respiratory motion displacement over 5.0 s. Activity concentrations were derived from a static [<sup>18</sup>F]-FDG patient scan. The FOV included the base of the lungs, diaphragm, and the top of the liver with a 40.0 mm diameter spherical lesion placed in the right lung.

#### 3.3.2.2 PET Data Simulation

PET acquisitions were simulated using Software for Tomographic Image Reconstruction (STIR) (Thielemans et al. 2012) through Synergistic Image Reconstruction Framework (SIRF) (Ovtchinnikov et al. 2020) to forward project the input data to sinograms using the geometry of a GE Discovery 710 (0.21306 cm pixel width and 0.327 cm slice width, a image matrix size of 256 by 256 by 47, and 11 timing positions, 23 segments, 288 views, and 381 tangential coordinates). Attenuation for every LOR was computed based on the mu-maps generated by XCAT. 2D simulations were used to reduce computation time (meaning the sinogram used now had 11 timing positions, one segment, 288 views, and 381 tangential coordinates). Where relevant, a TOF resolution of 375.0 ps (similar to the GE Signa PET/MR) was used (meaning the sinogram used now had 25 timing positions, one segment, 288 views, and 381 tangential coordinates). TOF mashing was incorporated to reduce computation time, resulting in 13 TOF time bins of size 376.5 ps. Attenuation was included in the simulation using the relevant  $\mu$ -Map generated by XCAT. Scatter and randoms were not taken into account in this section. Poisson noise realisations were generated to simulate an acquisition as if it had been gated into six bins over an acquisition of 120 s, with a true count rate of 60 KCPS (so 1,200,000 counts per bin). This emulates a standard single bed position acquisition but ignores any intra-gate motion.

#### 3.3.2.3 Image Reconstruction

Data were reconstructed with SIRF without attenuation correction using OSEM with two full iterations and 24 subsets (Hudson and Larkin 1994), see section 2.4 for more information. Volumes were post-filtered using a 3D Gaussian blurring with a kernel size of 6.4 mm FWHM.

#### 3.3.2.4 Motion Model Estimation

Following on from what is presented in section 2.6.5.1, a generalised framework unifying registration and MM estimation, NiftyRegResp (now named Surrogate Parametrised Respiratory Motion Modelling (SuPReMo)), was used to estimate the RCM, here using SSD as the objective function. The RCM was a linear 1D RCM fit on the SS values taken directly from XCAT. The default parameters of NiftyRegResp (SuPReMo) were used.

#### 3.3.2.5 Evaluation

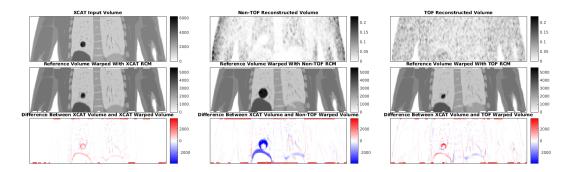
Three RCMs were compared. These three RCMs were calculated from the PET XCAT volumes (gold standard), Non-TOF NAC reconstructions, and TOF NAC reconstructions. To test the accuracy of the RCMs, while avoiding issues with incorporating the motion into the reconstruction, the three models were used to deform the PET volume generated by XCAT at the mean breathing position, to the position at each gate. The mean breathing position XCAT volume was generated by calculating the mean SS value and using this as an input to XCAT. A mean position volume was used because the RCM was fit with this as the reference position for the SS, as discussed in section 2.6.5. These estimated volumes were then compared to the original XCAT input volumes. Difference volumes were obtained by subtracting the original XCAT volume  $f_t$  and deformed volumes  $W(\alpha_t)f_{\rm ref}$  at the same gate. Mean Absolute Percentage Error (MAPE) was computed from these difference images follows:

$$M = \frac{\frac{1}{n} \sum_{n} |e_{n} - g_{n}|}{\frac{1}{n} \sum_{n} g_{n}} \times 100$$
 (3.1)

where in equation 3.1 n is the number of volumes,  $e_n$  are the estimated volumes and  $g_n$  are the ground truth volumes.

In addition, the COM of the lesion was tracked over the six gates. This was achieved by deforming a volume only including the lesion in the reference position, as above, and then computing the COM. The COM along each dimension is calculated using the following equation:

$$C = \sum_{\nu} r_{\nu} e_{n,\nu} / sum_{\nu} e_{n,\nu} \tag{3.2}$$



**Figure 3.1:** Coronal images illustrating results from the estimation of RCMs on NAC images. All volumes correspond to the end inhalation gate. First row from left to right: XCAT PET ground truth image, NAC Non-TOF reconstructed data, and NAC TOF reconstructed data. Second row, RCMs applied to mean position XCAT data with RCMs derived from XCAT PET data (left), NAC Non-TOF (middle), and NAC TOF (right) volumes. Colour map ranges are consistent for all images on this row. The third row from left to right, difference between the estimated volumes from the second row with the XCAT end inhalation volume. Colour map ranges are consistent for all images on this row.

**Table 3.2:** Comparison of the MAPE between the ground truth data and the volumes estimated from the XCAT based RCMs, the volumes estimated from the NAC Non-TOF based RCM, and the volumes estimated from the NAC TOF based RCM.

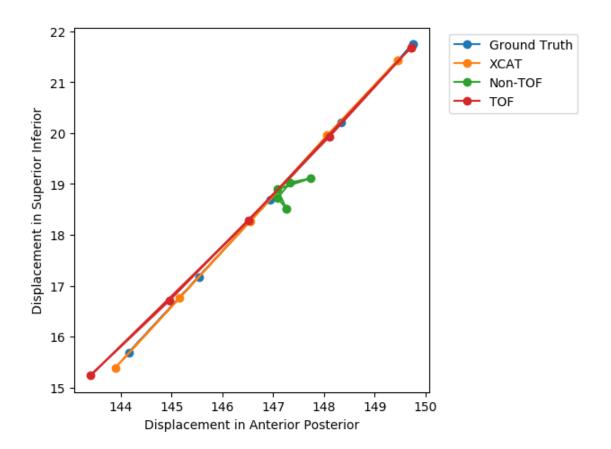
MAPE	XCAT	Non-TOF	TOF
1	1.95	8.35	4.18
2	1.59	1.61	1.84
3	2.06	9.91	5.23
4	1.97	6.15	3.68
5	1.65	4.45	2.52
6	1.95	8.35	4.18
Mean	1.86	6.47	3.60

where in equation 3.2 v indexes all voxels,  $r_v$  the coordinate of the voxel, and  $e_{n,v}$  the value of voxel v in  $e_n$ .

#### 3.3.3 Results

The reconstructed data, estimated and difference volumes can be seen in figure 3.1 and MAPE are in table 3.2. The mean MAPE was found to be lower for the NAC TOF data than for the NAC Non-TOF.

COM results can be seen in figure 3.2. The path of the NAC TOF data follows the ground truth path much closer than the NAC Non-TOF data, and is quite close to the gold standard XCAT-derived motion.



**Figure 3.2:** The path of the COM of the lesion, in voxel indices. Horizontal (respectively vertical) axis corresponds to motion in the AP (respectively SI) direction over the six gates. Different curves denote COM displacement for ground truth data, the estimated data from the XCAT based RCM, the estimated data from the NAC Non-TOF based RCM, and the estimated data from the NAC TOF based RCM.

#### 3.3.4 Discussion

MMs derived from NAC TOF volumes were found to be more robust than when using NAC Non-TOF, both visually and when comparing MAPE and COM. This was noticeable for the lung lesion in the thoracic cavity but also for other parts of the anatomy such as the liver. This is likely due to the improved image contrast of NAC TOF images.

A limitation of the work presented here is that it only evaluates the feasibility of performing a motion correction and makes no effort to apply this to an attenuation corrected reconstruction. This is addressed in the next section 3.4.

A further limitation of the work presented in this section (and more generally the chapter as a whole) is that the data and acquisition simulation used are quite naive, simplistic, and make some unrealistic choices. For instance, the lesion is quite big at

40.0 mm in diameter, it is unlikely that motion correction is necessary to locate such a large lesion. The motion used during data generation is more simple, it is driven by one signal and contains no hysteresis. There was no intra-gate motion and the SS used for fitting the MM was taken directly from XCAT.

These findings support the use of NAC TOF images for MM estimation in the remainder of this thesis (where the TOF resolution used in this section is maintained). As the motion correction pipeline becomes more complex and is tested under more realistic conditions in later chapters, maintaining robust motion estimation becomes increasingly important. The insights gained here directly inform the choice of reconstruction approach and motivate further investigation into the effect of noise levels, lesion size, and acquisition duration on MM performance, as discussed in Section 3.6.

## 3.4 PET/CT Respiratory Motion Correction With a Single Attenuation Map Using NAC Derived Deformation Fields

This section investigates the possibility of using MMs for inter-respiratory cycle motion correction in PET/CT, taking attenuation into account. It is not enough to fit a MM alone, some kind of motion compensated reconstruction should take place, as was discussed in section 2.6.6. In addition, this section explores whether iterative estimation of both the motion parameters and warping of a single  $\mu$ -Map, from any respiratory position, increases the accuracy of AC reconstruction. This section largely follows (Whitehead et al. 2020).

#### 3.4.1 Introduction

There are different strategies for handling attenuation correction in conjunction with motion correction. In clinical practice, usually a single  $\mu$ -Map is available, derived from CT in one respiratory state. This can introduce an unwanted bias (through misaligned anatomy) into the reconstructed images, even after motion correction, unless the deformation of the attenuation is taken into account. Other motion correction methods can incorporate, directly, both motion correction and  $\mu$ -Maps estimation into reconstruction. However,

this can be computationally expensive (Bousse et al. 2016b), (Bousse et al. 2016a).

This section expands upon the previous section 3.2 by addressing the problem of the misaligned  $\mu$ -Map in the final reconstruction. It uses the estimated MM to deform a coincidentally located  $\mu$ -Map (at the location of the reference position for the motion correction) to the position of the gates for attenuation correction. This then allows for a AC OSEM reconstruction of the original data with the deformed  $\mu$ -Map, hopefully with less artifacts as when using a static misaligned  $\mu$ -Map. This section then goes on to establish if it is worthwhile to fit a further MM on this new data or if there is no further improvement to the accuracy of the fit (on NAC data) MM.

#### 3.4.2 Methods

**Table 3.3:** This table shows a short summary of the highlights of the methods used in section 3.3 and section 3.4. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	<b>Motion Estimation</b>
Details			
Section 3.3	• XCAT:  – 40.0 mm lesion	• NAC OSEM (2 iterations, 24 subsets)	Jointly estimating registration and MM
	- 1D XCAT respiratory traces - No hysteresis	• Gaussian post-smoothing (6.4 mm FWHM)	<ul><li>Non-diffeomorphic</li><li>1D SS</li><li>SSD objective function</li></ul>
	<ul><li>6 volumes</li><li>No intra-gate motion</li><li>Simulation:</li></ul>		• Default parameters (CPG spacing 5 times voxel size, 5e-3 bending energy
	<ul> <li>Simulation:</li> <li>2D Simulation</li> <li>120 s simulation</li> <li>American clinical count rate (True 60 KCPS)</li> <li>Non-TOF and TOF</li> </ul>		weight)
	<ul> <li>- 375.0 ps TOF resolution</li> <li>- No scatter or random events</li> <li>• Gating:</li> <li>- XCAT SS</li> <li>- 6 respiratory bins</li> </ul>		

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.4	XCAT:     20.0 mm lesion     1D RPM respiratory traces     No hysteresis     240 volumes     Intra-gate motion      Simulation:     2D Simulation     120 s simulation     American clinical count rate (True 60 KCPS)     375.0 ps TOF resolution     No scatter or random events      Gating:     1D PCA respiratory SS     10 respiratory bins	Reconstruction for motion estimation:     NAC OSEM (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm FWHM)      Motion compensated reconstruction:     AC OSEM (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm FWHM)	<ul> <li>Activity MM:         <ul> <li>Jointly estimating registration and MM</li> <li>Non-diffeomorphic</li> <li>1D respiratory SS</li> <li>SSD objective function</li> <li>Tuned CPG spacing (12 times voxel size) and bending energy weight (9e-6)</li> </ul> </li> <li>Attenuation to activity registration:         <ul> <li>Non-diffeomorphic</li> <li>NMI objective function</li> <li>Tuned CPG spacing (5 times voxel size) and bending energy weight (9e-2)</li> <li>Multi-resolution registration (4 levels)</li> </ul> </li> <li>Attenuation deformed twice</li> </ul>

The method proceeds as follows: NAC TOF volumes are reconstructed using OSEM, section 3.4.2.3. These volumes are used as input for MM estimation, section 3.4.2.4. A single  $\mu$ -Map is then warped to the volumes, using the MM, and the volumes are AC, section 3.4.2.5. Following on from this, another motion estimation and correction cycle is performed, section 3.4.2.7.

For evaluation, XCAT simulations are used, for one bed position, with a FOV including the base of the lungs and the diaphragm. Compared to the previous section, the simulations were more realistic by taking intra-cycle variations into account, as described in section 3.4.2.1 and 3.4.2.2. The output from the proposed method is evaluated against a non-motion corrected reconstruction of the same data both visually, using a profile, as well as with SUV analysis, seen in section 3.4.2.7.

#### 3.4.2.1 XCAT Volume Generation

Volume generation follows the same basic procedure as presented in section 3.3.2.1. However, here 240 volumes were generated over a 120s respiratory trace (with inter-

respiratory cycle variation) derived from data captured using a RPM. The max displacement of AP and SI motion was set to 1.2 cm and 2.0 cm respectively. The FOV included the base of the lungs, diaphragm, and the top of the liver with a 20 mm diameter spherical lesion placed into the centre of the right lung.

#### 3.4.2.2 PET Acquisition Simulation

The PET simulation (and reconstruction) follows the procedure described in section 3.3.2.2. However, here the 240 volumes were forward projection (including attenuation). Noise realisations were used to simulate an overall acquisition over 120 s, emulating a standard single bed position acquisition.

A respiratory SS was then generated using PCA (Thielemans et al. 2011), for more information see section 2.6.3. This was used to gate the data into 10 respiratory bins using displacement gating. For the purpose of the MM fitting, SS values were obtained for the post-gated data by taking an average of the SS values of the data in each bin.

#### 3.4.2.3 Non-Attenuation Corrected Image Reconstruction

As before, data were reconstructed, without attenuation correction using OSEM with two full iterations and 24 subsets (Hudson and Larkin 1994). Volumes were post-filtered using a 3D Gaussian blur with a kernel size of 6.4 mm full width half maximum.

#### 3.4.2.4 Motion Model Estimation

This section improved on the method described in section 3.3.2.4. SSD was used as the similarity measure and bending energy was used as a penalty. The CPG spacing and penalty weight (bending energy) were tuned using a grid search.

#### 3.4.2.5 Attenuation Map Warping

A  $\mu$ -Map close to the mean respiratory position was selected from the  $\mu$ -Maps generated by XCAT. This  $\mu$ -Map was then registered (using NMI) to the mean position NAC motion compensated image (generated using the MM). In order to ensure that the selected  $\mu$ -Map is as close as possible to the mean respiratory position the SS value of the selected  $\mu$ -Map was subtracted from all SS values, effectively biasing the MM to the position of the selected  $\mu$ -Map. The MM was then used to generate DVFs for the SS values of each bin, which were then used to warp the  $\mu$ -Map from the mean respiratory position to each bin.

**Table 3.4:** Comparison of SUV<sub>max</sub>, SUV<sub>median</sub> and SUV<sub>peak</sub> between ungated (CINE-CT), ungated (static CT), NAC MM, and AC MM.

SUV	Max	Median	Peak
<b>Ungated (CINE-CT)</b>	4.63	2.73	3.39
<b>Ungated (static CT)</b>	4.66	3.05	3.54
NAC MM	5.56	3.18	4.07
AC MM	5.43	3.18	4.00

### 3.4.2.6 Motion Corrected Image Reconstruction with Attenuation Correction

Data were re-reconstructed, with attenuation correction, using the  $\mu$ -Maps from section 3.4.2.5. The same reconstruction parameters as in section 3.4.2.6 were used. These data were then either motion corrected using the original NAC MM or a new MM was fit on the AC volumes as in section 3.4.2.4.

#### 3.4.2.7 Evaluation

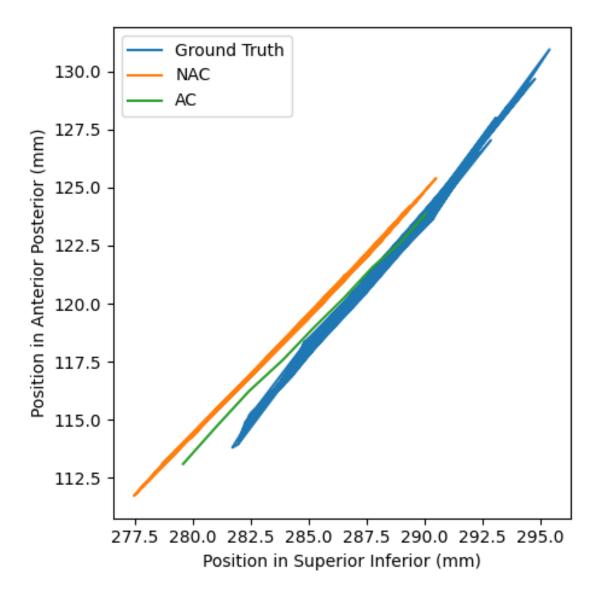
To evaluate the validity of the MM results, the COM of the lesion, over time, was tracked for both NAC and AC reconstructions as in section 3.3.2.5.

In addition to the reconstructions performed in section 3.4.2.6, data were also reconstructed after simply summing all gates together. This was done using either a sum of all  $\mu$ -Maps (to emulate an averaged CINE-CT) or one  $\mu$ -Map, positioned close to the mean respiratory position. This process matches current clinical practice.

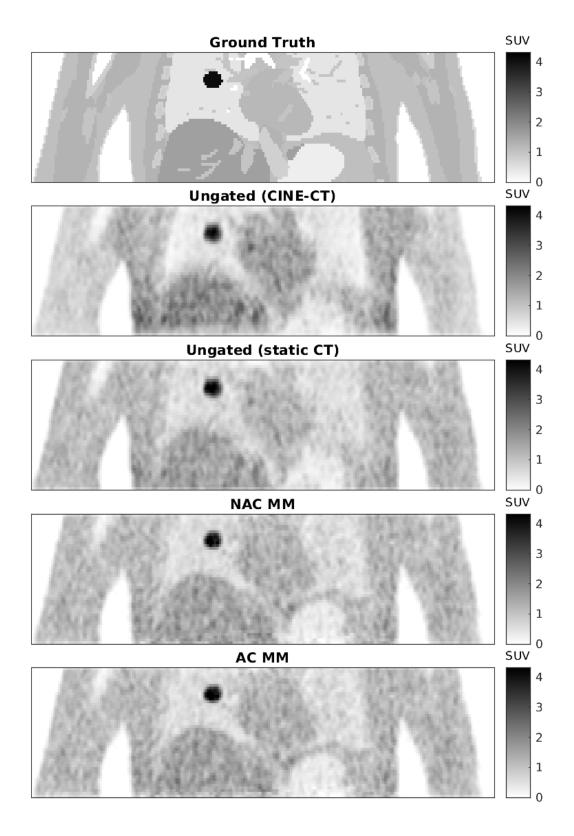
Comparisons used included a profile over the lesion and  $SUV_{max}$ ,  $SUV_{median}$ , and  $SUV_{peak}$ .  $SUV_{peak}$  here was defined following European Association of Nuclear Medicine (EANM) guidelines (Boellaard et al. 2015).

#### 3.4.3 Results

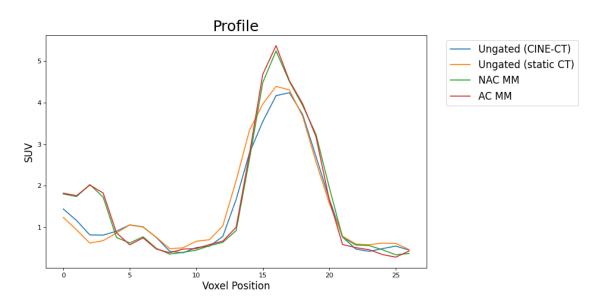
COM results can be seen in figure 3.3. The slope of the line of the COM of both the NAC and AC matches closely the ground truth COM. However, both results appear to be biased slightly by approximately the same value. This is potentially caused by the subtraction of the  $\mu$ -Map SS value from all SS values. Furthermore, attention should be paid to the fact that the minimum resolution of this figure is the size of one voxel, thus small variations could be discretised to a whole voxel away.



**Figure 3.3:** The path of the COM of the lesion, in voxel indices. Horizontal (respectively vertical) axis corresponds to motion in the SI (respectively AP). Different curves denote COM displacement for ground truth data, the estimated data from the NAC based MM, and the estimated data from the AC based MM.



**Figure 3.4:** Ground truth and reconstructions using ungated (CINE-CT), ungated (static CT), NAC MM, and AC MM. Colour map ranges are consistent for all images.



**Figure 3.5:** A profile across the lesion for ungated (CINE-CT), ungated (static CT), NAC MM, and AC MM.

The ungated and the MM results for a visual analysis can be seen in figure 3.4. When compared visually structures can be seen, less blurred, in the MM data that cannot be seen in the ungated data. For instance, structures at the boundary between the diaphragm and the lung. The different levels of blurring in the ungated (CINE-CT) and static CT could be attributed to the constraint put on the reconstruction by having a sharp  $\mu$ -Map in one respiratory position in the static CT case. Additionally the lesion itself can be seen to be more homogeneous, this can be further observed in the profile across the lesion in figure 3.5.

SUV results can be seen in table 3.4 and consistently show that SUVs are greater for the MM over the ungated method.

#### 3.4.4 Discussion

Results from both a visual analysis, a comparison of profiles and SUVs show that the MM provides volumes more free from blurring and less susceptible to artefacts when compared to the ungated data. Results also indicate that the NAC MM provides similar results as when re-estimating the motion from the AC volumes, while not requiring the additional computation of the AC MM. MMs fit on either the NAC or MLACF reconstruction are used for the remainder of the thesis. Results indicate that motion correction of interrespiratory cycle motion is possible with this method, while accounting for attenuation

deformation. However, the work presented here has many limitations, as will be discussed below.

Compared to the previous section, the work presented here improves the data and the acquisition simulations applied to this data to be more realistic. The size of the lesion is reduced by half, a realistic respiratory signal is used to drive the motion of XCAT, and the data is actually gated into respiratory bins using a PCA derived SS (rather than being simulated at a respiratory bin without intra-bin motion).

A limitation of the work presented in this section (and more generally the chapter as a whole) is that the data and acquisition simulation continues to make some unrealistic choices. For computational purposes the acquisition is still simulated as a 2D acquisition. Although an RPM signal is used to drive the XCAT motion it is still very simple and continues to be driven by one signal and thus contains no hysteresis.

A major limitation of the work here is the position of the  $\mu$ -Map used for attenuation registration. It is not likely that the available  $\mu$ -Map from CT is coincidentally located close to the mean respiratory position. This assumption was made here to give the method the best chance of working. It is more likely that a  $\mu$ -Map from CT will be located more closely to max inhalation, due to the breath hold procedure used in common clinical practise. This work also deforms the  $\mu$ -Map twice, once via a registration to the mean respiratory position and once to the position of each gate (by inverting the deformation from the MM for each gate). This is suboptimal as there is the potential for artifacts to be introduced by deforming twice. This is addressed in the next section.

A further limitation of the work presented here (and by extension the work presented in section 3.3) is that no comparison with other motion estimation is provided. MM based methods are common in CT but not in PET, therefore an evaluation of the improvement gained by forgoing straight pair- or group-wise and applying a MM based approach is necessary. This will be presented in the next section.

# 3.5 Comparison of Motion Correction Methods Incorporating Motion Modelling for PET/CT Using a Single Breath Hold Attenuation Map

This section completes the feasibility study in this chapter by comparing different motion correction techniques, with and without incorporating MMs. Compared to the previous section, lower count data are simulated. In addition, the  $\mu$ -Maps used are at end inhalation, to better reflect clinical practice. The work is this section is based on (Whitehead et al. 2021).

# 3.5.1 Introduction

Previously in this chapter no evaluation of the benefits of including a MM has been provided. This section seeks to demonstrate the regularisation benefit of using a MM, where the registration and MM are not jointly fit. The reason for not using the same joint estimation of deformations and MM as previous sections (section 3.2 and section 3.4) is twofold. Firstly, it is not fair to demonstrate a difference between a straight registration and a registration augmented with a MM unless the only change in the second case is the inclusion of a MM. Secondly, by swapping to a more modular framework it allows for all aspects of the motion correction to be easily changed. For instance, in this section we also examine pair- and group-wise registration strategies. This modular framework also allows for the easy use of NMI for all registrations (and therefore the sharing of parameters between the activity and attenuation registration). To not fit the MM and the registration jointly means that rather than fitting directly a MM which parametrises the DVFs we iteratively swap between fitting the DVFs and the MM (on the DVFs).

Furthermore, the simulation used in this section is again an improvement of the data used in the previous section. Here, the count rate is lowered to an approximate European clinical count rate, pseudo-scatter and random events are included in the simulation, and the size of the lesion is reduced.

A method incorporating MMs for dynamic PET/CT, was proposed and tested on clinical data in (Chan et al. 2018). The work presented here differentiates itself, specifically from previous work in motion modelling for PET/CT (Chan et al. 2018), by firstly using a

2D SS, rather than a 1D SS. Thus both irregular breathing and hysteresis can be included in the model, at the expense that each gate contains fewer counts, for more information see section 2.6.5. Additionally, the group-wise method, presented here, makes use of an iterative motion correction algorithm rather than using only a pair-wise method.

# 3.5.2 Methods

**Table 3.5:** This table shows a short summary of the highlights of the methods used in section 3.4 and section 3.5. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.4	XCAT:     20.0 mm lesion     1D RPM respiratory traces     No hysteresis     240 volumes     Intra-gate motion      Simulation:     2D Simulation     120 s simulation     American clinical count rate (True 60 KCPS)     375.0 ps TOF resolution     No scatter or random events      Gating:     1D PCA respiratory SS     10 respiratory bins	Reconstruction for motion estimation:     NAC OSEM (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm FWHM)      Motion compensated reconstruction:     AC OSEM (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm FWHM)	<ul> <li>Activity MM:         <ul> <li>Jointly estimating registration and MM</li> <li>Non-diffeomorphic</li> <li>1D respiratory SS</li> <li>SSD objective function</li> <li>Tuned CPG spacing (12 times voxel size) and bending energy weight (9e-6)</li> </ul> </li> <li>Attenuation to activity registration:         <ul> <li>Non-diffeomorphic</li> <li>NMI objective function</li> <li>Tuned CPG spacing (5 times voxel size) and bending energy weight (9e-2)</li> <li>Multi-resolution registration (4 levels)</li> </ul> </li> <li>Attenuation deformed twice</li> </ul>

3.5. Comparison of Motion Correction Methods Incorporating Motion Modelling for PET/CT Using a Single Breath Hold Attenuation Map

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.5	XCAT:     12.0 mm lesion     2D MR respiratory traces     Hysteresis     240 volumes     Intra-gate motion      Simulation:     2D Simulation     120 s simulation     European clinical count rate (True 48 KCPS)     375.0 ps TOF resolution     Pseudo scatter (triangle filter 22 voxels, 39 KCPS) and random (56 KCPS) events      Gating:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>NAC OSEM (2 iterations, 24 subsets)</li> <li>Replication of end slices to size of μ-Map</li> <li>Gaussian post-smoothing (2 times voxel size)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing (6.39 mm FWHM transverse, 3.27 mm axial)</li> </ul> </li> </ul>	Registration:  Non-diffeomorphic  NMI objective function  Tuned CPG spacing (12 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (8)  Multi-resolution registration (4 levels)  Pair- and group-wise registration  Activity and attenuation DVF composed  MM:  With and without MM  MM fit only at highest resolution level  Weighted (counts) linear regression MM  2D respiratory SS

As before, XCAT ground truth volumes were generated but now with more realistic respiratory motion input, section 3.5.2.1, used as input for the simulation and reconstructed without AC, section 3.5.2.2. These volumes are used as input for registration (both pairand group-wise) and MM estimation, section 3.5.2.3 and section 3.5.2.4. The reference PET volume from the previous step, is then registered to the single  $\mu$ -Map, at end inhalation. The resultant MM and its inverse are used to perform a motion compensated AC reconstruction, section 3.4.2.5 and section 3.5.2.6. This process continues iteratively: performing an activity registration, resampling a reference volume, performing an attenuation registration, composing the DVFs, fitting a MM fitting a MM.

The output for pair- and group-wise registration (both with and without an additional MM fitting step) is evaluated against a non-motion corrected reconstruction of the same data (using a single static  $\mu$ -Map and a  $\mu$ -Map average from CINE-CT for attenuation correction): visually, using a profile, as well as with SUV analysis, section 3.4.2.7.

# 3.5.2.1 XCAT Volume Generation

Volume generation follows the same basic procedure as presented in section 3.3.2.1. However, to increase realism, the SS used to drive XCAT was derived from MR navigator patient data. To acquire this signal a pencil shaped bar is placed over an edge which undergoes motion in the AP direction and another in the SI direction. The location of the edge along the pencil shaped bar is determined using an edge detection algorithm (looking for sudden changes in the value along the pencil shaped bar). The position of this edge is then tracked over the frames and the position is used as a SS by taking it's absolute position. Moreover, a smaller spherical lesion of 12 mm diameter (smaller than the max displacement due to respiratory motion) was placed into the base of the right lung (within the max displacement, due to respiratory motion, of the diaphragm).

# 3.5.2.2 PET Acquisition Simulation and Non-Attenuation Corrected Image Reconstruction

Simulation procedure was similar as in section 3.4.2.2. Pseudo-randoms and scatter were added. Randoms were added by summing the scaled mean value to each voxel of each volume prior to forward projection. Pseudo-scatter was added by summing the scaled and smoothed mean  $\mu$ -Map prior to forward projection. The smoothing parameter was optimised to give scatter which tapered at the same rate as in clinical data. A full scatter simulation was not performed due to software limitations.

Noise was simulated, such that data matched an acquisition over 120 s, emulating a standard single bed position acquisition. The count rate was lower here than in the section 3.3 and section 3.4 (48 KCPS compared to 60 KCPS). This count rate was determined from a static clinical acquisition on the GE Discovery 710.

A respiratory SS was generated using PCA (Thielemans et al. 2011). The magnitude of this signal and its gradient, was used to gate data into 30 respiratory bins using displacement gating (10 amplitude and three gradient bins), this was as discussed in section 2.6.2. Gates with fewer than 0.42 % of the counts were discarded. For the purpose of the MM fitting, SS values were determined for the post-gated data by taking an average of the SS values of data in each bin.

Data were reconstructed without attenuation correction, using OSEM with two full

iterations and 24 subsets as before.

# 3.5.2.3 Registration

Before being registered, each volume underwent pre-processing, including replication of end-slices, transformation to be approximately normally distributed histograms (Johnson et al. 2013) and post-smoothing. This pre-processing was only applied to intermediate data and was not used for the final output of the method. Because a breath hold  $\mu$ -Map is the final target position for the motion correction, 10 repeating slices are added to the top and bottom of each volume to allow space for the volumes to be registered into.

Two registration methods were examined in this work. Firstly, pair-wise registration, where the reference position was selected as the gate with the highest number of counts. All other gates were registered to it. Secondly, group-wise registration, where after an initial pair-wise registration step, the DVFs generated had the inverse mean of all DVFs composed with them, before a new reference volume was resampled. Registration to the new reference volume, followed by the inverse mean composition and resample, continued for a set number of iterations. NiftyReg (Modat et al. 2010) was used to perform registrations using a B-spline parametrisation, a multi-resolution registration was used with four levels. The Gaussian smoothing FWHM, CPG spacing of the B-spline coefficients, bending energy regularisation term weight, and number of iterations were tuned using a grid search.

## 3.5.2.4 Motion Model Estimation

When using a MM for motion estimation, it was fit as a direct RCM on the DVFs from section 3.5.2.3 and the SS from section 3.5.2.2. A weighted linear regression was used, where the weighting was taken based on the number of counts in each gate. Note that this functionality was not available in the NiftyRegResp/SuPReMo software used in previous sections. Once a MM was fit, new DVFs were generated for each gate, using the SS values used to fit the MM. For group-wise registration, MM fitting occurred between iterations, the DVFs generated by the MM were used to resample the new target volume at each iteration.

# 3.5.2.5 Attenuation Map Warping

A  $\mu$ -Map at end inhalation was selected from the  $\mu$ -Maps generated by XCAT. The PET volume from the previous step was then registered to this  $\mu$ -Map. The resulting DVFs were composed with the DVFs from the last iteration of the motion estimation method, and a new volume resampled. The inverse of these DVFs were then used to warp the  $\mu$ -Map to each gate.

# 3.5.2.6 Motion Corrected Image Reconstruction with Attenuation Correction

Data were re-reconstructed with attenuation correction, using the  $\mu$ -Maps from section 3.5.2.5. The same reconstruction parameters as in section 3.5.2.6 were used. Motion correction was then applied to data following section 3.5.2.3, section 3.5.2.4, and section 3.5.2.5. Volumes were post-filtered using a Gaussian smoothing, with a FWHM of 6.39 mm in the transverse plane (equivalent to three voxels) and 3.27 mm (equivalent to one voxel) in the axial direction. The target of applying the method like this is to have a single motion corrected volume at the position of the  $\mu$ -Map, where motion estimation was from NAC volumes. MCIR was implemented by deforming the post-reconstructed image and summing.

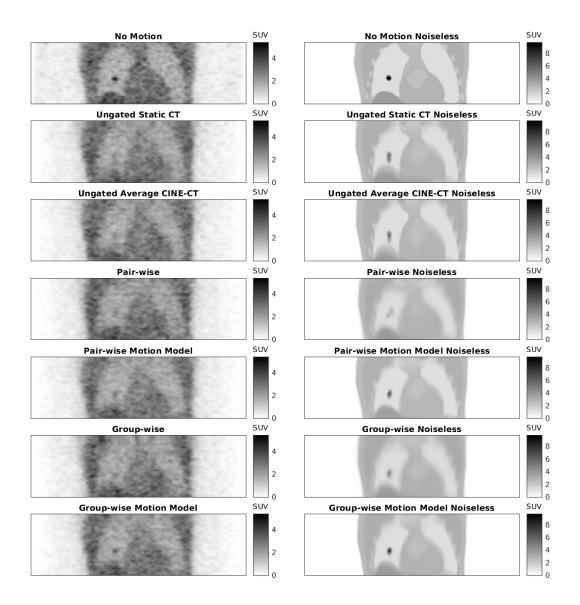
## 3.5.2.7 Evaluation

In addition to the reconstructions performed in section 3.5.2.6, data were also reconstructed without motion correction, using either a sum of all  $\mu$ -Maps (to emulate an averaged CINE-CT), or the end inhalation  $\mu$ -Map. For the present evaluation, the volumes without motion correction were then registered to the position of the end inhalation  $\mu$ -Map. Additionally, DVFs generated by each method were also applied to noiseless data for visual analysis.

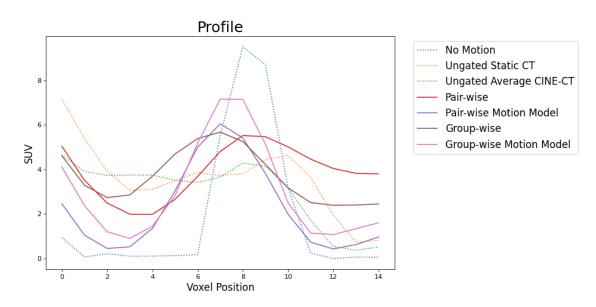
Comparisons used included the following. A visual analysis, a profile over the lesion, and  $SUV_{max}$  and  $SUV_{peak}$  (defined following EANM guidelines (Boellaard et al. 2015)).

## **3.5.3** Results

A visual comparison of the reconstructed images (see figure 3.6) shows that the motion makes the lesion almost disappear in the ungated reconstruction. Motion of the diaphragm



**Figure 3.6:** Coronal slices for various methods. First column contains AC motion corrected reconstructions and the second column contains the result of applying the final motion correction on the original XCAT images (for easier assessment of the accuracy of the estimated DVFs). This is for ungated static CT, ungated averaged CINE-CT, pair-wise, pair-wise MM, group-wise, and group-wise MM. Colour map ranges are consistent for all images in each column.



**Figure 3.7:** A profile across the lesion for ungated static CT, ungated averaged CINE-CT, pairwise, pair-wise MM, group-wise, and group-wise MM.

**Table 3.6:** Comparison of SUV<sub>max</sub> and SUV<sub>peak</sub> between ungated static CT, ungated averaged CINE-CT, pair-wise, pair-wise MM, group-wise, and group-wise MM.

SUV	Max	Peak
No Motion	9.50	9.06
Ungated Static CT	5.25	5.15
<b>Ungated averaged CINE-CT</b>	5.38	5.07
Pair-wise	4.21	3.92
Pair-wise MM	6.63	6.07
Group-wise	4.42	4.21
Group-wise MM	7.64	7.03

is somewhat recovered in the registration methods. However, the methods that incorporate a MM visually outperform the other methods, with the lesion recovered best by the groupwise MM method.

Similar conclusions can be drawn from the profiles through the lesion (see figure 3.7). The peak of is greater for MM methods than for MM-free methods. However, the latter still have a larger peak than the ungated methods.

SUV results consistently show that including MMs increases the SUV when compared to when one is not used (see table 3.6).

# 3.5.4 Discussion

Results from a visual analysis, a comparison of profiles, and SUV, shows that adding a MM to any motion correction method (tested here) improved the quality of volumes produced. This indicates that for very noisy data, MMs are almost required.

An advantage of the modular framework for motion correction presented here, when compared to the joint estimation as was used in section 3.3 and section 3.4, is that it allows for the easy incorporation of the μ-Map alignment. In NiftyRegResp/SuPReMo (McClelland et al. 2017) the input is a series of floating images and one static image (plus SS), natively the framework would perform a pair-wise registration/MM (if the framework was run multiple times where the static image was replaced with a resampled image then it would be a group-wise version of the same). However, it doesn't seem obvious how to then incorporate the registration to the  $\mu$ -Map and the composition (giving one MM which parametrises DVFs from the floating images to the  $\mu$ -Map). It may be possible to use the  $\mu$ -Map as the static input image to NiftyRegResp/SuPReMo. However, it is not expected that this registration would perform well, for one the only objective function available in the software package is SSD, which would fail for different modalities with different scales and distributions. Furthermore, the work presented in this section uses a weighted MM, where more weighting is given to volumes which contain more counts, which also does not exist in NiftyRegResp/SuPReMo. To only enforce the  $\mu$ -Map as the reference position for the MM by using its SS value (as was done in section 3.4) was shown to be insufficient by the bias noticeable in figure 3.3 of section 3.4.3 (furthermore, it is not obvious how to extract the SS of a CT volume from clinical data). While it may be technically possible to implement a joint estimation of DVFs and MM, simultaneously for both the registration of the activity volumes and the registration to the position of the  $\mu$ -Map, it is not a straightforward endeavour (and would need to change drastically if any part of the motion correction pipeline was to change). Therefore, where we want to provide an extensive evaluation of different combinations of aspects of the motion correction framework (as will appear in chapter 4) we will not jointly estimate the DVFs and MM.

A major disadvantage of the implementation used in this section is that a diffeo-

morphic registration method was not used. By using NiftyReg (Modat et al. 2010) through SIRF (Ovtchinnikov et al. 2020) (or through NiftyRegResp/SuPReMo (McClelland et al. 2017) as in section 3.3 and section 3.4) diffeomorphic registration was not possible due to software limitations. This is a major disadvantage as non-diffeomorphic registrations can cause folding, which is non-anatomical and makes inverting the resulting DVF problematic, while inverting the DVF is a core aspect of the methods presented here.

While this section used more realistic simulations than the previous ones, see table 3.5, the data and acquisition simulation continues to make some unrealistic choices. For computational purposes the acquisition is still simulated as a 2D acquisition. The TOF resolution is higher than the simulated GE Discovery 710 (although PET/CT scanners with even higher TOF resolution do exist). This will be rectified in chapter 4.

Despite these limitations, the findings from this section validate the potential of MMs-based motion correction and establish a foundation for extending these methods to more realistic and clinically relevant scenarios.

# 3.6 Discussion and Conclusion

The methods implemented in this chapter ignore the existence of methods such as MLAA and MLACF, as discussed previously in section 2.2.3.1 and section 2.4.3. Relatively satisfactory results were obtained on the data as simulated here. This does not mean though that motion correction would not benefit from an improved reconstruction, especially in situations where the quality of the data are worse. This will be further investigated in the next chapter.

Furthermore, the SS used during the XCAT simulation, for most of this work, is identical (but flipped) for both the AP and SI motion (where the XCAT and RPM respiratory traces are used in section 3.3.4 and section 3.4). This means that there is no hysteresis in the motion, which is unrealistic. This could be solved by giving separate signals (a 2D signal, for instance as in section 3.5 where a 2D respiratory trace was acquired from MR) for both the AP and SI motion, which are captured using an MR scanner, rather than a 1D signal captured from a RPM. This has begun to be addressed by using an MR derived signal in section 3.5.

The lesion used to demonstrate the improvement by the motion correction method

was unreasonably large. Because of the size of the lesion motion correction wouldn't be needed to identify and locate it. In order to better highlight the advantage of the motion correction method a smaller lesion placed closer to the diaphragm should be used. This also has begun to be addressed in section 3.5.

The method to simulate the PET acquisition in this chapter has some limitations. Firstly, in the work presented here, a relatively high count rate is used for the noise caused by the PET acquisition physics. In the next chapter we will examine how the motion correction methods perform under more strenuous circumstance, for instance, research scan count rate levels. Additionally, in both section 3.3.2.2 and section 3.4.2.2, no random or scatter simulation were used, while in section 3.5 pseudo-random and scatter were used. The reason for a lack of random or scatter simulation was because the software available did not provide this functionality when using TOF simulations.

Furthermore, as discussed in section 3.5.4, the TOF resolution used in this work, so far, does not reflect any current scanner. This will be addressed in the following chapter.

In the initial experiments, the registrations performed as part of the motion modelling were fit using SSD. This objective function is sensitive to differences in the magnitude of the input. If gates are registered to a mean volume, as in group-wise registration seen earlier in section 2.6.4, then the magnitude of the inputs will almost never be equal, although this problem could be solved by preprocessing the data. Additionally, if SSD was used to register a PET activity distribution to a CT  $\mu$ -Map, it would also fail. This was addressed in section 3.5 where NMI was used. However, the method used did not fit the motion correction and MM simultaneously.

Alternatively, because a MM does not need to be fit on the data it is applied to, the MM could be fit on coarsely gated data and applied to finely gated data. This is an advantage unique to MMs and is also something that, as far as we know, has not yet been attempted in PET. This is advantageous because for a 1D MM technically only two data points are required to fit it (as it is fitting a line), for a 2D MM three data points are required to fit is (as it is fitting a plane). Using 30 data points, as in section 3.5, is redundant an could actually lower the quality of the MM due to the noise in the reconstructions. However, more gates are advantageous when forming a final motion compensated reconstruction. In an ideal world, the maximum number of gates should

be used during the motion compensated reconstruction such that all motion, which can be differentiated based on the resolution of the scanner, is accounted for. MMs hold the unique advantage that they can be fit on and applied to different numbers of data points, the only requirement is that the SS values remain consistent.

There are many different variations on motion correction, we have shown in this chapter that we have comparable results to other motion correction work, as presented in section 2.6.7. We will go on to demonstrate in the next chapter that we have continued development of the general modular framework, first presented in section 3.5. This will allows for the fair comparison of aspects of different motion correction methods.

Further limitations and future developments are discussed in chapter 6.

# **Chapter 4**

# **Applying Motion Correction Utilising Motion Modelling to Realistic Simulations**

# 4.1 Overview

This chapter presents the next stage in the development of motion correction methods for PET/CT, building on the feasibility study introduced in Chapter 3. That preliminary study was valuable for testing core components of the framework-such as DVF estimation and MM fitting-but relied on simplified simulation data. At the beginning of the previous chapter these idealised conditions did not reflect the complexity of real clinical scenarios, overlooking key challenges such as variability, noise, TOF resolution, and (to a degree)  $\mu$ -Map mismatch. Advances were made towards the end of the previous chapter but there was still room for improvement.

To address these limitations, this chapter presents an enhanced evaluation of the motion correction framework under more realistic settings. While still based on synthetic data, simulations are designed to better reflect patient-like conditions by incorporating realistic respiratory motion, attenuation effects, and scanner-specific characteristics. Furthermore, this chapter introduces the use of MLACF reconstruction in conjunction with the existing MM. This enables joint estimation of activity and attenuation in the presence of motion, allowing the framework to be assessed under conditions that more closely resemble real-world PET/CT imaging.

The first section of this chapter introduces work to investigate the impact of using MMs to allow for motion correction to be performed on coarsely gated data and then applied to finely gated data. This addresses a practical constraint in clinical workflows, where high gating resolution is often limited by count statistics. This is also a novel approach that to our knowledge has not been attempted in PET/CT (Whitehead et al. 2022c).

The second section of this chapter highlights an evaluation of the impact of 2D and 1D gating as well as showing the impact of using MLACF or OSEM for reconstruction prior to motion correction. These comparisons help assess how reconstruction method and gating strategy affect the accuracy and robustness of motion correction under more realistic conditions.

The third section of this chapter utilises the same framework as the second section of this chapter, however it seeks to evaluate the method further. Specifically, this section investigates the impact of lesion size on motion correction capabilities. Here we attempt to find if the motion correction fails for much smaller lesions and what impact this has on the final reconstructed image.

Finally, the fourth section of this chapter discusses the problems or limitations of the work performed in the previous three sections with reference to where problems and limitations highlighted in section 3.6 have been addressed.

# 4.2 Introduction

In the previous chapter, a flexible framework for motion correction of high TOF resolution PET data was introduced, where the reference was set to the position of a breath hold  $\mu$ -Map. Preliminary experiments indicated that the combination of both the MM and TOF, allowed to estimate the motion on volumes reconstructed without AC, while obtaining good alignment with the  $\mu$ -Map, increasing resolution and quantification accuracy (for simulations with high TOF resolution and count levels). However, in preliminary investigations of situations with low contrast lung tumours, we observed that the registration of the NAC volumes was very challenging. This chapter therefore seeks to extend upon the previous chapter by incorporating an approach suggested by (Lu et al. 2018): as opposed to estimating the motion on NAC OSEM volumes, Lu et al. used MLACF reconstructions

for the motion estimation. MLACF (Nuyts et al. 2012) provides volumes with greater contrast than NAC, without introducing bias due to misalignment with a fixed  $\mu$ -Map, see section 2.2.3.1, section 2.4.3, and section 3.5.4, while still offering relatively fast computation time.

However, as MLACF attempts to find an ACF for every sinogram bin, it requires data with more counts to function. To avoid intra-gate blurring, it is advantageous to use a high number of gates (e.g., 30). When using MLACF, a majority of the gates would have too few counts to produce meaningful reconstructions. This is where a further benefit of the MM method shines. As discussed in section 2.2.3.1 and in section 4.1, it is possible to fit a MM on different data to which it is applied to (as long as the SS remains consistent). Therefore, it may be possible to fit a MM on data which has been coarsely gated into few bins, and then apply this MM to data which has been finely gated.

This chapter extends the method of the previous chapter in several ways. Firstly, by incorporating MLACF as an initialisation technique (for the motion estimation). Secondly, low noise low temporal/gate resolution data is used to fit the MM, while high noise high temporal/gate resolution data is used for the output. This potentially allows for reduced mis-registration, as well as improving computation time. Thirdly, a diffeomorphic symmetric velocity field parameterised registration was used, which provided DVFs free from folding. In addition, the methods are tested on realistic simulations and count levels (similar to patient data from research study performed at our institution)..

A method incorporating MMs for dynamic PET/CT, was proposed in (Chan et al. 2018). Additionally, a method incorporating MLACF for PET/CT, was presented in (Lu et al. 2018). The work displayed here, differentiates itself, not only by the ways mentioned above (for instance, course fitting of the MM, fine application of the MM), but also by using a 2D SS (rather than a 1D SS). Thus, allowing more general parametrisation of the motion (here in terms of both displacement and velocity). Additionally, the group-wise method, presented here, makes use of an iterative motion correction algorithm, rather than relying only on a pair-wise method, as was used in the work mentioned above.

# 4.3 PET/CT Motion Correction Exploiting Motion Models Fit on Coarsely Gated Data Applied to Finely Gated Data

# 4.3.1 Introduction

As discussed above, it is desirable to use a motion estimation method which uses MLACF reconstructions as part of the MM fitting (Whitehead et al. 2022c), but this exacerbates problems with low count data. This section investigates the impact of gating coarseness on motion estimation performance, with respect to the noise present in each gate.

In preparation for a more complete evaluation in the next section, this section also addresses some of the limitations highlighted in section 3.5.4 and section 3.6.

Firstly, the data used in this work is made more realistic and challenging by comparing to the late time point of the dynamic PET research data used in chapter 5. Therefore, the count rates (trues, scatter and randoms) are taken as an average of the count rates for the patients used in chapter 5. Furthermore, the pseudo-scatter simulation was improved such that the profile of the scatter is matched as closely as possible to the patient data. This is achieved by smoothing the  $\mu$ -Map using a Gaussian smoothing with a 128.0 mm FWHM and forward projecting.

Secondly, the software was improved such that a fully diffeomorphic (velocity field parametrised and fit symmetrically, from floating to reference frame and from reference to floating frame) registration method is used throughout.

## 4.3.2 Methods

# 4.3. PET/CT Motion Correction Exploiting Motion Models Fit on Coarsely Gated Data Applied to Finely Gated Data

**Table 4.1:** This table shows a short summary of the highlights of the methods used in section 3.5 and section 4.3. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	<b>Motion Estimation</b>
Details			
Section 3.5	XCAT:     12.0 mm lesion     2D MR respiratory traces     Hysteresis     240 volumes     Intra-gate motion      Simulation:     2D Simulation     120 s simulation     European clinical count rate (True 48 KCPS)     375.0 ps TOF resolution     Pseudo scatter (triangle filter 22 voxels, 39 KCPS) and random (56 KCPS) events      Gating:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>NAC OSEM (2 iterations, 24 subsets)</li> <li>Replication of end slices to size of μ-Map</li> <li>Gaussian post-smoothing (2 times voxel size)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing (6.39 mm FWHM transverse, 3.27 mm axial)</li> </ul> </li> </ul>	Registration:  Non-diffeomorphic  NMI objective function  Tuned CPG spacing (12 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (8)  Multi-resolution registration (4 levels)  Pair- and group-wise registration  Activity and attenuation DVF composed  MM:  With and without MM  MM fit only at highest resolution level  Weighted (counts) linear regression MM  2D respiratory SS

# 4.3. PET/CT Motion Correction Exploiting Motion Models Fit on Coarsely Gated Data Applied to Finely Gated Data

Method	Data	Reconstruction	Motion Estimation
Details			
Section 4.3	XCAT:     12.0 mm lesion     2D MR respiratory traces     Hysteresis     480 volumes     Intra-gate motion      Simulation:     2D Simulation     240 s simulation     European (late dynamic) research count rate (41 KCPS)     375.0 ps TOF resolution     Pseudo scatter (Gaussian smoothing 128.0 mm FWHM, 29 KCPS) and random events (57 KCPS)      Gating for motion estimation:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient) or     4 radial pseudo-phase gates      Gating for motion compensated reconstruction:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC MLEM initialisation (1 iteration)</li> <li>MLACF (7 activity iterations, 24 subsets, 9 attenuation iterations)</li> <li>Quadratic prior (1.0)</li> <li>ACF initialised with ones</li> <li>Activity volume and ACF standardised between iterations</li> <li>Replication of end slices to size of μ-Map</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	Registration: Diffeomorphic (velocity field parametrised, symmetric)  NMI objective function Tuned CPG spacing (8 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (5)  Multi-resolution registration (5 levels) Pair-wise registration initialisation Group-wise registration Activity and attenuation DVF composed  MM: MM fit between each resolution level Weighted (counts) linear regression MM  ID or 2D respiratory SS

General overview of the method: MLACF (coarsely gated), motion estimation and alignment of  $\mu$ -Map, final AC OSEM reconstruction incorporating motion and  $\mu$ -Map alignment (finely gated).

AC volumes are reconstructed using MLACF (both 'phase' gated into four and amplitude gated into 30 bins), as seen in section 4.3.2.1, section 4.3.2.2, and section 4.3.2.3. These volumes are used as input for group-wise registration and MM estimation, as seen in section 4.3.2.4. The reference PET volume, from the previous step, is then registered to a single  $\mu$ -Map, at end inhalation. The resultant MM and its inverse are used to perform a motion compensated AC reconstruction, as seen in section 4.3.2.5.

For validation, XCAT simulations are used, for one bed position, with a FOV includ-

ing the base of the lungs and the diaphragm. The output for data 'phase' gated into four and amplitude gated into 30 bins is evaluated against the following. Firstly, the above procedure was additionally performed (for both data 'phase' gated into four and amplitude gated into 30 bins) where the registration and MM fitting was performed on noiseless data. Secondly, a non-motion corrected reconstruction of the same data (using a single static  $\mu$ -Map and a  $\mu$ -Map averaged from CINE-CT for attenuation correction). Evaluations performed included a visual analysis (Structural Similarity Index Measure (SSIM) (Wang et al. 2004)), an examination of profile, and SUV analysis, seen in section 4.3.2.6.

# 4.3.2.1 XCAT Volume Generation

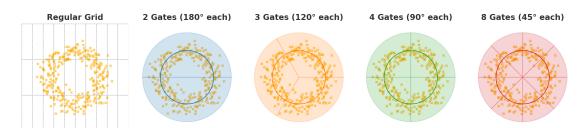
XCAT (Segars et al. 2010) was used to generate 480 volumes over a 240 s period, using respiratory traces, derived from 2D MR patient data, as was previously discussed in section 3.5.4 and section 3.6. This means that because the XCAT receives two different signals for the AP and SI motion then they can move independently from one another (in a realistic way), which introduces hysteresis into the data. The maximum displacement of AP and SI motion, was set to 1.2 cm and 2.0 cm respectively. Activity concentrations were derived from a static [<sup>18</sup>F]-FDG patient scan. The FOV included the base of the lungs, diaphragm, and the top of the liver. A 12 mm diameter spherical lesion (smaller than the max displacement due to respiratory motion) was placed into the base of the right lung (within the maximum displacement due to respiratory motion of the diaphragm).

# 4.3.2.2 PET Acquisition Simulation

PET acquisitions were simulated (and reconstructed), using STIR (Thielemans et al. 2012), (Efthimiou et al. 2019b) through SIRF (Ovtchinnikov et al. 2020), to forward project data using the geometry of a GE Discovery 710. 2D simulations were used to reduce computation time. TOF mashing was incorporated to reduce computation time, resulting in 13 TOF time bins of size 376.5 ps. Attenuation was included using the relevant  $\mu$ -Maps generated by XCAT. For information related to the geometry of the scanner and the specifics of the volumes and sinograms used please see section 3.3.2.2.

Pseudo-randoms and scatter were added. Randoms were added by summing a scaled uniform forward projection to the simulation sinogram to obtain the same randoms-to-trues ratio as observed in the clinical data. Pseudo-scatter was added by summing a

### Comparison of Gating Strategies on 2D Surrogate Signal



**Figure 4.1:** This figure shows an example of the amplitude and gradient 2D SS gating technique as compared to examples of the pseudo-phase based gating technique. On the right of this image is a more traditional gating technique, here using 10 amplitude gates and three gradient gates. The next four images show the pseudo-phase based gating with two, three, four, and eight radial bins, plus an example of how a circular separation can be used for deep and shallow breaths.

scaled and smoothed forward projection of the activity data at each time point to the simulation sinogram. The smoothing parameter was optimised to give scatter with a similar distribution to a late time point dynamic PET data. A full scatter simulation was not performed due to software limitations.

Noise was simulated, such that data matched a single bed position acquisition over 240 s. The count rate was selected to match that of research scans (41 KCPS true counts), below that of diagnostic clinical scans. This count rate was selected as a 'worst case scenario'.

A respiratory SS was generated using PCA (Thielemans et al. 2011). The value of this signal, and its gradient, was used for gating. For MM estimation, data were initially pseudo-phase gated. The data was displacement and gradient gated into four bins, where each bin was a quadrant centred on the maximum or minimum of the displacement or gradient. This is useful as it gives a small number of gates while allowing to fit a 2D MM (a plane). For the purpose of the MM fitting, SS values were determined for the post-gated data by taking an average of the SS values in each bin.

The 'phase' gating mentioned above differs from traditional phase gating (assigning a percentage from end-inhalation to end-inhalation, for instance, as discussed in section 2.6.2). The inspiration for this method was, when plotting the 2D SS it was noted that it follows a circular path, where the width of this path is the breath to breath variability. If using a more traditional 2D gating technique, like in section 3.5, this is like imposing a

regular grid over the circular 2D SS. Therefore, it seems natural to use a gating technique based on bisecting this circle. The gates can be described by the angle of the circle that they cover, for instance if two gates are used they cover 180.0° each, three cover 120.0° each, four cover 90.0° each, and eight cover 45.0° each. It is also possible to extend this method to differentiate deep and shallow breaths by bisecting the circle of the SS with one or more additional circles. A graphical example of these concepts can be seen in figure 4.1

# 4.3.2.3 MLACF Image Reconstruction

Data were reconstructed using MLACF (described in section 2.4.3), with seven full iterations and 24 subsets for the activity update, and nine iterations for the attenuation update (Nuyts et al. 2012). MLACF was initialised using one iteration of MLEM, where the breath hold CT was used for attenuation correction (the ACF were still initialised with all ones), this gave a small improvement to results without introducing bias. Between iterations, the activity volume and ACF sinogram were normalised (scaled to the same magnitude as an equivalent volume/sinogram filled with ones), and a small value (epsilon) was added to each voxel, this was to aid with stabilising the reconstruction. A quadratic prior was included in the reconstruction (this is as was discussed in section 2.4.2), to promote smoothness of the ACF sinogram.

# 4.3.2.4 Registration and Motion Model Estimation

Before being registered, each volume underwent pre-processing, including replication of end-slices, smoothing, and standardisation. This pre-processing was only applied to intermediate data, and was not used for the final output of the method.

Group-wise registration was used, where after an initial pair-wise registration step (the reference volume for pair-wise registration was selected as the one which most data was binned into), a new reference volume was resampled. Registration to the new reference volume, followed by another resample, continued for a set number of iterations. NiftyReg (Modat et al. 2010) was used to perform registrations, using a diffeomorphic symmetric velocity field B-spline parametrisation. CPG spacing of the B-spline coefficients, bending energy regularisation term weight, and number of iterations were tuned using a grid search.

Between each iteration, the resampled volume was registered to the position of the  $\mu$ -Map. The DVFs for both the group-wise registration, and  $\mu$ -Map registration were composed together to form one final DVF. With the diffeomorphic nature of the registration, this would give a MM (after the MM is fit) which could generate DVFs to and from the gates to the  $\mu$ -Map.

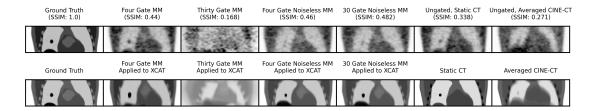
MMs were fit as a direct RCM on the DVFs, from section 4.3.2.4 and the SS from section 4.3.2.2. A weighted linear regression was used, where the weighting was taken based on the number of counts in each gate. Once a MM was fit, new DVFs were generated for each gate, using the SS values that were used to fit the MM. MM fitting occurred between iterations, the DVFs generated by the MM were used to resample the new target volume at each iteration.

A multi-resolution approach to registration is common (Modat et al. 2010). This is where the images are downsampled and a registration is applied to these low resolution images, then the resolution of both the images and the DVF are increased until they return to their original resolution. This is beneficial as it allows for the macro elements of the image to be aligned before focusing on higher frequency information. This provides benefits in situations where there is noise and also where there are flat regions of the image. Previously, in section 3.5 the MM was only fit on the highest level resolution DVF. Here however, the MM is fit at every resolution level of the registration.

# 4.3.2.5 Motion Corrected Image Reconstruction with Attenuation Correction

Data were re-gated, using the value of SS, and its gradient, to gate data into 30 respiratory bins (10 displacement and three gradient bins).

Data were re-reconstructed with attenuation correction, deforming the  $\mu$ -Maps, using the inverse of the DVFs determined by the MM from section 4.3.2.4. For reconstruction, OSEM was utilised, with two full iterations and 24 subsets (as in clinical practice) (Hudson and Larkin 1994). Motion correction was then applied, using DVFs determined by the MM from section 4.3.2.4. This is a standard method once motion and AC position is known (Kyme and Fulton 2021), (Fayad et al. 2015), (Meng et al. 2023), (Pösse et al. 2020), see section 2.6.6. Volumes were post-filtered as in clinical practice, employing



**Figure 4.2:** First row contains, AC motion corrected reconstructions (plus SSIM to the ground truth). Second row contains the results of applying the final MM on the original XCAT volumes. This is using a MM fit on four and 30 gate binned data applied to 30 gate binned data, a MM fit on noiseless four and 30 gate binned data applied to 30 gate binned data, and ungated data AC with a static  $\mu$ -Map at end inhalation and all  $\mu$ -Maps summed. Colour map ranges are consistent for all images in each row.

a Gaussian smoothing, with a FWHM of 6.4 mm in the transverse plane, and a 'normal' Z-filter, a GE specific post-processing step where the image is convolved in the axial direction with a kernel containing 1,4,1 (normalised).

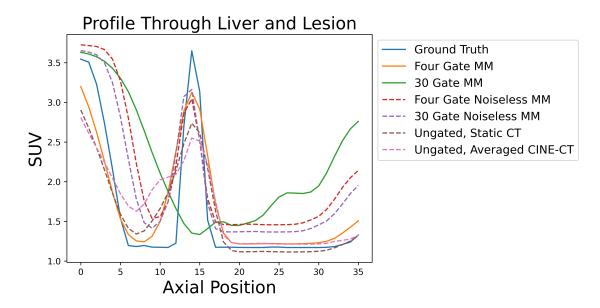
### 4.3.2.6 Evaluation

The motion estimation was evaluated on both noiseless and noisy data. This results in six cases to be evaluated: where data has been pseudo 'phase' gated into four bins or 30 bins (10 displacement and three gradient), both with and without noise, where gating (and therefore motion correction) has not been applied and AC has been applied with a static  $\mu$ -Map at end inhalation, and where gating (and therefore motion correction) has not been applied and AC has been applied with a  $\mu$ -Map which is an average of all  $\mu$ -Maps (to emulate an averaged CINE-CT). For the present evaluation, the volumes without motion correction were registered to the position of the end inhalation  $\mu$ -Map.

As in the previous chapter, DVFs generated by each method were also applied to the original XCAT volumes, for visual analysis. Other comparisons made included SSIM to the ground truth (Wang et al. 2004), (Wang and Bovik 2009), a profile over the lesion, and  $SUV_{max}$  and  $SUV_{peak}$  (defined following EANM guidelines (Boellaard et al. 2015)).

## 4.3.3 Results

A visual comparison of the reconstructed images (see figure 4.2) shows that the high noise high temporal/gate resolution method performs quite poorly, most probably due to the high level of noise apparent in the volumes. Conversely, the low noise low temporal/gate resolution data method appears to be able to motion correct the data without being too



**Figure 4.3:** A profile through the lesion in the SI direction, summed over a window in the AP and Lateral Medial (LM) directions, with median smoothing, for the ground truth XCAT data, a MM fit on four and 30 gate binned data applied to 30 gate binned data, a MM fit on noiseless four and 30 gate binned data applied to 30 gate binned data, and ungated data AC with a static  $\mu$ -Map at end inhalation and all  $\mu$ -Maps summed.

**Table 4.2:** Comparison of SUV<sub>max</sub> and SUV<sub>peak</sub>, for the ground truth XCAT data, a MM fit on four and 30 gate binned data applied to 30 gate binned data, a MM fit on noiseless four and 30 gate binned data applied to 30 gate binned data, and ungated data AC with a static  $\mu$ -Map at end inhalation and all  $\mu$ -Maps summed.

SUV	Max	Peak
Ground Truth	8.76	7.96
Four Gate MM	8.04	6.18
30 Gate MM	1.77	1.32
Four Gate Noiseless MM	8.05	6.24
30 Gate Noiseless MM	7.96	5.99
Ungated, Static CT	6.61	5.08
Ungated, averaged CINE-CT	5.65	4.44

adversely affected by the noise level.

The peak of the profile (see figure 4.3) for the four gate MM results is comparable to the noiseless results. In contrast, the 30 gate MM method fails on the noisy data. For all other motion correction methods, the peak is greater than without motion correction.

SUV (and SSIM) results confirm the above (see table 4.2).

# 4.3.4 Discussion

Results show that using a low number of gates for MM fitting drastically improves motion correction when there is a high level of noise in the gates, while having minimal impact at low noise. In addition, the execution time using a reduced number of gates is lower.

A key limitation of this section, as well as the preceding ones, stems from two principal issues related to the simulated data and acquisition process. Firstly, although the simulations were configured with a count rate and scanner geometry consistent with the GE Discovery 710 system, they employed an enhanced TOF resolution representative of a more modern PET/MR scanner. For the evaluation presented here to reliably inform the method's application to clinical data, the simulation conditions must accurately reflect the characteristics of the target system. Therefore, it is essential that the method be tested using data that incorporates the actual TOF resolution of the GE Discovery 710.

Secondly, all simulations conducted thus far have used 2D PET acquisition models. While this decision may be justified by the considerable computational cost of 3D simulation (particularly when MLACF is included), it is nevertheless important to demonstrate that the method remains valid under more realistic 3D conditions. Ensuring that the approach performs robustly in such settings is critical for supporting its eventual clinical translation.

A secondary limitation of the work so far is that there has been no example of the affect of the number of gates used or the dimensionality of the SS. Previously 30 bins (10 amplitude and three gradient) were selected because 10 amplitude bins are sufficient to cover the range of motion present in the data based on the resolution of the scanner, and three gradient bins cover where the signal is increasing in value, decreasing in value, and is stationary. Four pseudo-phase bins were selected as they were a convenient number for both the pseudo-phase gating and the MM fitting. However, it has not been demonstrated

if this is a sensible assumption. Furthermore, it has been claimed that a 2D SS is necessary to model hysteresis in the MM, however it has not been shown what impact this has on the motion correction and the reconstruction.

These limitations will be addressed in the next section.

# 4.4 Evaluation of PET/CT Motion Correction Incorporating Motion Models Using MLACF and Complex Gating Schemes

# 4.4.1 Introduction

This section evaluates the method presented in section 4.3 on simulation data that are made as similar as possible to GE Discovery 710 patient data mentioned in section 4.3.2, specifically by using a 3D PET simulation and the TOF resolution of the GE Discovery 710 (550.0 ps) (as opposed to the higher resolution used in previous results).

Where not otherwise stated, the methods used here followed section 4.3. However, to be able to handle 3D PET data, the MLACF and registration methods were made more robust and stable, compared to what was used in section 4.3. The improvements made to MLACF here are specifically tailored to motion correction and were evaluated by the improvement they afforded to the registration.

The section presents an evaluation of the impact of including MLACF (vs NAC or AC OSEM) as well as various gating schemes (including the use of 2D SSs).

## 4.4.2 Methods

**Table 4.3:** This table shows a short summary of the highlights of the methods used in section 4.3 and section 4.4. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	Motion Estimation
Details			
Section 4.3	XCAT:     12.0 mm lesion     2D MR respiratory traces     Hysteresis     480 volumes     Intra-gate motion      Simulation:     2D Simulation     240 s simulation     European (late dynamic) research count rate (41 KCPS)     375.0 ps TOF resolution     Pseudo scatter (Gaussian smoothing 128.0 mm FWHM, 29 KCPS) and random events (57 KCPS)      Gating for motion estimation:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient) or     4 radial pseudo-phase gates      Gating for motion compensated reconstruction:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC MLEM initialisation (1 iteration)</li> <li>MLACF (7 activity iterations, 24 subsets, 9 attenuation iterations)</li> <li>Quadratic prior (1.0)</li> <li>ACF initialised with ones</li> <li>Activity volume and ACF standardised between iterations</li> <li>Replication of end slices to size of μ-Map</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	Registration: Diffeomorphic (velocity field parametrised, symmetric)  NMI objective function Tuned CPG spacing (8 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (5)  Multi-resolution registration (5 levels) Pair-wise registration initialisation Group-wise registration Activity and attenuation DVF composed  MM: MM fit between each resolution level Weighted (counts) linear regression MM  1D or 2D respiratory SS

# 4.4. Evaluation of PET/CT Motion Correction Incorporating Motion Models Using MLACF and Complex Gating Schemes

Method	Data	Reconstruction	Motion Estimation
Details Section 4.4			
Section 4.4	• XCAT:	• Reconstruction for motion es-	• Registration:
	– 12.0 mm lesion	timation:	- Diffeomorphic (velocity
	- 2D MR respiratory traces	- AC OSEM initialisation (2	field parametrised,
	- Hysteresis	iterations, 24 subsets)	symmetric)
	– 480 volumes	- MLACF (3 activity iterations, 24 subsets, 1 atten-	- NMI objective function
	- Intra-gate motion	uation iterations)	- Tuned CPG spacing (13 times voxel size) and
	Simulation:	– ACF initialised with $\mu$ -Map	bending energy (1e-1 at-
	- 3D Simulation	- Activity volume scale set	tenuation)
	– 240 s simulation	to OSEM reconstruction and	- Multi-resolution registra- tion (4 levels)
	- European (late dynamic) research count rate (41 KCPS)	– ACF scale set to μ-Map between iterations	- Iterations used at resolution level: 2, 2, 2, 5
	– 550.0 ps TOF resolution	<ul> <li>Re-interpolation of outliers between iterations</li> </ul>	<ul> <li>Pair-wise registration initialisation</li> </ul>
	- Pseudo scatter (Gaussian smoothing 128.0 mm FWHM, 29 KCPS) and random events (57 KCPS)	<ul> <li>Smoothing of endplanes between iterations</li> <li>Median smoothing (kernel size 3) between iterations</li> </ul>	<ul><li>Group-wise registration</li><li>Activity and attenuation</li><li>DVF composed</li></ul>
	<ul> <li>Gating for motion estimation:</li> <li>1D or 2D PCA respiratory</li> <li>SS</li> <li>30 respiratory bins (10 dis-</li> </ul>	<ul> <li>Sinogram mashing in attenuation update (1, 7, 2, 3 mashing factor respectively)</li> <li>Filled with Not a Numbers</li> </ul>	MM:     MM fit between each resolution level     Weighted (counts) linear regression MM
	placement, 3 gradient) or	(NaNs) to size of $\mu$ -Map	- 1D or 2D respiratory SS
	- 5 respiratory bins (5 displacement, 1 gradient) or	Motion compensated reconstruction:	
	- 3, 4, 8, or 12 radial pseudo- phase gates	- AC OSEM (2 iterations, 24 subsets)	
	Gating for motion compensated reconstruction:     1D or 2D PCA respiratory SS	- Gaussian post-smoothing transverse (6.4 mm FWHM) - 'Normal' Z-filter axially	
	- 30 respiratory bins (10 displacement, 3 gradient)	(1.0, 4.0, 1.0 convolution)	

General overview of the method: MLACF (coarsely gated), motion estimation and alignment of  $\mu$ -Map, final AC OSEM reconstruction incorporating motion and  $\mu$ -Map alignment (finely gated). A high level overview of the method as well as pseudocode will be presented with the final evaluation of the method in section 4.5.

AC volumes are reconstructed using MLACF (as well as using AC and NAC OSEM for comparison) with improvements for stability ('phase' gated into three, four, eight, and 12 bins and amplitude gated into five and 30 bins), as seen in section 4.4.2.1 and

section 4.4.2.2. These volumes are taken as input for a multi-resolution group-wise registration and MM estimation. The reference PET volume, from the previous step, is registered at each iteration to a single  $\mu$ -Map, at end inhalation. The resultant MM and its inverse are used to perform a motion compensated AC reconstruction, all as seen in section 4.4.2.3.

For validation, XCAT simulations are used, for one bed position, with a FOV including the base of the lungs and the diaphragm. The output for data 'phase' gated into three, four, eight, and 12 and amplitude gated into five and 30 bins are evaluated against the following. Firstly, a ground truth XCAT volume. Secondly, a non-motion corrected reconstruction of the same data (using a single static  $\mu$ -Map for attenuation correction). Evaluations performed included a visual analysis (Perception Based Image Quality Evaluator (PIQE) (Venkatanath et al. 2015)), an examination of profile, and SUV analysis, shown in section 4.4.2.4.

# 4.4.2.1 XCAT Volume Generation and PET Acquisition Simulation

XCAT volumes were acquired in the same way as in section 4.3.2.1.

PET acquisitions were simulated (and reconstructed) similarly to in section 4.3.2.2. However, a great deal of care was taken to ensure the simulated data matched real GE Discovery 710 patient data as closely as possible. Count levels were set to those from a research study (which used the GE Discovery 710, the same data discussed in section 4.3.2 and used in chapter 5). The TOF resolution used for simulation was the same as the GE Discovery 710 PET/CT (550.0 ps), rather than the improved TOF resolution from the GE Signa PET/MR (375.0 ps, as was used in all previous work). Pseudo-randoms and scatter were added in a similar way as in section 4.3.2.2. Furthermore, for the first time, 3D acquisition data was used rather than 2D. The change to 3D acquisition data may seem negligible for the evaluation of a motion correction algorithm, however, when MLACF is introduced, the change to 3D data becomes more significant. Previously, 2D simulations and reconstruction were used purely for computational efficiency. Beyond the importance of aligning the simulations with real GE Discovery 710 patient data, ensuring that the simulated conditions closely reflect those of the scanner increases the likelihood that the method will generalise effectively to clinical data.

A respiratory SS was generated employing PCA (Thielemans et al. 2011). The value of this signal, and its gradient was used for gating (where appropriate). Multiple gating schemes were utilised for MM estimation. These gating schemes were as follows.

- 'Phase' gating follows the same scheme as used in section 4.3.2.2. In the case of Four Phase, the data was displacement and gradient gated into four bins, where each bin was a quadrant centred on the maximum or minimum of the displacement or gradient. However, there do not necessarily have to be exactly four bins. Thus, for Eight Phase and Twelve Phase it is as if each quadrant from Four Phase is split equally into two or three further bins respectively. For Three Phase the 'circle' of the 2D SS is instead split into three one third sectors (or 120.0° sectors).
- Amplitude gating follows the same scheme as used in section 3.5.2.2. In the case of Thirty Amplitude, the magnitude of the SS and its gradient, was used to gate data into 30 respiratory bins using displacement gating (10 amplitude and three gradient bins). In the case of Five Amplitude, five respiratory bins were used (five amplitude and one gradient bin). For Five Amplitude both a 2D and 1D SS was used. The 2D SS was utilised to mimic Thirty Amplitude. However, in the Five Amplitude case, only one gradient bin was used. Thus, the mean value for the gradient bin was used for each displacement bin in the Five Amplitude 2D case. In the Five Amplitude 1D case, the gradient value was set to zero for each displacement bin (for both fitting and during use). This is the most simple and fair way to implement a 1D SS (so as not to change the gating and MM fitting code too drastically). However, this does mean that the MM regression has an additional degree of freedom (the linear regression uses an extra parameter which is always the same value).

# 4.4.2.2 Image Reconstruction

MLACF reconstruction mainly followed section 4.3.2.3. Three full iterations and 24 subsets for the activity update, and one iteration for the attenuation update were used (Nuyts et al. 2012). One iteration for the attenuation update was found to be sufficient. This is the case because the first update of the ACFs was large and subsequent updates were small. More joint iterations of activity and attenuation could be performed, in equal or

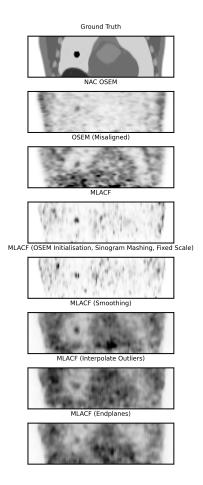


Figure 4.4: A visual analysis which shows the results of performing different reconstruction methods on XCAT data with a 12 mm lesion placed into the base of the right lung (the left of this image). Where noise is present then this data represents a 60 s acquisition on a GE Discovery 710 (roughly equivalent to one gate from a clinical acquisition). At the top of this image is the Ground Truth image (without motion or reconstruction artifacts), second from the top of this image is a NAC OSEM reconstruction, third from the top of this image is a AC OSEM reconstruction using a  $\mu$ -Map at a different respiratory position than the Ground Truth, fourth from the top of this image is a MLACF reconstruction following the literature (Nuyts et al. 2012). From the fifth image from the top to the bottom of this image all images represent adding additional processing steps to attempt to make the reconstructions better for motion correction. Fifth from the top of this image is the result of initialising MLACF with a misaligned AC OSEM reconstruction, using sinogram mashing, and fixing the scale of the image to the scale of the misaligned AC OSEM reconstruction. Sixth from the top of this image is the result of applying a median smoothing (with a kernel size of three) after each activity update (and before every attenuation update). Seventh from the top of this image is the result of interpolating any value outside three times the Inter Quartile Range (IQR) of the image after each activity update using a linear interpolation. The bottom of this image is the result of heavily smoothing the endplanes. All reconstructed images include scatter and random events. All OSEM reconstructions follow standard clinical practise (see chapter 3). For the NAC OSEM and MLACF following the literature (Nuyts et al. 2012) the output was rescaled to the AC OSEM. Colour map ranges are consistent for all images.

less time, if fewer updates of the ACFs were used for each activity iteration (the updates of the ACFs are quite time consuming).

The activity volume was initialised from a AC OSEM reconstruction (two iterations, 24 subsets) and the ACFs were initialised using the breath hold CT. Both of these initialisations gave a small improvement to results, compared to either initialising with a MLEM iteration or without initialising, without introducing bias as can be seen in figure 4.4.

Between iterations, the scale of the ACF sinogram was set to the scale of the original breath hold CT ACF, this is in order to aid with stabilising the reconstruction. This is similar to work which (was performed simultaneously with this work) where two ACFs were used. Here, ACFs were estimated as a correction of the CT derived ACF (Jolivet et al. 2022). In this other work a regularisation term was also used to keep the scale of the corrected ACFs close to the original ACFs while allowing some variation (Jolivet et al. 2022). Although the other implementation allowed the scale of the ACFs to vary slightly, while the work here constrained it, this is not necessarily an advantage. The activity volume would experience cross talk from the ACFs and as such could compensate for the fact that the scale of the ACFs was fixed. Additionally, fixing the scale of the ACFs removes a hyperparameter (and a regularisation term) which could otherwise cause instability. Regardless, the inclusion of this step, in both works, attempts to address the scale issue of the original implementation of MLACF.

To aid in situations of high noise sinogram mashing was using during the attenuation update, the mashing factors used were one, seven, two, and three in the radial positions, angles, transaxial plane, and TOF dimensions respectively.

To further combat the instability of the MLACF algorithm (specifically when applied to motion correction), processing of the activity volume was performed between each iteration:

• The first step of this processing was the removal and re-interpolation of outlying values. This is due to the fact that outlying values were seen, between iterations, which caused the optimisation to fail. Outliers were defined as being any value outside three times the IQR of the data. New values were interpolated using linear interpolation.

- The second step was to smooth the endplanes of the activity volume, due to the endplanes exhibiting significantly more noise and outlying values than any other slice of the volume. This noise in the endplanes was most likely caused by the reduced sensitivity of the scanner in these areas. To aid in smoothing the endplanes an 'endplanes' volume was created (this was a copy of the activity volume). Outliers were removed from the 'endplanes' volume (in the same way as previously but using only one times the IQR) before a max filter was applied across each slice (with a kernel size of three). The 'endplanes' volume was then summed with the activity volume using a ramp filter for the first and last five slices.
- The final step of the processing was to apply median filtering (with a kernel size of three), in order to promote smoothness of the activity volume.

Reconstructions were also performed in the same way as in chapter 3 for comparison purposes. AC and NAC OSEM (using the breath hold CT) were used with two full iterations and 24 subsets (Hudson and Larkin 1994). For the 'phase' gated methods only MLACF was used. However, for the amplitude gated methods both MLACF and OSEM were utilised.

# 4.4.2.3 Registration, Motion Model Estimation, and Motion Corrected Image Reconstruction with Attenuation Correction

Registrations performed was almost identical to section 4.3.2.4. However, in this work without replication of the endplanes, smoothing, or standardisation. Instead the volume was expanded with slices filled with NaNs. NaNs were used, rather than replication of the endslice, because voxels filled with NaNs would not contribute to the objective function calculation. Pre-smoothing was found to be detrimental to the accuracy of the registration (regularisation of the optimisation was found to be sufficient). Standardisation provided no benefit (as NMI was used as the objective function).

Group-wise registration was also used with a pair-wise initialisation. However, in contrast to in section 4.3.2.4, the reference volume for pair-wise registration was selected by finding the minimum NMI to the breath hold CT. This method for selecting the initial reference volume was chosen as it should give both (or either, depending on the noise levels of the input) the volume which is closest (in its respiratory position) to the  $\mu$ -Map,

**Figure 4.5:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. 'Phase' gated and MLACF reconstructed data with three, four, eight and 12 bins respectively. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row.

or it will give the volume with the least amount of noise. Both of these properties are beneficial for achieving a high quality end result.

The MM was applied to the task of the final motion correction in exactly the same way as in section 4.3.2.5.

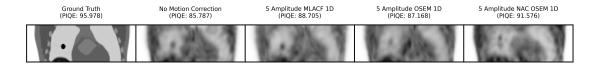
## 4.4.2.4 Evaluation

In addition to the reconstructions performed in section 4.4.2.3, data were also reconstructed without motion correction, using the end inhalation  $\mu$ -Map. For the present evaluation, the volume without motion correction was registered to the position of the end inhalation  $\mu$ -Map. Additionally, DVFs generated by each method were also applied to the original XCAT volumes, for visual analysis.

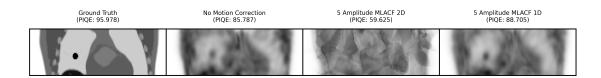
Comparisons made included a visual analysis, PIQE (Venkatanath et al. 2015), (Chan and Goldsmith 2000), a profile over the lesion, and  $SUV_{max}$ ,  $SUV_{median}$ , and  $SUV_{peak}$  (defined following EANM guidelines (Boellaard et al. 2015)). PIQE is a noreference image quality score, which is correlated to the perceptual quality of the image. A high score value indicates high perceptual quality and a low score value indicates low perceptual quality (Venkatanath et al. 2015).

### 4.4.3 Results

A visual analysis of the 'phase' gated results with noise can be seen in figure 4.5 and without noise in figure 4.10. The 'phase' based gating approaches using MLACF universally provided high quality results. The lesion appears more homogeneous and circular in all motion corrected results than in the non-motion corrected result. The Three Phase and Four Phase results are comparable and superior to the Eight Phase and Twelve Phase results. From what was observed in section 4.3, it is to be expected that with fewer gates



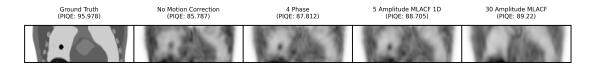
**Figure 4.6:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row.



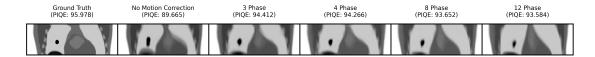
**Figure 4.7:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF. Here, in the 2D case a 2D SS and in the 1D case a 1D SS were used. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row.



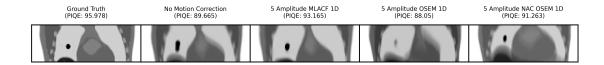
**Figure 4.8:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row.



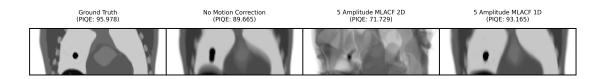
**Figure 4.9:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The methods presented here are a combination of the best one method from figure 4.5, figure 4.6, and figure 4.8. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Colour map ranges are consistent for all images in each row.



**Figure 4.10:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. 'Phase' gated and MLACF reconstructed data with three, four, eight and 12 bins respectively. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row.

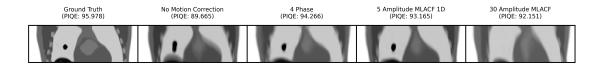


**Figure 4.11:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row.



**Figure 4.12:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF. Here, in the 2D case a 2D SS and in the 1D case a 1D SS were used. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row.

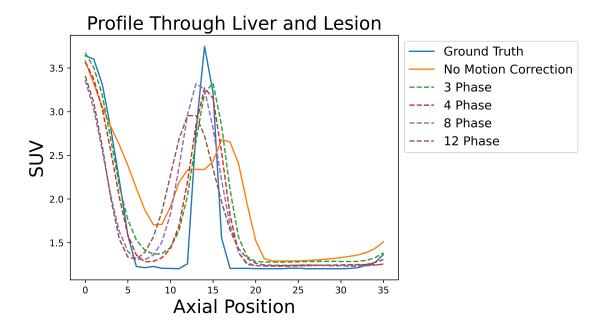
**Figure 4.13:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row.



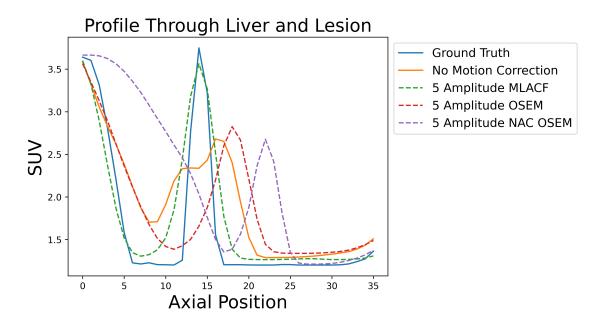
**Figure 4.14:** AC motion corrected reconstructions (plus PIQE), where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The methods presented here are a combination of the best one method from figure 4.10, figure 4.11, and figure 4.13. Here, results are shown applied to volumes without noise. This is due to it potentially being easier to see smaller differences between motion correction methods when noise is removed. However, noise is still used during the motion correction process. Colour map ranges are consistent for all images in each row.

lead to better motion correction. The PIQE results support this, the values for the Three Phase and Four Phase methods are greater than the Eight Phase and Twelve Phase methods. The PIQE values for the Three Phase and Four Phase are also greater than where no motion correction is used and are closest to the ground truth PIQE values.

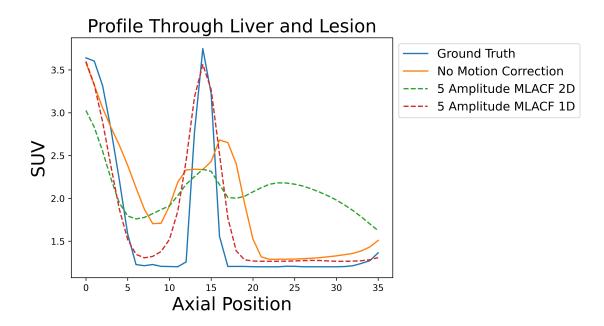
A visual analysis of the amplitude gated results using five bins with noise are shown in figure 4.6 and without noise in figure 4.11. Here, it can be seen that MLACF (when used with amplitude gating of five bins) significantly improves results when compared to AC and NAC OSEM. The lesion in the MLACF example is comparable to the lesions present in the 'phase' based results. However, for both the AC and NAC OSEM results the lesion does not appear in the same location as in the ground truth and also looks very non-homogeneous. The results from both OSEM methods actually seem to be worse than doing no motion correction at all. The PIQE values somewhat support this, however PIQE appears to have scored the NAC OSEM results highly. This scoring could be attributed to



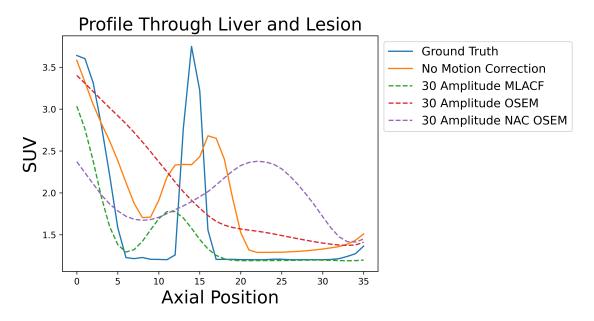
**Figure 4.15:** A profile through the lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. This is for where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. 'Phase' gated and MLACF reconstructed data with three, four, eight and 12 bins respectively. A MM was also fit using a single bin, this is practically the equivalent of doing no motion correction.



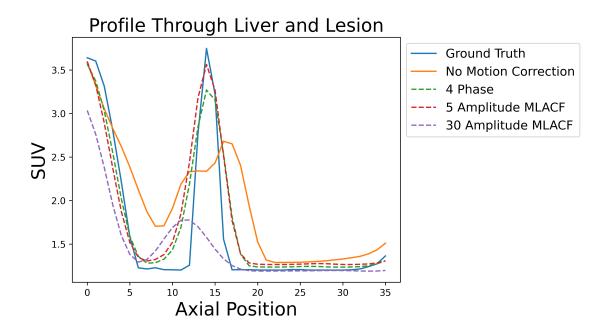
**Figure 4.16:** A profile through the lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. This is for where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction.



**Figure 4.17:** A profile through the lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. This is for where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with five bins reconstructed using MLACF. Here, in the 2D case a 2D SS was used and in the 1D case a 1D SS was used. A MM was also fit using a single bin, this is practically the equivalent of doing no motion correction.



**Figure 4.18:** A profile through the lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. This is where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. Amplitude gated data with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction.



**Figure 4.19:** A profile through the lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. This is where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The methods presented here are a combination of the best one method from figure 4.15, figure 4.16, and figure 4.18. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction.

**Table 4.4:** Comparison of SUV<sub>max</sub>, SUV<sub>median</sub>, and SUV<sub>peak</sub>, where different data was used to fit the MM and the final MM was applied to 30 gate binned data. The data used to fit the MM was as follows. 'Phase' gated and MLACF reconstructed data with three, four, eight and 12 bins respectively. Amplitude gated data with five bins reconstructed using MLACF, AC OSEM, and NAC OSEM (Here, in the 2D case a 2D SS was used and in the 1D case a 1D SS was used.). Amplitude gated data with 30 bins reconstructed using MLACF, AC OSEM, and NAC OSEM. A MM was also fit using a single bin, which is practically the equivalent of doing no motion correction.

SUV	Max	Median	Peak
Ground Truth	9.03	2.47	8.20
No Motion Correction	5.69	2.12	4.31
Three Phase	8.09	2.62	5.78
Four Phase	7.48	2.65	5.75
Eight Phase	7.76	1.42	5.60
12 Phase	7.68	1.30	4.53
Five Amplitude MLACF 2D	4.06	1.75	3.33
Five Amplitude MLACF 1D	8.50	2.31	6.63
Five Amplitude OSEM 1D	4.16	1.40	2.16
Five Amplitude NAC OSEM 1D	2.64	1.34	2.24
30 Amplitude MLACF 1D	3.41	1.21	1.71
30 Amplitude OSEM 1D	2.43	1.38	1.86
30 Amplitude NAC OSEM 1D	5.64	1.28	3.37

this image being smoother and thus potentially more perceptually pleasing.

A visual analysis of the difference between the amplitude gated results with five bins using a 2D and a 1D SS with noise can be seen in figure 4.7 and without noise in figure 4.12. The 2D SS was detrimental to motion correction when compared to a 1D SS. This could be by a lack of constraint along the gradient dimension, as the 2D SS is used for all 'phase' based methods, which show good results. There does not appear to be a fair way to compare where a 1D and a 2D SS are used. In order to incorporate the 2D SS it would mean making fundamental changes to the gating of the 1D SS. If the second dimension of the SS was only gated into two bins then it means the counts in each bin approximately half (five amplitude bins and one gradient bin compared to five amplitude bins and two gradient bins, five total bins compared to 10 total bins). However, if the number of amplitude bins was halved to compensate for this then it also means a fundamental change to the amplitude gating (six amplitude bins and one gradient bin compared to three amplitude bins and two gradient bins).

A visual analysis of the amplitude gated results using 30 bins with noise can be seen in figure 4.8 and without noise in figure 4.13. All methods which utilised 30 bins failed, as in section 4.3. As with the PIQE values for the five bin methods, the PIQE values here are high. This again could be caused by the smoothness of the images where motion correction has failed (perceptually pleasant natural images are smooth, PIQE was not designed for use with medical images).

A visual analysis of the best results from the previous analysis with noise can be seen in figure 4.9 and without noise in figure 4.14. In this case, the Four Phase and Five Amplitude methods perform relatively comparably (and significantly better than the Thirty Amplitude results where motion correction has failed). The lesion and diaphragm in the Four Phase and Five Amplitude cases both show a significant improvement over where no motion correction is used. The Four Phase and Five Amplitude methods both contain a lesion in the same location, and of a similar shape, when compared to the ground truth. The results potentially lead to the conclusion that, with the data used here, it is more important to use fewer bins when gating than it is to use a 2D SS (the XCAT potentially does not have adequately complex motion to highlight the benefit of a 2D SS, for instance lacking sliding motion). PIQE values appear potentially inconclusive when the values

from the Thirty Amplitude method are taken into account.

A profile across the lesion for the 'phase' gated results are presented in figure 4.15. Here, again it can be seen that the Three Phase and Four Phase methods provide profiles with the highest peaks and which most closely match the ground truth. When compared to the visual analysis, it is more obvious that there is a degradation in quality as the number of bins used increases. The Twelve Phase method has a much lower peak which is shifted to the left.

A profile across the lesion for the amplitude gated results using five bins is shown in figure 4.16. The MLACF based method provides an exceptionally good profile. More specifically, the MLACF based method has a profile where the peak almost exactly matches the ground truth. However, the width of the peak for the MLACF method is much greater than the ground truth. The other methods shown here appear to have completely failed, their peak is heavily shifted to the right and is much lower than the MLACF based method.

A profile across the lesion showing the difference between the amplitude gated results with five bins using a 2D and a 1D SS can be seen in figure 4.17, confirming that the 2D SS method failed. Although the peak of the 2D SS method is in the same place as in the ground truth, there is an additional erroneous peak to its right.

A profile across the lesion for the amplitude gated results using 30 bins is presented in figure 4.18, confirming that all methods utilising a large number of bins have failed entirely. If there is a peak in the profile, then it does not match the scale or location of the ground truth at all.

A profile across the lesion showing best results from the previous profiles can be seen in figure 4.19. This shows that the Four Phase and Five Amplitude methods provide very similar profiles. Although the peak of the Five Amplitude method may be greater than the peak of the Four Phase method, the width of the peak for the Four Phase method is less than that of the Five Amplitude method. Additionally, the Four Phase method very slightly matches the ground truth more closely outside the peak, for instance, where the diaphragm is.

 $SUV_{max}$ ,  $SUV_{median}$ , and  $SUV_{peak}$  results can be seen for all methods in table 4.4. An analysis of SUV results shows the following.

• For SUV<sub>max</sub> results, the Three Phase and Five Amplitude (using MLACF) methods

have values which most closely match the ground truth when compared to other methods. All Thirty Gate and OSEM methods have a very low  $SUV_{max}$  value. Strangely, the Four Phase method has a lower  $SUV_{max}$  than the Eight Phase and Twelve Phase methods. This could be attributed to the fact that  $SUV_{max}$  is very susceptible to noise.

- For SUV<sub>median</sub> results, the Three Phase and Four Phase methods have exceptionally high values. These values are even higher than the ground truth, potentially leading to the conclusion that the 'phase' based methods have shrunk the size of the lesion when compared to the ground truth. The Five Amplitude method (using MLACF) has a lower SUV<sub>median</sub> but one which is close to the ground truth.
- For SUV<sub>peak</sub> results, the Five Amplitude (using MLACF) method has a very good score which is closest to the ground truth. The Three Phase and Four Phase methods have values which are very close to each other but they lag behind the Five amplitude (using MLACF) method.

#### 4.4.4 Discussion

The analysis of the results indicates that the Four Phase and Five Amplitude (using MLACF) methods provided the best results, with little differentiation between those two methods. It should be noted that both methods significantly outperformed other methods. For instance, all OSEM and all methods with a large number of bins unanimously failed.

When compared to the methods which failed, it is more interesting to highlight what the two best methods share in common, rather than what separates them. The two best methods both used MLACF and a small number of bins for the estimation of the MM (30 total bins, 10 amplitude and three gradient, were used for MCIR regardless of gating method for the MM fitting). It appears from experiments here, but also when comparing results from this chapter and chapter 3 that when low count data, or a lower resolution for TOF is used, MLACF is almost a necessity for accurate motion correction. Furthermore, the fact that a low number of bins is beneficial when applying motion correction shows a benefit for MMs. This is because using a low number of bins for motion correction but then applying this motion correction to a high number of bins would not be possible without motion modelling.

However, when it comes to the benefits of 2D gating, results here are inconclusive. It could be argued that the simulated data used here was not complex enough to show this (the XCAT does not contain sliding motion and other complex motion that would highlight the benefits of 2D gating).

While PIQE was introduced in an attempt to provide an objective measure of improvement in image quality without the need for ground truth data, results shown here appeared to be not helpful. As was discussed briefly in section 4.4.3, PIQE is a method designed for use in natural images (and not medical images) where smoother images are often more perceptually pleasing. When evaluating patient data a ground truth does not exist for which to use most scoring metrics against, therefore it is imperative to a quality evaluation to develop a metric which does not require a ground truth image. One potential way to achieve this would be to use a NN trained to detect good images. However, a problem with this approach would be that it would be very time consuming to acquire data on which to train and additionally it's not guaranteed that this method would work with different scanners, tracers, or even when using different a FOV.

One limitation of the evaluation performed here is that the size of the lesion remains consistent across all data shown. This is addressed in the next section.

# 4.5 Evaluation of the Effect of Lesion Size on the Quality of Motion Correction

# 4.5.1 Introduction

This section aims to further evaluate the methods presented in section 4.3 and section 4.4. Here, the exact same method as in section 4.4 is used. However, the input data is different. To establish the effect that variable lesion size has on the performance of the method, in addition to the 12.0 mm spherical lesion used in the previous section, we used a 10.0 mm, 8.0 mm, and a 6.0 mm spherical lesion. Smaller lesion sizes were not considered due the resolution of the scanner.

The location of the lesion remains consistent, this is because it is more fair to compare only the difference in the size of the lesion without also changing its location. Also, the location chosen previously is a region which experiences the largest amount of motion in

the lung, and as such is a challenging location for the legion to be in. Therefore, if the method works here then it could be assumed that it will work similarly in other locations in the lung for respiratory motion.

As well as showing the results of the method applied to data with a 12.0 mm, 10.0 mm, 8.0 mm, and a 6.0 mm spherical lesion lesion (with noise and motion), we will also demonstrate for the 8.0 mm lesion the effect of motion (without correction) and what a perfectly motion corrected volume would look like. This is useful as it shows the benefit of applying our motion correction clearly. 8.0 mm was selected for this as it proved to be an exceptionally challenging case, but where it may still be possible to detect the lesion, even with noise.

To evaluate the results a visual analysis is presented, for all lesion sizes, both where the method has been applied to reconstructed data as well as where the MM fit on the reconstructions is applied to noiseless data (to better show only the benefit of the motion correction, unadulterated by artifacts of noise). Furthermore, a profile across the lesion is shown as well as  $SUV_{max}$  and  $SUV_{peak}$  values.

#### 4.5.2 Methods

**Table 4.5:** This table shows a short summary of the highlights of the methods used in section 4.4 and section 4.5. This table is useful for quickly referencing the changes between sections. To see the development of the method throughout the thesis please see table A.1 in appendix A.

Method	Data	Reconstruction	Motion Estimation
Details			
Section 4.4	<ul> <li>XCAT: <ul> <li>12.0 mm lesion</li> <li>2D MR respiratory traces</li> <li>Hysteresis</li> <li>480 volumes</li> <li>Intra-gate motion</li> </ul> </li> <li>Simulation: <ul> <li>3D Simulation</li> <li>240 s simulation</li> <li>European (late dynamic) research count rate (41 KCPS)</li> <li>550.0 ps TOF resolution</li> <li>Pseudo scatter (Gaussian smoothing 128.0 mm FWHM, 29 KCPS) and random events (57 KCPS)</li> </ul> </li> <li>Gating for motion estimation: <ul> <li>1D or 2D PCA respiratory SS</li> <li>30 respiratory bins (10 displacement, 3 gradient) or</li> <li>5 respiratory bins (5 displacement, 1 gradient) or</li> <li>3, 4, 8, or 12 radial pseudophase gates</li> </ul> </li> <li>Gating for motion compensated reconstruction: <ul> <li>1D or 2D PCA respiratory SS</li> <li>30 respiratory bins (10 displacement, 3 gradient)</li> </ul> </li> </ul>	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC OSEM initialisation (2 iterations, 24 subsets)</li> <li>MLACF (3 activity iterations, 24 subsets, 1 attenuation iterations)</li> <li>ACF initialised with μ-Map</li> <li>Activity volume scale set to OSEM reconstruction and</li> <li>ACF scale set to μ-Map between iterations</li> <li>Re-interpolation of outliers between iterations</li> <li>Smoothing of endplanes between iterations</li> <li>Median smoothing (kernel size 3) between iterations</li> <li>Sinogram mashing in attenuation update (1, 7, 2, 3 mashing factor respectively)</li> <li>Filled with NaNs to size of μ-Map</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	<ul> <li>Registration:         <ul> <li>Diffeomorphic (velocity field parametrised, symmetric)</li> <li>NMI objective function</li> <li>Tuned CPG spacing (13 times voxel size) and bending energy (1e-1 attenuation)</li> <li>Multi-resolution registration (4 levels)</li> <li>Iterations used at resolution level: 2, 2, 2, 5</li> <li>Pair-wise registration initialisation</li> <li>Group-wise registration</li> <li>Activity and attenuation DVF composed</li> </ul> </li> <li>MM:         <ul> <li>MM fit between each resolution level</li> <li>Weighted (counts) linear regression MM</li> <li>1D or 2D respiratory SS</li> </ul> </li> </ul>

Method	Data	Reconstruction	Motion Estimation
Details 4.5			
Section 4.5	XCAT:     12.0 mm, 10.0 mm, 8.0 mm, or 6.0 mm lesion     2D MR respiratory traces     Hysteresis     480 volumes     Intra-gate motion      Simulation:     3D Simulation     European (late dynamic) research count rate (41 KCPS)     550.0 ps TOF resolution     Pseudo scatter (Gaussian smoothing 128.0 mm FWHM, 29 KCPS) and random events (57 KCPS)      Gating for motion estimation:     2D PCA respiratory SS     4 radial pseudo-phase gates      Gating for motion compensated reconstruction:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC OSEM initialisation (2 iterations, 24 subsets)</li> <li>MLACF (3 activity iterations, 24 subsets, 1 attenuation iterations)</li> <li>ACF initialised with μ-Map</li> <li>Activity volume scale set to OSEM reconstruction and</li> <li>ACF scale set to μ-Map between iterations</li> <li>Re-interpolation of outliers between iterations</li> <li>Smoothing of endplanes between iterations</li> <li>Median smoothing (kernel size 3) between iterations</li> <li>Sinogram mashing in attenuation update (1, 7, 2, 3 mashing factor respectively)</li> <li>Filled with NaNs to size of μ-Map</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	Registration: Diffeomorphic (velocity field parametrised, symmetric)  NMI objective function Tuned CPG spacing (13 times voxel size) and bending energy (1e-1 attenuation)  Multi-resolution registration (4 levels) Iterations used at resolution level: 2, 2, 2, 5 Pair-wise registration initialisation Group-wise registration Activity and attenuation DVF composed  MM: MM fit between each resolution level Weighted (counts) linear regression MM  2D respiratory SS

The exact same method is used here as was used in section 4.4, except for the difference in input data, as was discussed in section 4.5.1. Therefore, this section will act as a final full high level overview of the method used in order to clarify in one section the exact procedure used.

• XCAT is used to generate input volumes for the method. The volumes use the geometry of the GE Discovery 710. The FOV includes the top and base of the lung, including some of the liver. Either a 12.0 mm, 10.0 mm, 8.0 mm, or a 6.0 mm

```
Algorithm 1: Re-interpolate Outliers
  Data: outlierVolume, iqrMultiplier
  Result: outlierVolume
1 interQuartileRange = IQR of outlierVolume
3 for currentValue in outlierVolume do
      if currentValue >median of outlierVolume +
        interQuartileRange × iqrMultiplier then
                        2.0
          currentValue in outlierVolume = NaN
 5
 6
      if currentValue < median of outlierVolume -
        interQuartileRange × iqrMultiplier then
          currentValue in outlierVolume = NaN
 8
9 end
10
11 whileBool = true
  while while Bool do
13
      for currentValue in outlierVolume do
14
          if currentValue is NaN then
15
              if values surrounding currentValue is not NaN then
16
                 currentValue in outlierVolume = linear interpolation of values
17
                   surrounding currentValue
      end
18
19
      if all currentValue in outlierVolume is not NaN then
20
          whileBool = false
21
22 end
```

spherical lesion is placed into the base of the right lung, close to the liver (so as to experience maximum respiratory motion and also so that the liver makes the registration more difficult). A 2D respiratory trace is used to drive the motion of the XCAT, derived from MR navigator data. Because two signals are used then there is hysteresis present in the data. 480 volumes are produced over 240 s (sampling at  $2.0 \, \text{Hz}$ ).  $\mu$ -Maps were also generated at the position of all 480 volumes, these were used to add attenuation during the simulation step. The  $\mu$ -Map with the maximum respiratory value from the MR respiratory trace was selected to be used as the misaligned  $\mu$ -Map for the rest of the method.

## **Algorithm 2:** Smooth Endplanes

```
Data: smoothEndplanesVolume
  Result: outputVolume
1 outputVolume = smoothEndplanesVolume
3 smoothEndplanesVolume = Re-interpolate Outliers of smoothEndplanesVolume
    with iqrMultiplier = 1.0
5 smoothEndplanesVolume = max filter kernel size 3 of smoothEndplanesVolume
7 \ rampFilter = 1.0, 0.8, 0.6, 0.4, 0.2
s inverseRampFilter = 0.0, 0.2, 0.4, 0.6, 0.8
10 for sliceIndex in 0, 1, 2, 3, 4 do
      smoothEndplanesSlice = smoothEndplanesVolume axial slice at sliceIndex
11
        \times rampFilter at sliceIndex
      outputVolumeSlice = outputVolume axial slice at sliceIndex \times
12
       inverseRampFilter at sliceIndex
13
      outputVolume axial slice at sliceIndex = smoothEndplanesSlice +
14
       outputVolumeSlice
15
      endSliceIndex = end slice index of outputVolume - sliceIndex
16
17
      smoothEndplanesSlice = smoothEndplanesVolume axial slice at
18
       endSliceIndex \times rampFilter at sliceIndex
      outputVolumeSlice = outputVolume axial slice at endSliceIndex \times
19
       inverseRampFilter at sliceIndex
20
      outputVolume axial slice at endSliceIndex = smoothEndplanesSlice +
21
       outputVolumeSlice
22 end
```

# Algorithm 3: MLACF for Motion Correction

```
Data: sinogram, attenuationVolume
  Result: activityVolume
1 attenuationVolumeACF = forward project attenuationVolume
2 mlacfACF = attenuationVolumeACF
4 osemActivityVolume = AC OSEM on sinogram with
   attenuationCorrectionFactors = attenuationVolumeACF and iterations = 2 and
   subsets = 24
5 activityVolume = osemActivityVolume
7 for mlacfIterations in 0, 1, 2 do
      activityVolume = Re-interpolate Outliers of activityVolume with
8
       iqrMultiplier = 3.0
      activityVolume = Smooth Endplanes of activityVolume
9
      activityVolume = median filter kernel size 3 of activityVolume
10
      activityVolume = rescale activityVolume to scale of osemActivityVolume
11
12
      mlacfACF = attenuation update on mlacfACF with sinogram = sinogram and
13
       activityVolume = activityVolume
      mlacfACF = rescale mlacfACF to scale of attenuationVolumeACF
14
15
      ctivityVolume = AC OSEM on sinogram with attenuationCorrectionFactors
16
       = mlacfACF and iterations = 1 and subsets = 24
17 end
18
19 activityVolume = Re-interpolate Outliers of activityVolume with iqrMultiplier =
20 activityVolume = Smooth Endplanes of activityVolume
21 activityVolume = median filter kernel size 3 of activityVolume
22 activityVolume = rescale activityVolume to scale of osemActivityVolume
```

- STIR (Thielemans et al. 2012), (Efthimiou et al. 2019b) through SIRF (Ovtchinnikov et al. 2020) is used to perform a 3D TOF (550.0 ps) PET simulation. The count rate was established from the last two minutes of a dynamic scan (the data from chapter 5). 127 KCPS total counts were used, 41 KCPS true counts, 29 KCPS scatter counts, and 57 KCPS random counts. A full scatter simulation could not be performed due to software limitations so a forward projection of the smoothed (Gaussian 128.0 mm FWHM)  $\mu$ -Map was used instead.
- A 2D SS was acquired using PCA to be used during gating and MM fitting. Two

#### **Algorithm 4:** Motion Modelling

```
Data: activityVolumes, surrogateSignals, attenuationVolume
  Result: motionModel, motionCorrectedActivityVolume, attenuationVolumes
1 for volume in activityVolumes do
      currentNMI = NMI between volume and attenuationVolume
2
3
      if currentNMI < currentMinNMI then
4
          currentMinNMI = currentNMI
5
          motionCorrectedActivityVolume = volume
 6
7 end
9 for level in 0, 0, 0, 1, 1, 2, 2, 3, 3, 3, 3, 3 do
      for volume in activityVolumes do
10
          activityDVFs = register volume to motionCorrectedActivityVolume with
11
           and cpgVoxels = 13 and bendingEnergyWeight = 0.0 and
           registrationLevel = level and initialisation = activityDVFs
      end
12
13
      motionModel = weighted linear regression on activityDVFs and
14
       surrogateSignals with weighting = counts in activityVolumes
      activityDVFs = motionModel on surrogateSignals
15
16
      motionCorrectedActivityVolume = resample activityVolumes with
17
       activityDVFs
18
      attenuationDVF = register motionCorrectedActivityVolume to
19
       attenuation Volume and cpgVoxels = 13 and bendingEnergyWeight = 1e-1
       and registrationLevel = level and initialisation = attenuationDVF
20
      for dvf in activityDVFs do
21
         dvf = \text{compose } dvf \text{ with } attenuationDVF
22
      end
23
24
      motionModel = weighted linear regression on activityDVFs and
25
       surrogateSignals with weighting = counts in activityVolumes
      activityDVFs = motionModel on surrogateSignals
26
27
      motionCorrectedActivityVolume = resample activityVolumes with
28
       activityDVFs
29 end
31 for attenuationDVF in inverse of activityDVFs do
      attenuationVolumes append resample attenuationVolume with
       attenuationDVF
33 end
```

gating schemes were used, one for the motion correction and one for the motion compensated reconstruction:

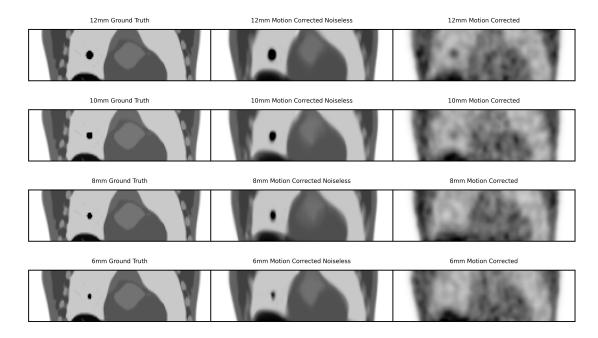
- For the motion correction, four radial pseudo-phase gates were used, as discussed in section 4.3.2.2.
- For the motion compensated reconstruction, 30 respiratory bins were used (10 displacement and three gradient bins). This number was selected due to the resolution of the scanner.
- A modified version of MLACF was used as the initial reconstruction method for the volumes to be motion corrected. First the data was reconstructed using AC OSEM (two iterations, 24 subsets), using the misaligned μ-Map from end inhalation. The result of the AC OSEM reconstruction was input to MLACF, after every activity update the new activity volume was rescaled to the AC OSEM volume. MLACF used three iterations, 24 subsets, and one attenuation update for each activity update. The misaligned μ-Map was used to initialise the ACF and the scale of the ACF was rescaled to the misaligned μ-Map ACF following every attenuation update. Sinogram mashing was used in the attenuation update with a mashing factor of one, seven, two, and three in the radial positions, angles, transaxial plane, and TOF dimensions respectively. Three aditional processing steps were performed on the activity volume after every update, these were added specifically to ensure the volume was smooth for motion correction and were as follows:
  - Firstly, values outside three times the IQR are re-interpolated, using a linear interpolation, as can be seen in algorithm 1.
  - Secondly, the endplanes of the volume were extremely noise and as such required extensive smoothing if they are to register without artifacts. To achieve this the current activity volume was copied into another volume, known as the endplanes volume. This endplanes volume has outlying values re-interpolated where they fall outside the IQR, this can be seen in algorithm 1. A max filter is then applied with a kernel size of three. This endplanes volume is then summed with the activity volume using a ramp filter. The ramp filter starts at

one and linearly reduces to zero over five values, the same ramp filter can be used at the top and the bottom of the volume if it is flipped so that the value of one is always along the edge of the volume. If the endplanes volume is multiplied by the ramp filter and the activity volume is multiplied by the inverse of the ramp filter then when they are summed their scale remains consistent. This process can be seen in algorithm 2.

- Finally, a median smoothing with a kernel size of three is applied.

Pseudocode for the modified version of MLACF can be seen in algorithm 3.

- Between the reconstruction and motion correction stages the activity volumes have their axial size expanded to the size of the μ-Map by appending whole slices filled with NaNs. Often the FOV of the μ-Map is greater than any one bed position of the PET. Furthermore, by appending with NaNs we allow the activity volumes to be registered outside of the FOV of the original acquisition, this could be important if the activity volumes and μ-Map are extremely misaligned. NaNs are used as they will not affect the objective function calculation of the registration.
- For registration NiftyReg (Modat et al. 2010) was used with a symmetric velocity field parametrised diffeomorphic implementation. The initial reference frame for registration was selected as the one with the highest NMI to the misaligned μ-Map. Registration of the activity volumes came first (in a group-wise manor) before the reference volume was resampled. This new reference volume was then registered to the misaligned μ-Map. DVFs were composed from the two registrations. For registration, four resolution levels were used, a weighted linear regression (weighted on the counts in each activity volume) MM was fit after each iteration, each iteration was initialised with the DVFs from the previous iteration. Three iterations of the lowest resolution level were used, two of the next, two of the next, and five of the highest resolution level (three of the lowest level were used as the first one substituted a pair-wise initialisation). A NMI objective function was used for all iterations. The CPG was spaced 13 voxels apart and bending energy was used for the attenuation registration with a weighting of 1e−1. Pseudocode for the motion correction method can be seen in algorithm 4.

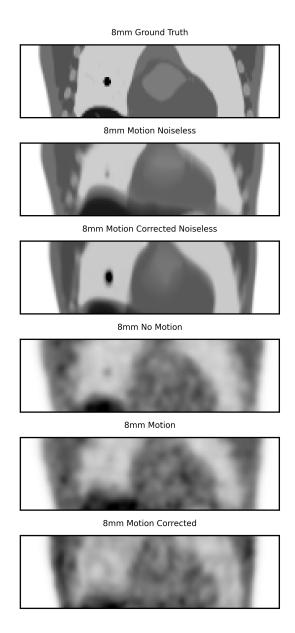


**Figure 4.20:** A visual analysis which shows the results of performing the method on XCAT data with a 12 mm, 10 mm, 8 mm, or 6 mm lesion placed into the base of the right lung (the left of this image). On the left of each row is the Ground Truth image (without motion or reconstruction artifacts), in the centre of each row are the results of applying a motion correction fit on reconstructed data to noiseless data, and on the right of each row is the result of applying the motion correction to data which has been reconstructed. The top row is where a 12 mm lesion was used, the second row down is where a 10 mm lesion was used, the third row down is where a 8 mm lesion was used, and the bottom row is where a 6 mm lesion was used. In this case the input to the motion correction was the 4 Phase MLACF reconstructions and the motion correction has been applied to the 30 bin OSEM reconstructed images. Colour map ranges are consistent for all images.

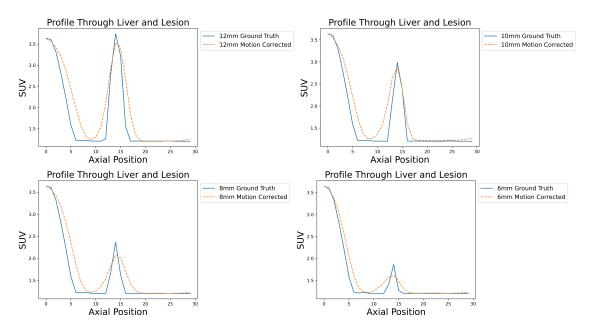
• A final motion compensated reconstruction is performed using the method whereby the activity volumes are reconstructed, deformed using the DVF, and then summed together. Here the data gated into 30 respiratory bins were used (10 displacement and three gradient bins) is used. AC OSEM with two iterations and 24 subsets is performed using the deformed μ-Map from the registration step. The usual Gaussian post-smoothing (with a 6.4 mm FWHM) is used transverse and a 'normal' Z-filter is used axially (a convolution with a kernel containing 1.0, 4.0, 1.0).

#### 4.5.3 Results

Figure 4.20 shows the results from data using a different lesion size. From the second column it can be seen clearly that the method is capable of correctly estimating the motion



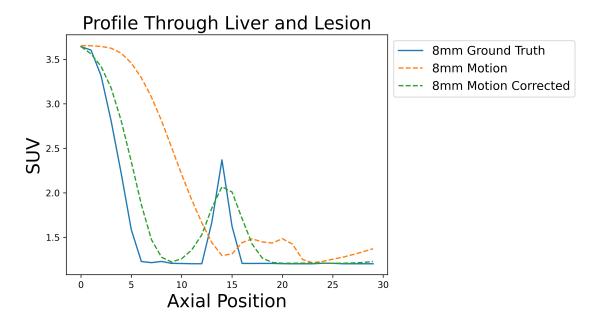
**Figure 4.21:** A visual analysis which shows the results of performing the method on XCAT data with a 8 mm lesion placed into the base of the right lung (the left of this image). At the top of this image is the Ground Truth image (without motion or reconstruction artifacts), second from the top of this image is where no motion correction has been applied to noiseless data, third from the top of this image is the results of applying a motion correction fit on reconstructed data to noiseless data, fourth from the top of this image is the result of reconstructing an image without motion which has been reconstructed, fifth from the top of this image is the result of reconstructing an image with motion which has been reconstructed, and at the bottom of this image is the result of applying the motion correction to data which has been reconstructed. In this case the input to the motion correction was the 4 Phase MLACF reconstructions and the motion correction has been applied to the 30 bin OSEM reconstructed images. Colour map ranges are consistent for all images.



**Figure 4.22:** A profile through the 12 mm, 10 mm, 8 mm, or 6 mm lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. The blue line shows the profile from the Ground Truth image with a 12 mm, 10 mm, 8 mm, or 6 mm lesion and the dashed orange line shows the result of applying a motion correction on reconstructed data with a 12 mm, 10 mm, 8 mm, or 6 mm lesion. The top left of this image is where a 12 mm lesion was used, the top right of this image is where a 10 mm lesion was used, the bottom left of this image is where a 8 mm lesion was used, and the bottom right of this image is where a 6 mm lesion was used. In this case the input to the motion correction was the 4 Phase MLACF reconstruction and the motion correction has been applied to the 30 bin noiseless images.

**Table 4.6:** Comparison of  $SUV_{max}$  and  $SUV_{peak}$  over the lesion, where different lesion sizes were used for the XCAT simulation. The different lesion sizes used were 12 mm, 10 mm, 8 mm, and 6 mm. In the case of the 8 mm lesion the results without motion correction are also reported.

SUV	Max	Peak
12 mm <b>Ground Truth</b>	9.08	8.24
12 mm <b>Motion Corrected</b>	7.85	6.25
10 mm <b>Ground Truth</b>	9.08	7.61
10 mm <b>Motion Corrected</b>	7.73	5.86
8 mm Ground Truth	9.08	5.18
8 mm <b>Motion</b>	1.97	1.29
8 mm <b>Motion Corrected</b>	5.48	3.87
6 mm Ground Truth	9.08	3.80
6 mm Motion Corrected	3.65	2.14



**Figure 4.23:** A profile through the 8 mm lesion, in the SI direction, summed over a window in the AP and LM directions, with median smoothing. The blue line shows the profile from the Ground Truth image with a 8 mm lesion, the dashed orange line shows the result of not applying motion correction, and the dashed green line shows the result of applying a motion correction on reconstructed data with a 8 mm lesion. In this case the input to the motion correction was the 4 Phase MLACF reconstruction and the motion correction has been applied to the 30 bin noiseless images.

of the liver and the lesion even when the size of the lesion is reduced. Obviously, the motion correction appears to struggle more in the case where the lesion is 6 mm than it does where it is 12 mm, but the motion correction does not fail and the position of the liver and the lesion matches closely what is shown in the ground truth column. The third column generally corroborates the interpretation of the second column. However, in the third column it becomes difficult to make out the lesion in the 8 mm case and the 6 mm case (it is clearly possible to identify the lesion in the 12 mm case and the 10 mm case). This is potentially caused by the noise present in the reconstruction as well as partial volume effects due to the scanner resolution. It must be stressed that, for the GE Discovery 710, an 8 mm diameter spherical lesion is just less than two voxels across.

A further visual analysis of the 8 mm data can be seen in figure 4.21. The top image shows the ground truth image. The second image shows the 480 input volumes to the method for the 8 mm averaged together, illustrating the extent of the motion in the input data. The third image shows the final MM from the motion estimation method applied

to data without artifacts from simulation and reconstruction (the data has however been gated). The fourth image shows the result of the reconstruction of the noisy data if no motion is present, giving an indication of the best possible motion correction. The lesion can just be seen in this image, however it has much less contrast than in the ground truth image, as expected. The fifth image shows the result of the reconstruction of the noisy data without applying any motion correction. The lesion cannot be seen at all in this image. The final image shows the output of the method. The lesion can just be made out in the same position as in the first image.

A profile over the lesion for each lesion size can be seen in figure 4.22. It should be highlighted that for each axial position values in a region have been averaged together, to eliminate noise. This region was kept consistent for each sub-figure and represented a 12 mm bounding box around the lesion. This explains why as the size of the lesion decreases so does the size of the peak in the profile. For each lesion size both the ground truth (in blue) and the results of the method are shown (in broken orange). All profiles appear to match the ground truth well, the peak of all profiles coincides with the peak of the ground truth, as does the position of the liver. The height of the peak is consistently lower for the motion correction, but not significantly. The difference in the height of the peak appears increases slightly as the size of the lesion decreases.

As for the visual analysis above so too is this analysis performed in more detail for the 8 mm lesion, in figure 4.23. Here, not only is the ground truth (in blue) and the result of the method (in broken green) shown but also the result of not applying motion correction is shown (in broken orange). Here it can be seen that the position of the lesion without motion correction is entirely different from the position of the lesion in the ground truth (and motion corrected example). The profile of the lesion where no motion correction has occurred is almost fully outside where the lesion appears in the ground truth case. Without motion correction, the lesion almost overlaps with the edge of the liver of the volume, highlighting the difficulty of motion correcting the data with the lesion in this position, of this size, and with this extent of motion.

The  $SUV_{max}$  and  $SUV_{peak}$  for a region containing the lesion is shown for all results (from above) in table 4.6. The  $SUV_{max}$  of the ground truth compared to the  $SUV_{max}$  of the motion corrected results represents a 1.16, a 1.17, a 1.66, and a 2.49 times reduction

for the data containing the 12 mm, the 10 mm, the 8 mm, and the 6 mm lesion respectively. To put this into context, the  $SUV_{max}$  of the ground truth compared to the  $SUV_{max}$  of the non motion corrected result represents a 4.61 times reduction for the data containing the the 8 mm lesion. Similar observations can be seen in the  $SUV_{peak}$  values. There is a noticeable drop off in the SUV values for lesion of size 8 mm and below. This is likely caused by less accurate motion correction as well as partial volume effects. However, the SUV results without motion correction for the 8 mm lesion are much lower than with motion correction.

#### 4.5.4 Discussion

Results from visual analysis, comparison of profiles, and SUV on realistic and challenging data, both with and without noise, and with and without motion correction for lesions of size 12 mm, 10 mm, 8 mm, and 6 mm diameter shows that the method presented both here and in 4.4 both works well and is robust.

For a further discussion and conclusion of the method presented in this section please see section 4.4.4 and section 4.6.

# 4.6 Discussion and Conclusion

The results of the previous chapter demonstrated proof-of-concept performance. However, it relied on idealised assumptions that limited its clinical relevance. The work presented in this chapter addresses several of those limitations by introducing more realistic simulation conditions, closely matching data late time frame data from a research study performed at our institution (see the next chapter), including variable respiratory patterns, lower TOF resolution, pseudo-random and scatter components, lower counts than in routine clinical practice and different lesion sizes.

Notably, this chapter also marks a significant step forward in integrating attenuation correction into the motion correction workflow. The mismatch between emission and attenuation data, previously a major limitation, was mitigated through deformation of the  $\mu$ -Map using DVFs estimated by the MM. This improvement brings the framework closer to applicability in real clinical scenarios.

These developments represent a meaningful step toward applying the motion correction framework to real patient data. The simulation and reconstruction pipelines have been

refined to support this transition, with adaptations for increased realism and complexity.

Several technical hurdles emerged during preliminary attempts to apply the method to these challenging simulation data. The transition from 2D to 3D simulation, lower counts and reduced TOF resolution introduced new complexities. Additionally, the reconstruction pipeline had not previously been validated with realistic inputs such as randoms, scatter, and normalisation sinograms, which further exposed limitations in robustness. These challenges underscore the need for continued refinement of the implementation to ensure stable and accurate performance under clinical conditions.

The modifications introduced in Section 4.4 enabled the method to run end-to-end using real patient data acquired on a GE Discovery 710 scanner. The primary remaining barrier to full patient data evaluation in the current implementation is the need to divide the acquisition into 0.5 Hz sinograms. While this straightforward in principle, it raises practical issues as generating such a large number of sinograms, particularly in dynamic studies, requires substantial storage and processing capacity. However, this could be addressed by using the gating capabilities available in STIR, which have not yet been ported to SIRF, the software used in our current implementation, at the time of writing). Nevertheless, the framework has approached readiness for clinical testing.

The XCAT simulations used in this chapter uniformly use SSs derived from MR to drive the respiratory model. This addresses the limitation mentioned previously where there was no breath to breath variability present in the data.

Compared to the previous chapter, the lesion was made much smaller. The lesion was small enough that from the minimum displacement to the maximum displacement, due to respiration, the lesion would not overlap with itself. Furthermore, the lesion was placed close enough to the diaphragm that from the minimum displacement to the maximum displacement, due to respiration, the lesion is fully within the area where the diaphragm was at the opposite respiratory state.

As was discussed in section 4.4.2.1, the method to simulate the PET acquisition is now significantly more realistic. A relatively low number of counts are used, at the level of a research scan count rate (41 KCPS true counts). A pseudo-random and scatter simulation was used to emulate the randoms and scatter profile seen in patient data. This is not a true random and scatter simulation, due to software limitations when using

TOF. However, care was taken to make the simulations as realistic as possible given the constraints.

Again, as was discussed in section 4.4.2.1, the TOF resolution for the PET acquisition simulation was reduced to that of the GE Discovery 710. This did have a noticeable effect on the quality of the reconstructed and motion corrected volumes. However, the method was improved so that similar performance was seen as before the resolution was reduced.

The registrations used universally in this chapter employed NMI as an objective function. This even included the pair-wise registration to the reference volume in section 4.4.2.3. In the current application, NMI appears to be superior to SSD or MSE, especially when performing registration between two modalities (like PET and CT). If NMI is used for all registrations, including where the same modality is used, it makes the hyperparameter tuning easier due to the parameters being similar for both types of registration.

Furthermore, although the work in this chapter continued to fit the registration and MM separately, it never appeared to be the limiting factor for the quality of the motion correction.

The MMs used in this chapter continued to be fit as a weighted linear regression. It would be possible to consider non-linear regression. However, this would only be possible when more bins are used for the fitting, which would need careful treatment of the increased noise.

The work presented in this chapter demonstrated that motion correction of PET/CT data with a single CT using a fitted motion model is feasible under increasingly complex imaging conditions, even without having an external measurement of the respiration available, making the method attractive for clinical translation. However, these methods were tested within the context of static PET acquisitions. In dynamic PET, where SNR ratios are lower and the underlying activity distribution changes over time, a new challenge arises: how to extract a reliable respiratory SS directly from the PET data itself. The next chapter addresses this by introducing a DD approach for SS extraction, enabling the extension of motion correction techniques to dynamic imaging protocols in the future.

Furthermore, please see appendix B for work which was inspired by the methods here (Whitehead et al. 2022a). In appendix B (Whitehead et al. 2022a) a NN DIP (Uly-

anov et al. 2020) based method for the de-noising of dynamic PET frames is presented. The preprocessing of the data during the MLACF reconstruction used in this section is complex, what is required is a smooth and high contrast input image for motion correction. Therefore, this method was developed and tested as a preprocessing step for dynamic PET to be used with kinetic modelling. It stands to reason that it may be possible to replace the preprocessing of the data for MLACF, and motion correction, with this method. This is especially true if this motion correction method is used on dynamic PET, where the de-noising method has already been evaluated.

Further limitations and future developments are discussed in chapter 6.

# **Chapter 5**

# Data Driven Surrogate Signal Extraction for Dynamic PET

# 5.1 Overview

This chapter presents work on the extraction of respiratory surrogate signals from dynamic PET data. In previous chapters, respiratory motion correction was developed for static PET data, for instance, assuming that the tracer metabolism is stabilised. The focus was on model-based deformation fields and gated reconstructions. However, dynamic PET acquisitions introduce additional challenges, particularly due to time-varying tracer kinetics and lower signal levels in early frames. A particular example is the extraction of a SS for the motion from the PET data. In previous chapters, PCA was used, which was based on the assumption that motion is the main cause for change in the data, which is no longer true for dynamic PET.

To address this limitation, this chapter investigates DD methods for extracting surrogate signals for the respiratory motion directly from the PET data itself, a critical step towards enabling motion correction in dynamic imaging without reliance on external hardware. In particular, this section will introduce novel methods using PCA, with several approaches, to improve its robustness to variation caused by radiotracer kinetics. It also includes a comparison to another methods from the literature which uses SAM. The chapter is based on (Whitehead et al. 2022b), (Whitehead et al. 2024).

# 5.2 Introduction

Methods of motion correction are commonly based on a respiratory trace. To acquire these respiratory traces, an external device, like the RPM, or a DD method based directly on the PET data can be used. DD methods have the advantage that they are non-invasive, requite no set-up time during the acquisition and can be performed post-acquisition. However, current DD methods have the disadvantage that they are adversely affected by the tracer kinetics of a dynamic acquisition. For a full description of the background of respiratory SS extraction (especially in static PET) please see section 2.6.3.

Current DD methods are adversely affected by the radiotracer kinetics of a dynamic acquisition, where the tracer is injected after the beginning of the scan. As an example, methods that use dimensionality reduction (such as PCA) are hampered by the fact that at the start of the scan, rapid redistribution of the radiotracer (rather than the respiratory motion) causes more variance in the data. Previously, work was performed to extend the SAM method to be robust to radiotracer kinetics. This work proposed the use of Short Time Fourier Transform (STFT) to generate masks for SAM (rather than a static mask for all time intervals), this was called Kinetic Respiratory Gating (KRG) (Schleyer et al. 2014). STFT operates by splitting the data into windows, and applying a FFT on them independently. However, this method was unable to extract a usable signal at very early time intervals (after tracer injection).

The aim of the current chapter is to propose several adaptions of the PCA method, with which it can be used with dynamic data, and compare their performance with a method based on KRG. The methods explored in this section include; the use of a moving window, re-use of the PCs from a later time interval to estimate the SS for earlier time points, and automatic scoring, selection, and combination of multiple PCs, akin to SRF.

Firstly, in section 5.3, the data (including train and test splits) will be introduced, before moving on to the methods which are being proposed or compared. Next, in section 5.4, a more thorough description of the data, including how it is prepared, and the evaluation methods used to compare the methods are presented. In this section we also introduce potential post-processing techniques, which can be performed to SSs, to improve results generally. This is followed by section 5.5 which presents a comparison

of the methods defined in the previous section. The advantages and disadvantages of the methodology are discussed in section 5.6. Finally, in section 5.7, the arguments put forth are drawn together, before briefly pointing out the potential future directions for the work.

# 5.3 Methods

Here, we briefly describe the data acquisition, before describing methods that are either simple modifications of the conventional methods (based on KRG) or use a novel method to score, select, and combine signals. These methods can be used with SAM, but for simplicity, we will refer specifically to PCA.

The following subsections will address the method with respect to the diagram seen in figure C.1.

# 5.3.1 Data Acquisition and Train/Test Split

Data used was acquired from a research study with patients suffering from Idiopathic Pulmonary Fibrosis (IPF) (Emond et al. 2020). 22 dynamic [<sup>18</sup>F]-FDG acquisitions, with a FOV covering the upper lung and heart, were acquired on a GE Discovery 710 in list-mode. Data used in this study was from the first 14 min with the acquisition starting roughly 20 s before injection of the radiotracer. An external SS was acquired in parallel using an RPM (Oh et al. 2019).

These 22 acquisitions comprised of two scans each of 11 subjects. One scan was performed before and one scan was performed some time after an intervention was performed. In this case, the intervention was the administration of an anticoagulant drug. The first acquisition of each patient is called the baseline scan and the second acquisition is called the post treatment scan (Emond et al. 2020). Of these 22 acquisitions only 10 were suitable to be used as either part of the training or testing process. Five of the acquisitions could not be used as the list-mode itself could not be loaded by the software. The remaining acquisitions could not be used for either training or testing due to issues with the acquisition of the RPM. For instance, the acquired RPM could not be synced with the list-mode. These data without RPM could be used with the subsequent methods, just not evaluated, as can be seen in figure 5.13.

To form a train and test split the data was randomly split into three training data points and eight testing data points. More than one training data point was selected to attempt to prevent over-fitting on a single data point. Data points from both baseline and post treatment acquisitions were added to prevent over-fitting on type of acquisition. Train data points were selected such that the same patient did not appear in both the train and test dataset, to attempt to mitigate data leakage. A validation dataset was not utilised due to the low number of data points.

For selection of hyper-parameters, the method was applied to the train dataset, and the correlation coefficients with RPM were computed. The mean of the three resulting correlation coefficients was maximised while varying a hyper-parameter.

### 5.3.2 Conventional PCA

# **Algorithm 5:** Conventional Score

**Data:** timeSeriesSinograms, PC, respiratoryFrequencyWindow

**Result:** respiratoryScore

1 for sinogram in timeSeriesSinograms do

2 | respiratorySignal append  $sinogram \times PC$ 

3 end

4

5 *PSD* = absolute of FFT on *respiratorySignal* 

6

7 respiratoryScore = mean value of PSD within respiratoryFrequencyWindow

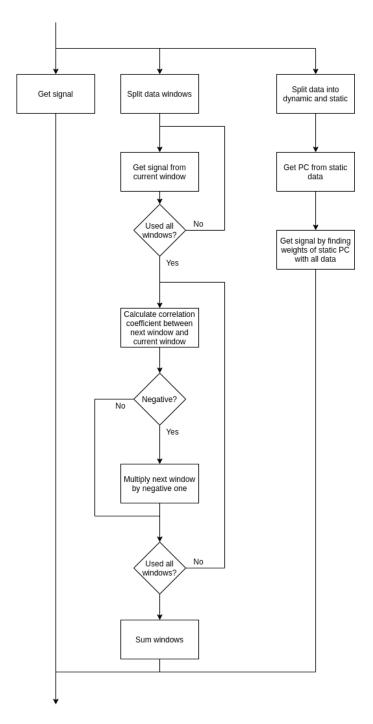
In this method, the "standard" PCA is applied to the entire data set in one go as it would be if the data were from a static acquisition.

As discussed in 5.1, the conventional method includes a selection mechanism using a score that maximises a signal with the appropriate respiratory features (for instance, a frequency matching the common human breathing). In this chapter, the score will be computed according to algorithm 5 (Bertolli et al. 2017b).

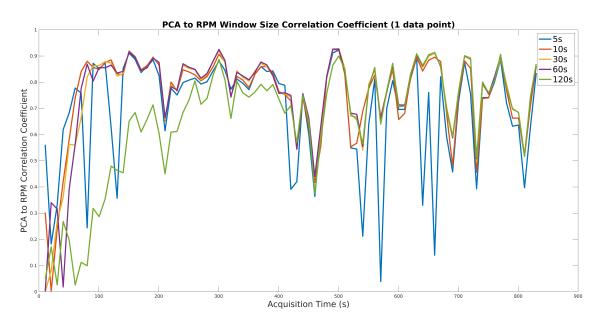
The generic equation for calculating the weights (or signal) from the PC and data is

$$W = PC \times D \tag{5.1}$$

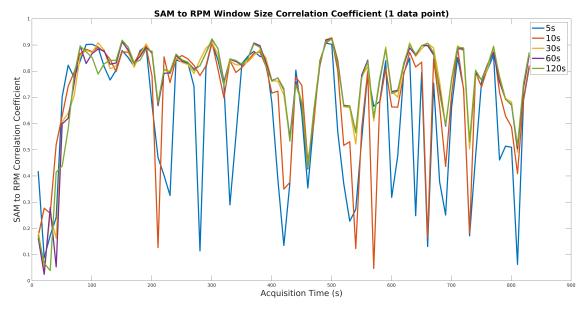
where in equation 5.1 PC represents the PC (which in this case is the shape of one sinogram) and S represents a time series of sinograms.  $\times$  denotes element-wise multiplication of the arrays (in this case multiplication of the one PC by each sinogram in the time series S), followed by summing. In fact, a similar equation is used by SAM, where a 'signed



**Figure 5.1:** This diagram shows the potential workflows for the different methods examined here. The path on the left shows only the signal (or signals) being extracted without any further processing (until the next steps). The path in the middle shows the process for extracting a signal using windows. The path on the right shows how one would extract a signal using a PC from a point in the data where it could be considered to be static.



**Figure 5.2:** A plot showing the moving window size optimisation for the PCA method. For different fixed window sizes, the correlation of the extracted signal to the RPM is shown for the windows sliding over the whole acquisition (taken for the first acquisition of patient one). Note that 0.5 s time frames were used.



**Figure 5.3:** A plot showing the moving window size optimisation for the SAM method. For different fixed window sizes, the correlation of the extracted signal to the RPM is shown for the windows sliding over the whole acquisition (taken for the first acquisition of patient one). Note that 0.5 s time frames were used.

mask' is multiplied with the data and summed.

# **Algorithm 6:** Moving Window Method

```
Data: timeSeriesSinograms, windowSizes
  Result: respiratorySignal
1 index = 0
2 whileBool = true
  while while Bool do
      if index >length of timeSeriesSinograms then
5
          index = length of timeSeriesSinograms - windowSize
 6
          whileBool = false
 7
8
      set windowSize to value at index of windowSizes
10
      windowSignal = fill with NaNs to index
11
      windowSignal append compute PC weight with PCA for data between index
12
       and index + windowSize
      windowSignal append NaNs to length of timeSeriesSinograms
13
14
      if windowSignal correlation with last signal in signals <0.0 then
15
          windowSignal = windowSignal \times -1.0
16
17
      signals append windowSignal
18
19
      index = index + \frac{windowSize}{2}
20
21
22 end
23
24 respiratorySignal = mean of signals ignoring NaNs
```

# **5.3.3** Moving Window Method

As shown in the algorithm 6, the data is split into a series of windows, where each subsequent window overlaps with the previous window by half its length. The motivation for attempting the Moving Window method is to increase the relative importance of motion vs kinetics. This is achieved through small windows being used at early time intervals, where the radiotracer kinetics are at their most severe, and longer windows can be used at later time intervals to reduce noise. If SAM is used rather than PCA, then the method approximates KRG (Schleyer et al. 2014).

The size of each window was predetermined on a small training data set. The PCA method was first run with a number of fixed window sizes. The window size which

gave the best signal (defined as the highest correlation coefficient with the RPM within each window) at each time interval was selected and recorded. Example results for this optimisation method for the moving window size can be seen in figure 5.2 and figure 5.3 for the PCA and SAM methods respectively.

For this method PCA (or SAM) is applied independently on each window, and the results are averaged together after sign correction. The latter is required as the sign of the signal from each window is arbitrary. The overlapping allows for a common sign to be found, by comparing the correlation coefficient of neighbouring windows, and flipping windows where the correlation coefficient have opposite sign. Other methods for sign correction are possible, for example see (Bertolli et al. 2017b), (Feng et al. 2018), as well as the sign choice in the algorithm 10.

#### **5.3.4** Late Time Interval Method

**Algorithm 7:** Late Time Interval Method

**Data:** timeSeriesSinograms, lateTimeIntervalCutoff

**Result:** respiratorySignal

1 lateTimeIntervalSeriesSinograms = split timeSeriesSinograms from lateTimeIntervalCutoff to end

2 lateTimeIntervalPC = PC from PCA for lateTimeIntervalSeriesSinograms

3

- 4 for sinogram in timeSeriesSinograms do
- s respiratorySignal append  $sinogram \times lateTimeIntervalPC$
- 6 end

Here, a PC from a late time interval is taken, and used with all data from all time intervals, which to our knowledge is a novel approach. The motivation behind this method was the hypothesis that; because the respiratory motion should be semi-consistent throughout the acquisition, then if a PC is capturing the respiratory motion at late time intervals, it should do the same at early time intervals as well. Indeed, we observed that PCs from late time interval data did not vary significantly when different windows were selected. However, this was not true for early time interval data.

The Late Time Interval PC method, as seen in the algorithm 7, splits the data into two channels, one which only contains later time interval data, where the radiotracer kinetics have diminished, and one which contains all the data. PCA is applied to the later

time interval data only. The PC from the later time interval data can then be taken and multiplied by the channel containing all of the data, to give the weights contributing to that PC for all time points.

The cut-off between early and later time interval data was determined on training data, by varying the cut-off point, and maximising the correlation coefficient between the output and RPM signal, for the first 120 s interval (between 20 s and 140 s). The cut-off determined here was 62 % or approximately 520 s from the start of the acquisition. It was noted however that later cut-offs gave similar results.

A flowchart of the above can be seen in figure 5.1.

# 5.3.5 Score, Select, and Combine Method

In this section, we describe a novel method based on a combination of previous work. The use of this method was inspired by the observation that signals with sizeable power in the frequency window of respiratory motion could be seen outside of the first few PCs. Additionally, a significant number of these had far less of a frequency contribution in a frequency window of the radiotracer kinetics. However, the information contained in these PCs is ignored if only one PC is used, as in (Thielemans et al. 2011), and (Bertolli 2018). This could lead to a reduced SNR. The method therefore uses a 'respiratory score', and orders and combines PCs to maximise this score. A flowchart of this can be seen in figure 5.4.

#### 5.3.5.1

Score and Select

In this step, the signals from each PC are scored and sorted. Signals with a too low score are rejected. Pseudo-code is in algorithm 8.

We developed several ways to calculate a score.

#### 5.3.5.1.1 Frequency Based

As seen in the algorithm 9, Power Spectral Density (PSD) analysis (Thielemans et al. 2011) used the PSD of the weights for each PC, to select for the PC with the highest contribution in the respiratory window. We extended this method to account for kinetic information. In our current implementation, these PSDs contain the frequency contribution

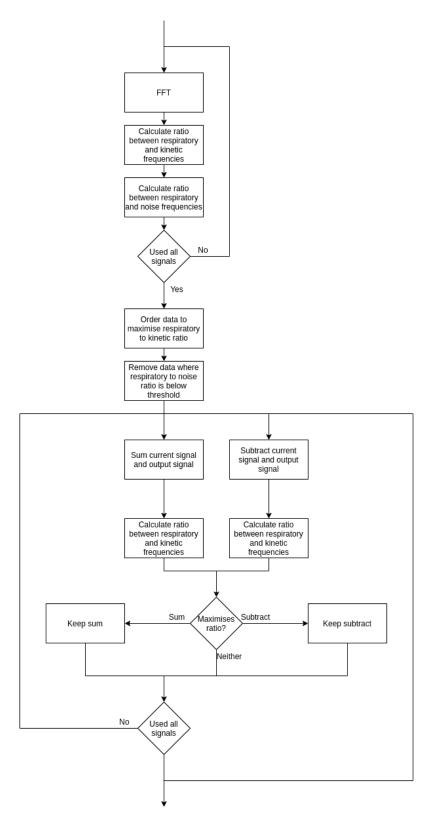


Figure 5.4: A diagram showing how PCs can be selected and combined.

#### **Algorithm 8:** Score and Select PCs

```
Data: timeSeriesSinograms, PCs, scoreThreshold
  Result: PCs
1 for PC in PCs do
      respiratoryScore = get respiratory score from PC and timeSeriesSinograms
      if respiratoryScore >scoreThreshold then
4
          respiratoryScoreList append respiratoryScore
5
      else
6
         remove PC from PCs
7
      end
9 end
10
11 sort PCs by respiratoryScoreList
```

#### Algorithm 9: Frequency Score

```
Data: timeSeriesSinograms, PC, kineticFrequencyWindow,
         respiratoryFrequencyWindow, noiseFrequencyWindow
   Result: respiratoryScore
 1 for sinogram in timeSeriesSinograms do
      respiratorySignal append sinogram \times PC
3 end
5 PSD = absolute of FFT on respiratorySignal
7 kineticContribution = mean value of PSD within kineticFrequencyWindow
8 respiratoryContribution = mean value of PSD within
    respiratoryFrequencyWindow
9 noiseContribution = mean value of PSD within noiseFrequencyWindow
10
11 respiratoryKineticRatio = \frac{respiratoryContribution}{\cdot \cdot \cdot}
                                kineticContribution
12
13 respiratoryNoiseRatio = \frac{respiratoryContribution}{.}
                               noise Contribution
14
15 respiratoryScore = respiratoryKineticRatio \times respiratoryNoiseRatio
```

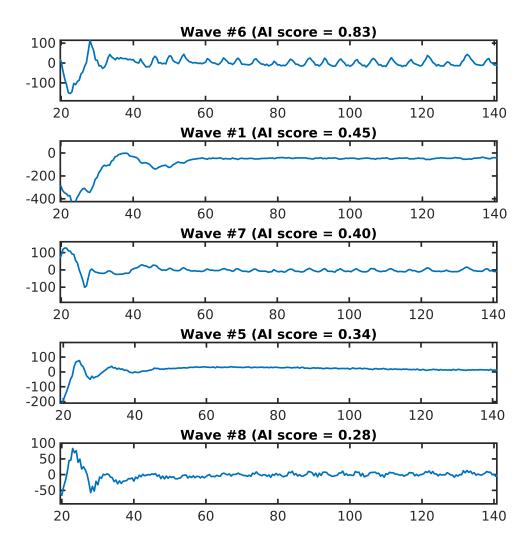
of each signal, between the frequencies of 0.0 Hz and 1.0 Hz (due to sampling the input data at 2.0 Hz and the Nyquist theorem (Whittaker 1915), (Nyquist 1928), (Shannon 1949)). Frequency windows representing the content of information related to radiotracer kinetics, respiratory motion, and noise are defined.

In an initial implementation they were defined as 0.0 Hz to 0.1 Hz, 0.1 Hz to 0.4 Hz, and above 0.4 Hz respectively (Bertolli et al. 2017b). However, it was found that the choice of respiratory window boundaries was limiting, it was both too wide (so as to encourage the mislabelling of noise), and not low enough (so as to fail on slow breathers). Thus, in the current implementation, the respiratory window is determined by first applying the Late Time Interval PC method, to acquire an initial estimate of the signal, and using this to estimate the window boundaries. A PSD of the initial estimate is acquired. The frequency which is at the mean value of the PSD is determined to be the centre of the window, and the boundaries are selected as being half the standard deviation of the PSD from this point. Half a standard deviation is used such that there is a full standard deviation between the upper and lower bounds of the window.

The contribution within each window is determined for each PC by finding the mean magnitude within the windows. Ratios are then calculated between the respiratory window and the kinetic window, and the respiratory window and the noise window, and a score determined by the product of these two values.

#### 5.3.5.1.2 Neural Network Based

A NN based scoring metric that was previously developed (Walker et al. 2020), and was tested here to remove complexity and increase robustness when compared to the frequency scoring method. The NN of (Walker et al. 2020) is a pre-trained model, designed to accept a signal as input and return a score between 0.0 and 1.0. Here a higher score indicates a more respiratory like signal. To achieve this, to avoid issues with signals of different lengths, features of the signal were extracted and used as input to the NN (rather than directly inputting the signal itself). For example, one potential feature could be the PSD of the signal. The network was trained on scores predetermined by clinicians. Specifically, two clinicians scored the signals used to train this model as either being a score of 0.0, 0.5, or 1.0, the mean of the scores from the two clinicians was used as target values. Examples



**Figure 5.5:** An example of the output for the NN can be seen here testing signals extracted from the first acquisition of patient one using PCA.

of the output of the NN can be seen in figure 5.5.

#### 5.3.5.2

#### Combine

```
Algorithm 10: Combining PCs
  Data: timeSeriesSinograms, PCs
   Result: respiratoryPC
1 respiratoryPC = first PC in PCs
2 remove respiratoryPC from PCs
3 respiratoryScore = get score from respiratoryPC and timeSeriesSinograms
5 for PC in PCs do
      currentRespiratoryScore = get score from PC and timeSeriesSinograms
6
 7
      scaledRespiratoryPC = respiratoryPC \times respiratoryScore
8
      scaledCurrentPC = PC \times currentRespiratoryScore
 9
10
      sumPC = scaledRespiratoryPC + scaledCurrentPC
11
      subtractPC = scaledRespiratoryPC - scaledCurrentPC
12
13
      sumRespiratoryScore = get score from sumPC and timeSeriesSinograms
14
      subtractRespiratoryScore = get score from subtractPC and
15
       timeSeriesSinograms
16
      if sumRespiratoryScore >respiratoryScore then
17
          respiratoryPC = sumPC
18
          respiratoryScore = sumRespiratoryScore
19
      else
20
          if subtractRespiratoryScore >respiratoryScore then
21
             respiratoryPC = subtractPC
22
             respiratoryScore = subtractRespiratoryScore
23
      end
24
25 end
```

After the PCs are sorted from a high to low score, they are then iterated over, both being summed and subtracted (with a weighting, the score), and a new score is found for both resulting signals. If one of the signals increases this score, then it becomes the new best PC, and goes forward to the next iteration. PCs are both summed and subtracted to handle the arbitrary sign problem mentioned earlier in section 5.3.3 (Bertolli et al. 2017b).

A similar method of combining signals can be seen in (Kesner and Kuntner 2010). However, the method presented here, attempts to improve on this by using a scoring

**Table 5.1:** This table shows the correlation coefficient between the result of the Score, Select, and Combine using NN scoring and the RPM for the train dataset where a number of preand post-processing methods have been applied to the data. Each new line represents the addition of this method, therefore the last element includes all previous pre- and post-processing. This is for the Freeman-Tukey transform, the Yeo-Johnson transform, the incorporation of a mask to remove low count areas of the sinogram, smoothing and downsampling the sinogram and smoothing the signal, and parallel compression.

Correlation Coefficient with RPM	20 s-840 s	20 s-140 s
Freeman-Tukey	0.783	0.703
Yeo-Johnson	0.790	0.747
Mask	0.826	0.824
Smoothing and Downsampling	0.826	0.825
Parallel Compression	0.830	0.839

metric whose behaviour better reflects the qualities of a 'good' signal, while in (Kesner and Kuntner 2010), the standard deviation is maximised. In addition, the proposed method compute the metric based on signals derived from PCs, as opposed to a single voxel or sinogram bin value (Kesner and Kuntner 2010), which should lead to noise reduction.

Please note that the algorithm 10 is described in terms of PCs. In fact, a simpler version just sums and subtracts the corresponding signals, this will give the same final signal.

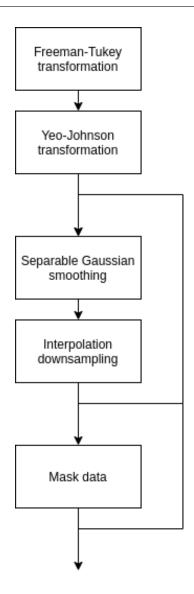
In the current implementation, this method is applied on all time intervals at once. It is possible to integrate this method with the Late Time Interval method in section 5.3.4 or the Moving Window method. However, this has not been demonstrated here.

## 5.4 Evaluation

Here, we discuss how the data was prepared for evaluation. We also present a suite of methods which can be applied generally to SSs (both from dynamic as well as potentially static acquisitions), in order to combat issues such as noise and outliers. Finally, in this section, we highlight how the methods in section 5.3 will be evaluated in section 5.5.

## **5.4.1** Data Preparation

TOF data were unlisted into low spatial resolution sinograms, each with a time frame duration of 500 ms, using the GE PetToolbox, following (Bertolli et al. 2017a), resulting in sinograms with dimensions  $95 \times 16 \times 47 \times 11$  (radial positions×angles×transaxial plane×TOF). To extract respiratory variation, the sampling rate of the PET sinograms



**Figure 5.6:** A diagram showing pre-processing performed.

was chosen as 2 Hz, so as to attempt to mitigate the effect of cardiac motion (Bertolli 2018).

Data was pre-processed element-wise by first applying a Freeman-Tukey transformation (Freeman and Tukey 1950), followed by a Yeo-Johnson power transformation (Yeo and Johnson 2000), to transform the Poisson distributed data to be more Gaussian-like.

In the current implementation, these transformations are applied element-wise on the (Poisson distributed) TOF sinogram  $S_p$ 

$$S_g = YJ_{\lambda}(FT(S_p)) \tag{5.2}$$

The Freeman-Tukey transformation is defined as

$$FT(X) = \sqrt{X+1} + \sqrt{X} \tag{5.3}$$

The Yeo-Johnson power transformation is defined as

$$YJ_{\lambda}(X) = \begin{cases} ((X+1)^{\lambda} - 1)/\lambda & \text{if } \lambda \neq 0, X \geq 0 \\ \log(X+1) & \text{if } \lambda = 0, X \geq 0 \\ -[(-X+1)^{(2-\lambda)} - 1)]/(2-\lambda) & \text{if } \lambda \neq 2, X < 0 \\ -\log(-X+1) & \text{if } \lambda = 2, X < 0 \end{cases}$$
(5.4)

The  $\lambda$  parameter is determined by minimising the Kullback Leibler Deviation (KLD) between normal distributions and the transformed distribution (Yeo and Johnson 2000). In this paper a single  $\lambda$  was determined from all of the data, although it would be feasible to find different  $\lambda$  values for every element in the sinogram.

The Yeo-Johnson power transformation was not included in previous papers. Its impact can be seen in table 5.1.

It was previously found through experimentation, that Gaussian smoothing of the resulting sinograms can improve results, especially in the case of the SAM method (Thielemans et al. 2013). In the current implementation, further downsampling was performed post-smoothing to reduce memory usage and increase computational speed. Linear interpolation was used as it was shown in a preliminary investigation to give satisfactory results at little computational cost. We use linear interpolation here to reduce the matrix size to a size necessary to represent the highest (spatial) frequency information present in the data post-smoothing. The effect this inclusion makes on the end result can be seen in table 5.1.

Finally, it has been found that the introduction of a mask to further aid in the reduction of noise is beneficial (Thielemans et al. 2011). The mask itself is defined as being true for any value, in the sinogram, above a predetermined threshold. Values not in the mask are removed prior to further execution, this is because these values can be assumed to mostly be noise. Note that a mask can also be used to eliminate parts of the data potentially affected by non-respiratory movement (Bertolli 2018), but this has not

been implemented in the current work. Again, the impact of this inclusion can be seen in table 5.1.

Values for the Gaussian smoothing, and the threshold of the mask, were determined using a grid search on a randomly selected subset of the data (specifically three patients). This data was then not used as part of any final evaluation, as was stated in section 5.3.1. The Gaussian smoothing sigma for the PCA based methods were 1.0, 0.5, and 2.0, for SAM based methods they were 1.0, 3.0, and 1.0 (this is for the sinogram radial positions, angles, and transaxial planes respectively). The mask threshold is selected such that it removes the bottom 5 % of values. The same mask is used for every time point. From an examination of the masks used for the test dataset, because a threshold value as low as 5 % is used, it appears that the mask is only removing parts of the background.

A flowchart of the pre-processing steps can be seen in figure 5.6.

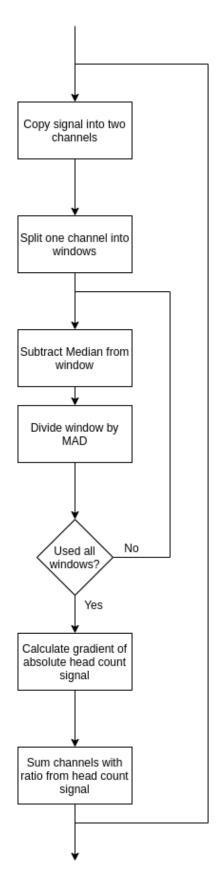
#### **5.4.2** Post-Processing

A flowchart of the method can be seen in figure 5.7 and figure 5.8.

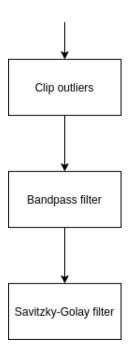
Regardless of the method used, there are still some effects of the radiotracer kinetics to be expected at early time intervals, and noise throughout. Thus, a method is proposed here to aid with the remaining radiotracer kinetics, and smoothing to help with noise in the extracted SS. The same post-processing is used regardless of the method to extract the SS.

## 5.4.2.1 Parallel Compression

Firstly there is, what shall be referred to as "parallel compression". This is a method borrowed from audio engineering (appearing notably in Dolby A noise reduction). The signal is split into two channels, one has its dynamic range reduced (through a process such as compression), while the other passes unchanged, before they are averaged back together (Izhaki 2012). This has the effect of reducing large differences in the dynamic range of the signal (for instance caused by tracer kinetics or some kind of drift), without losing a lot of breath to breath variability, compared to directly using the phase of the signal (if applied to respiratory SSs) (Lamare et al. 2022). For a more detailed explanation of the implementation please see the appendix E.



**Figure 5.7:** A diagram showing the parallel compression method.



**Figure 5.8:** A diagram showing the post-processing performed.

#### 5.4.2.2 Outlier Removal

Even though most of the large changes in intensity are remedied by parallel compression, some momentary spikes are still apparent. Thus, outliers are removed where they are outside a threshold of the quartile of the signal, and new values are interpolated.

## 5.4.2.3 Smoothing

Finally, smoothing is applied through the use of a bandpass filter (specifically a sinc filter), followed by a Savitzky-Golay filter (Savitzky and Golay 1964). A bandpass filter is used to remove frequencies in the signal outside of the respiratory window, and the Savitzky-Golay filter is used to promote local smoothness.

The bandpass filter is defined as

$$h_{BPF}(t) = 2B_H sinc(2B_H t) - 2B_L sinc(2B_L t)$$
(5.5)

where  $h_{BPF}$  is the bandpass function, t is the time distributed variable,  $B_H$  is the upper bound and  $B_L$  is the lower bound of the bandpass filter. The bandpass filter is implemented using a truncated sinc kernel. A polynomial order of three and a window length of five was used with the Savitzky-Golay filter, determined through a grid search on the training

**Table 5.2:** This table shows the p-values acquired as part of a statistical analysis of the test dataset using a mixed-effects model. Before applying the mixed-effects model the signals were first transformed to frequency space using FFT. One values between 0.05 Hz and 0.45 Hz were used. This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method.

Mixed-Effects Model	p-value
Static PCA	0.183
Moving Window PCA	0.483
Late Time Interval PCA	0.548
Score, Select, and Combine PCA PSD	0.649
Score, Select, and Combine PCA NN	0.908
Moving Window SAM	0.373

dataset.

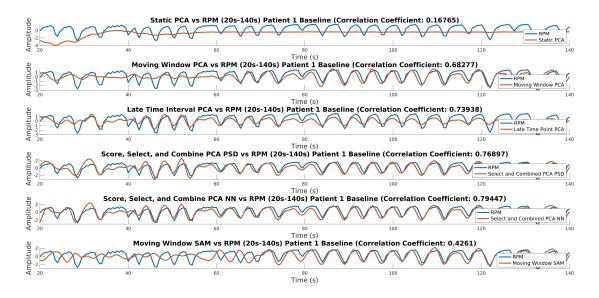
The impact of the inclusion of the above methods on the correlation of the Score, Select, and Combine method with NN based scoring with the RPM can be seen in table 5.1.

#### **5.4.3** Evaluation Methods

For evaluation of the results, the correlation coefficient of each SS between each method and the RPM, for all acquisitions in the test dataset, has been calculated. The correlation coefficient has been calculated for both the first 120 s (ignoring the first 20 s), and also the entire acquisition (between 20 s and 840 s). A statistical analysis using a mixed effects model has also been included.

All methods were compared to conventional PCA. We also included results for the Moving Window SAM method as this approximates KRG. While the Conventional and Late Time Interval methods can also be implemented using SAM, corresponding results are not shown here.

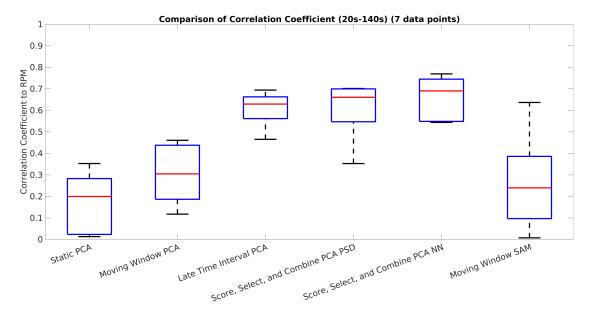
As stated earlier in section 5.3.1, parameters for the methods have been selected using a grid search on a randomly selected subset of the data (specifically three patients). This data was then not used as part of any final evaluation. The parameters were optimised by maximising the correlation coefficient between the SS and the RPM for the first 120 s of usable data (between 20 s and 140 s), due to there being initially no counts in the FOV.



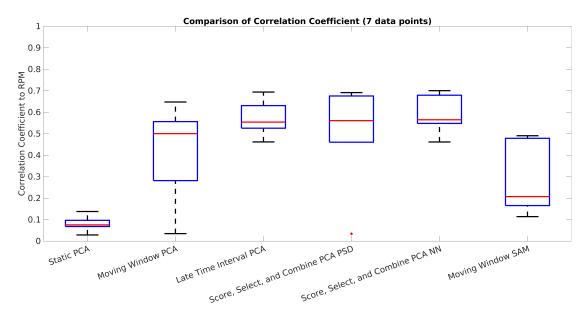
**Figure 5.9:** Output of each method compared to the RPM for the first usable 120 s (between 20 s and 140 s) (taken for the first acquisition of patient one). This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method.



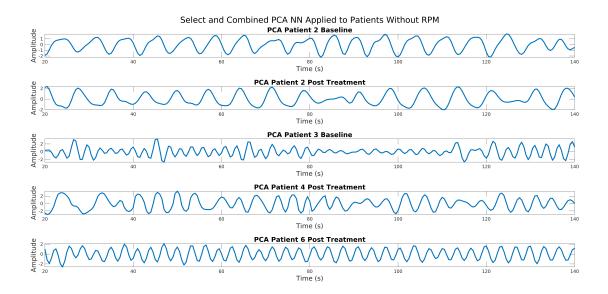
**Figure 5.10:** Output of each method compared to the RPM for the first usable 120 s (between 20 s and 140 s) (taken for the first acquisition of patient eight). This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method.



**Figure 5.11:** A box plot showing for each method its correlation coefficient to the RPM for the first usable 120 s (between 20 s and 140 s) (taken for the seven test acquisitions). This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method.



**Figure 5.12:** A box plot showing for each method its correlation coefficient to the RPM for the entire acquisition (taken for the seven test acquisitions). This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method.

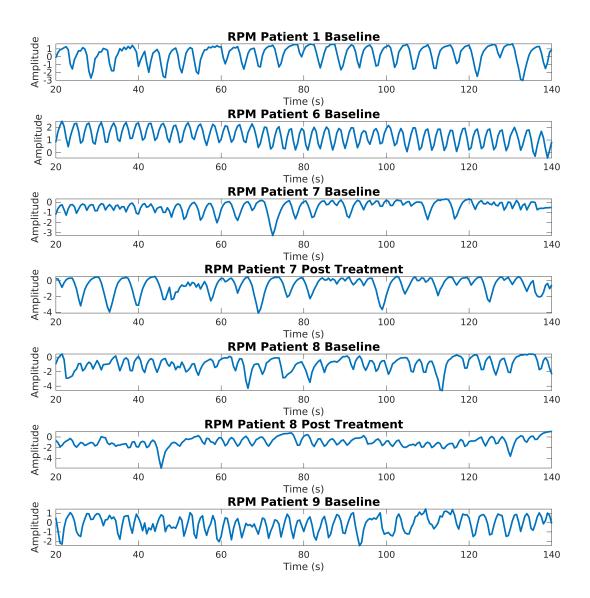


**Figure 5.13:** Signals for the first usable 120 s (between 20 s and 140 s) generated using the Select, Score, and Combine NN based scoring method. This is for the five acquisitions which did not have a usable RPM signal and as such could not be used for training or testing.

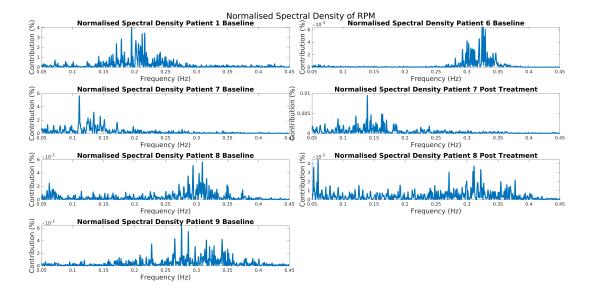
## 5.5 Results

A plot showing for each method its output, compared to the RPM, for the first 120 s (between 20 s and 140 s) (taken for the first acquisition of patient one), can be seen in figure 5.9. From a visual analysis, it can be observed that the Conventional PCA method has failed, post normalisation, it appears almost as if there is no variation in the signal at early time intervals. Both moving window methods show, towards the end of the plot, that they can extract a signal. However, it takes until between 60 s and 80 s for both methods to begin to pick up the signal. For the SAM based method, it appears as if the sign determination method has failed before 80 s, regardless though the method still cannot extract a signal before 60 s. The Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring methods, all appear to be able to extract a usable signal right down to 20 s (around when counts begin to appear in the FOV). The magnitude of the signal post 80 s more closely matches the RPM (or in comparison to before 80 s) for both Score, Select, and Combine methods than for the Late Time Interval PC method.

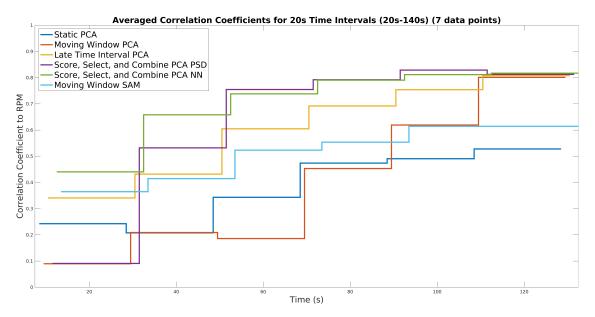
A plot showing for each method its output, compared to the RPM, for the first 120 s (between 20 s and 140 s) (taken for the first acquisition of patient eight) can be seen in



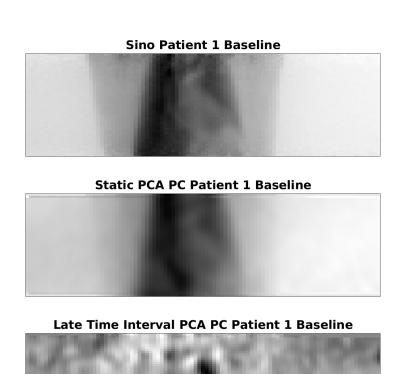
**Figure 5.14:** RPM signals for the first usable 120 s (between 20 s and 140 s) (for the seven test acquisitions). Notice that only the first acquisition of patient one and patient six shows a steady trace with an average frequency, every other trace shows variable breathing or artefacts.

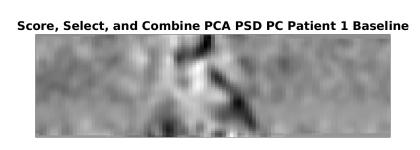


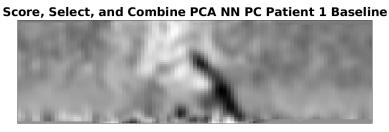
**Figure 5.15:** The result of applying FFT to the RPM signals for the first usable 120 s (between 0.05 Hz and 0.45 Hz) (for the seven test acquisitions). Notice that four of the seven have peaks very close to or outside the lower boundary of the resiratory window. Also notice that two of the seven have very wide frequency responses, which would be difficult for the automatic selection of a respiratory frequency window.



**Figure 5.16:** Correlation coefficients to the RPM for the first 120 s in 20 s intervals (between 20 s and 140 s) (taken as a mean for all data sets). This is for Conventional PCA, Moving Window PCA, Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring, and the Moving Window SAM method. The stair plots are staggered for the different methods for visual clarity.







**Figure 5.17:** A plot showing a single 'view' of the original PET data (top) as well as the PCs used to generate the output signal for the Conventional, Late Time Interval and the two Score, Select, and Combine methods (taken for the first acquisition of patient one).

figure 5.10. Similar results as for in figure 5.9 are repeated in figure 5.10, (although all methods match the RPM worse than in figure 5.9). This acquisition was selected to be shown due to it being a difficult trace to extract. Regardless, the Late Time Interval PC and Score, Select, and Combine methods appear to have extracted a signal early into the acquisition (from about 35 s on this patient and acquisition).

A box plot showing for each method its correlation coefficient to the RPM for both the first 120 s (between 20 s and 140 s), and also for the entire acquisition (taken for seven acquisitions), can be seen in figure 5.11 and figure 5.12. The correlation coefficient for the Conventional PCA method is low and the method is not usable. The correlation coefficients for both the moving window methods are roughly around 0.5, indicating that the Moving Window method is beneficial regardless of the method used to extract the signal for each window. However, again here, the correlation coefficient is lower than is acceptable. The results from the Late Time Interval PC, Score, Select, and Combine using frequency and NN scoring methods, all show correlation coefficients around 0.6 or higher, for both the early time interval as well as for all data. The Score, Select, and Combine methods show marginally higher correlation coefficient than the Late Time Interval PC method, and the NN shows slightly higher correlation coefficient than the frequency based scoring.

A plot showing for each method the evolution of the correlation coefficient with the RPM over time, for the first 120 s, by computing it in 20 s intervals, can be seen in figure 5.16. It can be observed, that on average across all data sets all methods struggle to produce usable results at the very beginning of the acquisition (around when counts begin to appear in the FOV). However, it is also apparent that on average both Score, Select, and Combine methods robustly begin to produce results, which closely match the RPM, as evidenced by a reasonable correlation past the first 40 s, on most acquisitions. The Moving Window method appears to perform well in the first interval.

A plot showing the components (or their combination) used to generate the output signal, can be seen in figure 5.17. One advantage of the sinogram-based methods is that the PC (or signed mask for SAM) can be visualised to see how it corresponds to anatomy and tracer uptake. In figure 5.17 it can be seen that the Conventional PCA method returns a PC which closely resembles the input data, leading to the conclusion that the variation

in the selected PC is dominated by the kinetics. The other methods produce a PC which is more related to edges of internal structures, where respiratory movement occurs. A visual inspection indicates that the least confounding variation, and noise, is included in the Score, Select, and Combine using the NN scoring method. Curiously, it appears that the Late Time Interval, and the Score, Select, and Combine using the NN scoring method return very similar distributions. However, the Score, Select, and Combine using the frequency scoring method also returns a high value region in tissue at the top of the image.

We utilised a mixed-effects model to analyse the differences between various methods and RPM measurements across all subjects. This model incorporated 'method' as a fixed effect and treated subjects as a random effect, to account for inter-subject variability. We found that the comparisons between each method and RPM were not statistically significant, with all p-values > 0.05, see table 5.2. While this may suggests that, within the constraints of our analysis and available data, there was no evidence to support significant differences between any methods in comparison to RPM measurements, given the small number of subjects, it's important to consider that this analysis may be considerably underpowered.

It is noteworthy to mention that we observed that the Late Time Interval PC and Score, Select, and Combine using frequency and NN scoring methods exhibited the highest p-values. This may suggest a lesser degree of difference from RPM. However, such interpretations should be approached with caution and not be taken as conclusive evidence of similarity.

Plots showing a comparison of results achieved with and without pre- and post-processing can be seen in appendix D.

Finally, results of applying the Score, Select, and Combine using NN scoring method to data where the RPM could not by synced with the list mode data can be seen in 5.13.

## 5.6 Discussion

This chapter introduces several methods for DD extraction of a respiratory signal from dynamic PET data. To the best of our knowledge, this appears to have only been attempted in (Schleyer et al. 2014). Data used here are from a [<sup>18</sup>F]-FDG study on patients with IPF,

while in the latter paper, Nitrogen-13 Ammonia ([<sup>13</sup>N]-NH3) data was used to evaluate the proposed KRG method. These tracers have different uptake patterns.

The work presented has several limitations. Firstly, the data used all originates from the same study, using the same procedure, the same radiotracer, and acquired on the same scanner. In order to better validate the generalisability of the method it would be positive to test on data acquired on different scanners and using different radiotracers. Additionally, from the data acquired, only a subset of this data is usable due to issues during acquisition. It would be beneficial to test the methods on a larger sample of patients. Furthermore, it would be beneficial for the data to include both a larger number of non-complex and complex breathers to better test the limitations of the methods.

An additional concern is the point at which the methods may fail, for patients who exhibit abnormal breathing patterns. For instance, extremely slow breathers will breathe at a rate less than 0.1 Hz, which in the case of the Score, Select, and Combine method using the frequency scoring method and a fixed frequency respiratory window would be considered to be radiotracer kinetics. Furthermore, when using a non-fixed respiratory window this method struggles with patients who breathe less regularly, as the window is expanded to include parts of the kinetics and noise (this is shown in figure 5.15). The discrepancy between the results for the SAM Moving Window method, presented here, and the KRG ones, shown in (Schleyer et al. 2014), could also be attributed to this complexity. Many of the patients breathed at varying rates, stopped breathing during acquisition, or breathed unusually fast or slow (this is shown in figure 5.14). This was probably due to the data being acquired for an IPF study.

The presented method includes several pre- and post-processing steps. Although these have been shown to be beneficial on the train data, table 5.1, the impact of some steps is relatively small. It remains to be determined on a larger dataset if some of these steps could be removed.

Furthermore, please see appendix F for work which was inspired by the methods here (Whitehead et al. 2023a), (Ferrante et al. 2022), (Ferrante et al. 2024). In appendix F (Whitehead et al. 2023a) a method to extract a TAC from dynamic PET is presented. In the work presented here we are attempting to extract a SS from dynamic data, the main other signal which is interfering with this is simply the TAC. Therefore,

in the work in appendix F conceptually we are attempting to perform the inverse of what is presented in this section. Which from the point of view of the method itself is quite a similar problem. However, in appendix F the TAC is extracted directly using a NN rather than using a NN scoring method to inform a more traditional signal processing.

In the future, research will focus on further development of the methods, including optimisation of the NN scoring method. In the next stage, these methods will be applied to the task of implementing advanced respiratory motion correction on dynamic PET data.

## 5.7 Conclusion

We have presented and evaluated several methods for extraction of a respiratory signal from dynamic PET data. Results from a visual comparison of early time interval output signals compared to the RPM, quality of PC, and correlation coefficient of the output signals to the RPM, indicates that the Late Time Interval PC and both Score, Select, and Combine methods are more robust and afford higher quality signals than moving window methods. The results also indicate that both Score, Select, and Combine methods can give a higher correlation coefficient earlier than the Late Time Interval PC method. Scoring using the NN shows slightly higher correlation coefficients than the frequency based scoring.

These findings represent a significant step toward enabling fully DD respiratory motion correction in dynamic PET.

## Chapter 6

## **Discussion and Conclusion**

#### 6.1 Overview

This chapter summarises the key findings from each stage of the work and outlines directions for future research.

Respiratory motion is a major challenge in PET imaging, with the potential to compromise diagnostic accuracy and clinical outcomes. In response, this thesis has developed and evaluated a motion correction framework aimed at reducing the impact of respiratory motion in both static and dynamic PET/CT imaging. A central focus of the work has been to ensure that the proposed methods are compatible with existing clinical practice, using a single CT, minimising the need for additional hardware, avoiding reliance on external surrogate signals, and maintaining manageable computational demands. The motion correction method (developed in Chapters 3 and 4) was developed by investigating several components of the framework on simulated data, using increasingly realistic data. Furthermore, a DD SS extraction method for dynamic PET (presented in Chapter 5) was developed and evaluated on patient data. Although the motion correction framework and the DD SS extraction method for dynamic PET were evaluated separately, they have not yet been fully integrated and tested as a single end-to-end pipeline. Nevertheless, the progress made in each area paves the way towards supporting motion correction in dynamic PET imaging.

### **6.2** Main Results

- Section 3.3 evaluated the impact of the introduction of TOF PET on the quality of motion estimation possible with NAC PET. This work was performed in order to assess if it was feasible to further develop a motion correction method which was unbiased by a static CT  $\mu$ -Map. If a static CT  $\mu$ -Map was used for AC, prior to motion correction, then it is likely that the artefacts added would cause the motion correction to fail. However, if NAC PET is used, it is likely that there is not enough contrast in the reconstructed NAC volumes within the lungs and with the diaphragm for motion correction to be possible. Preliminary experiments showed that TOF in the absence of a  $\mu$ -Map significantly improved the distribution of the signal in the NAC volumes. This work indicated that it was in fact possible to perform an adequate motion correction using NAC TOF PET. However, the simulations used in this section were simplified to enable an initial exploration. They excluded intragate respiratory motion, randoms, and scatter, and assumed a high TOF resolution. While this allowed focused testing of key algorithmic components, later studies showed that such assumptions limit the effectiveness of NAC TOF OSEM when applied to more realistic data.
- Section 3.4 introduced a motion correction approach aimed at addressing the challenge posed by the use of a single static μ-Map in the presence of respiratory motion. For accurate motion correction, each gated bin ideally requires attenuation correction using a μ-Map that matches its corresponding respiratory position. This section evaluated whether deformation of a single static μ-Map to the position of each bin is sufficient, and whether fitting a new MM on AC data offers any improvement. The results demonstrated that deforming the static μ-Map to match each gate produced motion-corrected images without introducing noticeable misalignment artefacts. Moreover, fitting a second MM on the already AC data did not improve the motion correction, suggesting that the additional computational burden of this step is unnecessary. These findings validate a more efficient approach to attenuation correction in motion-compensated PET. However, in addition to the still relatively simplified simulations, it should be noted that the static μ-Map was positioned at

the mean respiratory phase, arguably the most favourable condition for this method. This is clinically impossible to guarantee.

- Section 3.5 introduced a significantly more realistic simulation setup compared to previous sections. The SSs used to drive the XCAT phantom incorporated both hysteresis and breath-to-breath variability, better reflecting true respiratory motion. Additionally, the lesion model was reduced to a clinically realistic size. A background activity term was also introduced for the first time, further enhancing realism of the noise level in the data. This section focused on evaluating the impact of incorporating a MM within a modular motion correction pipeline, and comparing the performance of pair-wise versus group-wise registration schemes. Results demonstrated that the inclusion of a MM consistently improved lesion depiction over more traditional registration-based methods. Notably, the pair-wise registration approach failed to perform adequately without the MM, while group-wise registration produced better results. Nevertheless, in both cases, the use of a MM led to improved lesion homogeneity. Based on these findings, group-wise registration with MM fitting was adopted for all subsequent experiments.
- Section 4.3 exploited an important practical advantage of motion modelling: the ability to estimate motion fields from data gated at a different temporal resolution than the data to which the correction is applied. This is feasible as long as the space of the SS remains consistent. This was the first piece of work which made use of significantly more complex simulations, and also introduced MLACF as the reconstruction method (as opposed to NAC OSEM) for the initial volumes used for the motion estimation. With the more complex simulations, and with MLACF, the same gating scheme as was used in section 3.5 now resulted in a complete failure of the motion correction. However, with the new coarse gating scheme (applied to data binned in the same way) the results were significantly improved. With the coarse gating scheme, motion estimation was successful even on noisy data, with little detrimental impact observed on high count data. This section provided the final components of a flexible framework for motion correction of (static) PET/CT data.

- Section 4.4 presented an evaluation of the framework on realistic simulated data, together with some improvements to be able to handle such data. For instance, the simulation was swapped from a 2D to a 3D acquisition. Moreover, the TOF resolution was reduced to that one of the clinical scanners available at our institution. These changes caused MLACF to underperform. Because of this, the MLACF algorithm was adapted to make it more robust to the application of motion estimation. This was mainly achieved by initialising and constraining the ACFs using the CT μ-Map. The section then explored the impact of different gating schemes (phase vs amplitude, 1D vs 2D SSs), bin numbers, and reconstruction algorithms (MLACF, AC OSEM, and NAC OSEM). Results showed that for the MM estimation, using fewer bins was more effective, and MLACF consistently outperformed both types of OSEM when the same gating setup was used.
- Section 4.5 provided a thorough overview of the final method as well as further
  evaluating the method with data making use of various lesion sizes. This provided
  further evidence that the method is able to provide good results even for lesions of
  8.0 mm diameter, which is almost as small as the resolution of the scanner, as well
  as currently considered clinically relevant.
- Chapter 5, demonstrated work on DD SS extraction for dynamic PET. Here, a multitude of methods were presented and evaluated against each other to extract a respiratory SS from dynamic PET (where the tracer kinetics at the beginning of the scan obscure the respiration). These methods were compared to an existing method from the literature (Schleyer et al. 2014). Three methods proved to have generally better and robust performance when compared to the literature. The first method proposed taking a PC from a late time interval and applying it to the whole data set. The second and third methods used a frequency and NN based scoring metric to combine PCs. On the whole, the NN based method gave the best performance. However, the NN used in this method is proprietary. Importantly, this work establishes a pathway for extending motion correction techniques, previously limited to static imaging, to dynamic PET, marking a significant step toward broader clinical applicability.

### 6.2.1 Main Contributions and Knowledge Advancement

This thesis demonstrates that effective respiratory motion correction in PET/CT can be achieved using TOF PET data without reliance on a respiratory phase matched CT derived  $\mu$ -Map or external monitoring of the respiration. By leveraging TOF information and deformation of a static (end inhalation)  $\mu$ -Map, the need for multiple respiratory phase  $\mu$ -Map (or the artifacts associated with a misaligned static  $\mu$ -Map) were eliminated. The introduction of the MLACF reconstruction method, along with its tailored adaptation for motion correction, improved the motion estimation and hence image quality and outperformed conventional OSEM based approaches, especially under complex motion and gating conditions when TOF resolution was low.

A flexible and modular motion correction framework was developed and validated through increasingly realistic simulations. The framework incorporates a MM and supports the use of group-wise registration, which was shown to outperform pair-wise approaches, and uses a (count) weighted fitting. It incorporates alignment to a single CT. It was shown to improve lesion depiction and homogeneity in simulated data in realistic but challenging conditions. The method proved robust across a range of lesion sizes, including those near the resolution limits of the scanner.

A key practical advancement was the demonstration that MMs can be derived from coarsely gated data and applied to finely gated reconstructions, allowing for reduced noise and improved DVF estimation. To our knowledge this is a specific advantage of the MM method presented and has not been attempted before in PET/CT.

The extension of DD motion correction techniques to dynamic PET was enabled through DD SS extraction methods that remained effective despite tracer uptake variability. These advances pave the way for broader clinical translation of motion correction, with clearly defined strategies for gating, attenuation correction, and model integration.

### 6.3 Outlook

Of the various motion correction strategies explored in this thesis, the method presented in section 4.5 (which integrates gated PET based motion modelling with MLACF based attenuation correction) represents the most promising candidate for clinical implementation. In fact, as it offers a balance between robustness, quantitative accuracy, and computa-

tional feasibility, it has the potential to significantly advance static or even dynamic PET imaging and advance disease staging in clinical practice. However, for this method to be validated for clinical use, several main challenges remain to be addressed.

#### 6.3.1 Limitations

Despite the advances presented, there are several limitations of the work presented in this thesis:

- Firstly, the biggest limitation of the results presented in chapter 3 and chapter 4 is that the methods have only been tested on simulated data. The anatomy used for the XCAT simulation did not vary, it was an average male anatomy. In order to prove generalisability, and to avoid unnecessary bias, ideally the method should be tested on a multitude of anatomies including both male and female, big and small. The ground truth activity distributions were taken from XCAT simulated volumes consisting of organs and structures that are homogeneous, which is unrealistic. Note however that this actually poses a significant problem to most registration algorithms, where there would be no gradient across a homogeneous region.
- An additional issue posed by the use of XCAT simulations would be that the MM used to drive the simulation over time is rather simple, in particular it is a linear MM. Whereas the motion present in patients will probably never be able to be entirely represented by a linear system. The MM method used in the thesis also uses a linear MM, meaning that all of the motion present in the data could be captured by the method (under ideal circumstances), possibly giving misleading results. Nevertheless, the use of linear MMs is wide-spread in the literature, including for modalities such as CT and MR which have considerably higher spatial resolution than PET.
- The values used for the XCAT simulation were similar to the activity distribution of [18F]-FDG. While this is the most commonly used radiotracer, other radiotracer are of interest, for instance for cardiac imaging or prostate cancer. The results of the registration and therefore motion correction will depend on the activity distribution, which remains to be investigated.

- The objective function used for the MM fitting also used SSD which is sensitive to outliers and would cause a limitation during early iterations of the motion correction. This is when it is likely that the quality of the registration is poor and thus the DVF is also less than ideal. This limitation may be compounded if the methods were combined iteratively with PET reconstruction. This is because, at early time points not only would the registration be unsatisfactory, but the volumes provided to the registration by the reconstruction would also be poor. To mitigate this, an objective function which is robust to outliers could be used, such as Huber (Verducci et al. 2007), or a regression method such as Random Sample Consensus (RANSAC) (Fischler and Bolles 1981). Huber is defined as being equivalent to squared error below a threshold and absolute error above it (Verducci et al. 2007).
- The final images were obtained by gating the data into 30 bins, reconstructing the volumes (with attenuation matching), resampling the volumes to the reference position, and finally using a (weighted) average of the 30 images. It is also possible to use motion incorporated image reconstruction, where a single volume is reconstructed from all the gates by incorporating the DVFs into the forward model (Qiao et al. 2006), (Manjeshwar et al. 2006). This has been shown to have better noise properties as well as reduce bias due to the non-negativity constraint implied in OSEM (Polycarpou et al. 2012).
- Another limitation and a possible future direction for research would be that the objective function for the registration is currently calculated in image space. If the objective function was shifted to sinogram space then it would be much more likely that a combined iterative PET reconstruction and motion correction would converge to a good solution. This is because both methods are using objective functions calculated in the same space.

#### **6.3.2** Future Work

There are a number of potential future avenues for work, these include the following:

• A clear next step for future work is the full evaluation of the motion correction framework using patient data, as outlined in Section 4.6. The technical groundwork

for this has already been established, but patient evaluation presents new challenges, including the absence of a definitive ground truth. As such, new evaluation strategies will be required. One option is to incorporate expert clinical assessment as part of a qualitative evaluation, although this alone will not suffice. Quantitative image quality metrics like PIQE were explored in Section 4.4.2.4 due to their independence from a ground truth. However, results in Section 4.4.3 highlighted potential reliability issues with PIQE in this context. A NN based no reference image evaluation metric could be developed, if pairs of images and scores could be acquired.

- Chapter 5 opens up the possibility for the motion correction methods to be applied to dynamic PET. Aside from difficulties with low signal-to-noise ratios in this data, the main difficulty is that at early time frames, determining the respiratory motion will be difficult, as the tracer will still be mainly in the aorta. As a SS was successfully obtained for these early time frames in this chapter MMs may be useful to overcome this. If early time intervals are removed during the registration and MM fitting steps, then the SS could be used to generate DVFs for these time intervals from the MM (Emond et al. 2019).
- DIP, as seen in appendix B (Whitehead et al. 2022a), was presented in (Ulyanov et al. 2020) as a NN based method to solve general inverse problems. DIP could potentially aid in a NN based method to approximate a generalised framework unifying registration and MM estimation, NiftyRegResp (now named SuPReMo). If a NN was designed which took as input an *n*-dimensional SS and output a DVF then it could be trained in the same manner as DIP. This would result in the production of a NN model which is similar to a MM. If no activation functions were used then conceptually the NN would be almost identical to the model in NiftyRegResp (SuPReMo). However, if activation functions were added, the model would be nonlinear. In order to implement this, the most necessary step would be a re-sampler written in the NN framework of choice. Parametrisation and regularisation (such as those in NiftyReg) could be added as either an additional architectural requirement or summed with the loss.

- Appendix F (Whitehead et al. 2023a) presents a method using NNs to extract a TAC corresponding to the blood concentration from dynamic brain PET. When the FOV includes the thorax, then one of the signals which the method would be trying to remove (or ignore) may be respiration. The method developed in chapter 5 could be considered to be trying to remove the TACs from the respiratory SS (if the FOV had covered the thorax). It might therefore be possible to combine the two works to 'invert' the purpose of the NN from appendix F. This would result in a NN which can be used to extract respiratory SS from dynamic PET.
- Finally, there have been strides recently in the domain of DIP applied to the PET reconstruction process. However, there has not been seen recently a method to reconstruct 4D or dynamic PET using DIP. This would be a logical extension of the methods presented in appendix B. The main limitation to implementing this is that the forward model would be required to be implemented in the NN framework of choice.

## Appendix A

# **Motion Correction Methods Table**

**Table A.1:** This table shows a short summary of the highlights of all methods used in chapter 3 and chapter 4. This table is useful for quickly referencing the changes between sections and the development of the method throughout the thesis.

Method	Data	Reconstruction	<b>Motion Estimation</b>
Details			
Section 3.3	Four Dimensional Extended Cardiac Torso (XCAT):  - 40.0 mm lesion  - One Dimensional (1D) XCAT respiratory traces  - No hysteresis  - 6 volumes  - No intra-gate motion  Simulation:  - Two Dimensional (2D) Simulation  - 120 s simulation  - American clinical count rate (True 60 Kilo Counts Per Second (KCPS))  - Non-Time of Flight (Non-TOF) and Time of Flight (TOF)  - 375.0 ps TOF resolution  - No scatter or random events  Gating:  - XCAT Surrogate Signal (SS)  - 6 respiratory bins	Non-Attenuation Corrected (NAC) Ordered Subset Expectation Maximisation (OSEM) (2 iterations, 24 subsets)     Gaussian post-smoothing (6.4 mm Full Width at Half Maximum (FWHM))	<ul> <li>Jointly estimating registration and Motion Model (MM)</li> <li>Non-diffeomorphic</li> <li>1D SS</li> <li>Sum of Squared Differences (SSD) objective function</li> <li>Default parameters (Control Point Grid (CPG) spacing 5 times voxel size, 5e-3 bending energy weight)</li> </ul>

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.4	XCAT:	Reconstruction for motion estimation:  NAC OSEM (2 iterations, 24 subsets)  Gaussian post-smoothing (6.4 mm FWHM)  Motion compensated reconstruction:  Attenuation Corrected (AC) OSEM (2 iterations, 24 subsets)  Gaussian post-smoothing (6.4 mm FWHM)	<ul> <li>Activity MM:         <ul> <li>Jointly estimating registration and MM</li> <li>Non-diffeomorphic</li> <li>1D respiratory SS</li> <li>SSD objective function</li> <li>Tuned CPG spacing (12 times voxel size) and bending energy weight (9e-6)</li> </ul> </li> <li>Attenuation to activity registration:         <ul> <li>Non-diffeomorphic</li> <li>Normalised Mutual Information (NMI) objective function</li> </ul> </li> <li>Tuned CPG spacing (5 times voxel size) and bending energy weight (9e-2)</li> <li>Multi-resolution registration (4 levels)</li> <li>Attenuation deformed twice</li> </ul>

Method	Data	Reconstruction	Motion Estimation
Details			
Section 3.5	XCAT:     12.0 mm lesion     2D Magnetic Resonance     (MR) respiratory traces     Hysteresis     240 volumes     Intra-gate motion      Simulation:     2D Simulation     120 s simulation     European clinical count rate (True 48 KCPS)     375.0 ps TOF resolution     Pseudo scatter (triangle filter 22 voxels, 39 KCPS) and random (56 KCPS) events      Gating:     2D PCA respiratory SS     30 respiratory bins (10 displacement, 3 gradient)	<ul> <li>Reconstruction for motion estimation:         <ul> <li>NAC OSEM (2 iterations, 24 subsets)</li> <li>Replication of end slices to size of Attenuation Map (μ-Map)</li> <li>Gaussian post-smoothing (2 times voxel size)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing (6.39 mm FWHM transverse, 3.27 mm axial)</li> </ul> </li> </ul>	Registration:  Non-diffeomorphic  NMI objective function  Tuned CPG spacing (12 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (8)  Multi-resolution registration (4 levels)  Pair- and group-wise registration  Activity and attenuation Deformation Vector Field (DVF) composed  MM:  With and without MM  MM fit only at highest resolution level  Weighted (counts) linear regression MM  2D respiratory SS

Method	Data	Reconstruction	Motion Estimation
Details			
Section 4.3	XCAT:	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC Maximum Likelihood Expectation Maximisation (MLEM) initialisation (1 iteration)</li> <li>Maximum Likelihood Activity and Attenuation Correction Factors Estimation (MLACF) (7 activity iterations, 24 subsets, 9 attenuation iterations)</li> <li>Quadratic prior (1.0)</li> <li>Attenuation Correction Factor (ACF) initialised with ones</li> <li>Activity volume and ACF standardised between iterations</li> <li>Replication of end slices to size of μ-Map</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	<ul> <li>Registration:         <ul> <li>Diffeomorphic (velocity field parametrised, symmetric)</li> <li>NMI objective function</li> <li>Tuned CPG spacing (8 times voxel size), bending energy (1e-4 activity, 1e-2 attenuation), and number of iterations (5)</li> <li>Multi-resolution registration (5 levels)</li> <li>Pair-wise registration initialisation</li> <li>Group-wise registration</li> <li>Activity and attenuation DVF composed</li> </ul> </li> <li>MM:         <ul> <li>MM fit between each resolution level</li> </ul> </li> </ul> <li>MM fit between each resolution level</li> <li>Weighted (counts) linear regression MM</li> <li>1D or 2D respiratory SS</li>

Method	Data	Reconstruction	Motion Estimation
Details			
Section 4.4	• XCAT:  - 12.0 mm lesion  - 2D MR respiratory traces  - Hysteresis  - 480 volumes  - Intra-gate motion  • Simulation:  - Three Dimensional (3D)  Simulation  - 240 s simulation  - European (late dynamic)  research count rate  (41 KCPS)  - 550.0 ps TOF resolution  - Pseudo scatter (Gaussian  smoothing 128.0 mm  FWHM, 29 KCPS)  and random events  (57 KCPS)  • Gating for motion estimation:  - 1D or 2D PCA respiratory  SS  - 30 respiratory bins (10 displacement, 3 gradient) or  - 5 respiratory bins (5 displacement, 1 gradient) or  - 3, 4, 8, or 12 radial pseudophase gates  • Gating for motion compensated reconstruction:  - 1D or 2D PCA respiratory  SS  - 30 respiratory bins (10 displacement, 1 gradient) or  - 3, 4, 8, or 12 radial pseudophase gates	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC OSEM initialisation (2 iterations, 24 subsets)</li> <li>MLACF (3 activity iterations, 24 subsets, 1 attenuation iterations)</li> <li>ACF initialised with μ-Map</li> <li>Activity volume scale set to OSEM reconstruction and</li> <li>ACF scale set to μ-Map between iterations</li> <li>Re-interpolation of outliers between iterations</li> <li>Smoothing of endplanes between iterations</li> <li>Median smoothing (kernel size 3) between iterations</li> <li>Sinogram mashing in attenuation update (1, 7, 2, 3 mashing factor respectively)</li> <li>Filled with Not a Numbers (NaNs) to size of μ-Map</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	<ul> <li>Registration:         <ul> <li>Diffeomorphic (velocity field parametrised, symmetric)</li> <li>NMI objective function</li> <li>Tuned CPG spacing (13 times voxel size) and bending energy (1e-1 attenuation)</li> <li>Multi-resolution registration (4 levels)</li> <li>Iterations used at resolution level: 2, 2, 2, 5</li> <li>Pair-wise registration initialisation</li> <li>Group-wise registration</li> <li>Activity and attenuation DVF composed</li> </ul> </li> <li>MM:         <ul> <li>MM fit between each resolution level</li> <li>Weighted (counts) linear regression MM</li> <li>1D or 2D respiratory SS</li> </ul> </li> </ul>

Method	Data	Reconstruction	Motion Estimation
Details Section 4.5			
	XCAT:	<ul> <li>Reconstruction for motion estimation:         <ul> <li>AC OSEM initialisation (2 iterations, 24 subsets)</li> <li>MLACF (3 activity iterations, 24 subsets, 1 attenuation iterations)</li> <li>ACF initialised with μ-Map</li> <li>Activity volume scale set to OSEM reconstruction and</li> <li>ACF scale set to μ-Map between iterations</li> <li>Re-interpolation of outliers between iterations</li> <li>Smoothing of endplanes between iterations</li> <li>Median smoothing (kernel size 3) between iterations</li> <li>Sinogram mashing in attenuation update (1, 7, 2, 3 mashing factor respectively)</li> <li>Filled with NaNs to size of μ-Map</li> </ul> </li> <li>Motion compensated reconstruction:         <ul> <li>AC OSEM (2 iterations, 24 subsets)</li> <li>Gaussian post-smoothing transverse (6.4 mm FWHM)</li> <li>'Normal' Z-filter axially (1.0, 4.0, 1.0 convolution)</li> </ul> </li> </ul>	<ul> <li>Registration:         <ul> <li>Diffeomorphic (velocity field parametrised, symmetric)</li> <li>NMI objective function</li> <li>Tuned CPG spacing (13 times voxel size) and bending energy (1e-1 attenuation)</li> <li>Multi-resolution registration (4 levels)</li> <li>Iterations used at resolution level: 2, 2, 2, 5</li> <li>Pair-wise registration initialisation</li> <li>Group-wise registration</li> <li>Activity and attenuation DVF composed</li> </ul> </li> <li>MM:         <ul> <li>MM fit between each resolution level</li> <li>Weighted (counts) linear regression MM</li> <li>2D respiratory SS</li> </ul> </li> </ul>

### Appendix B

# Pseudo-Bayesian DIP Denoising as a Preprocessing Step for Kinetic Modelling in Dynamic PET

### **B.1** Abstract

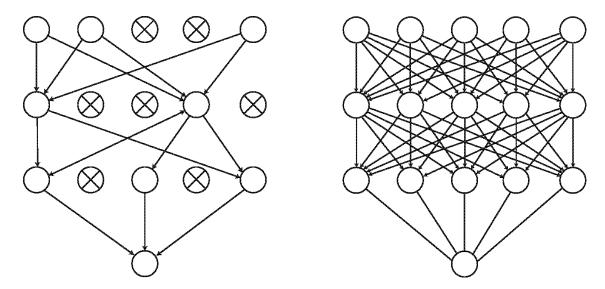
Noise (among other artefacts) could be considered to be the bane of Positron Emission Tomography (PET). Often, it causes what could otherwise be a more simple problem to explode in complexity. Many methods have been proposed to alleviate the worst annoyances of noise, however, not many take into account the temporal nature of dynamically acquired PET. Here, we propose an adaption of a method, which has seen increasing attention in more traditional imaging denoising circles. Deep Image Prior (DIP) exploits the initialisation of a carefully designed Neural Network (NN), so as to treat it as a bank of custom filters, which are to be trained and used afresh on each new image, independently. DIP has seen adaptation to PET previously (including dynamic PET), however, many of these adaptations do not take into account the large memory requirements of the method. Additionally, most previous work does not address the main weakness of the DIP, its stopping criteria. Here, we propose a method which is both memory efficient, and includes a smoothing regularisation. In addition, we provide uncertainty estimates by incorporating a Bayesian approximation (using dropout), and prototype a training scheme by which the model is fit on all data simultaneously. The denoised images are then used as input for kinetic modelling. To evaluate the method, dynamic Four Dimensional Extended Cardiac Torso (XCAT) simulations have been produced, with a Field of View (FOV) of the lung and liver. The results of the new methods (along with Total Variation (TV) and the old DIP) have been compared by a visual analysis, Structural Similarity Index Measure (SSIM), and  $K_i$  values. Results indicate that the new methods potentially outperform the old methods, without increasing computation time, while reducing system requirements (Whitehead et al. 2022a).

### **B.2** Introduction

Most machine learning or NN based methods, rely upon a workflow where a model is designed, trained, validated, and deployed (Kröse et al. 1993). This is logical as the intention with NN based methods is to treat them as a general function approximator (Kröse et al. 1993). This is where the function is determined by the relationship between the data used and the objective or loss function. However, in the domain of image denoising, the DIP method has received attention as training and inference are performed independently on each new image (Ulyanov et al. 2020). The DIP method could be considered to be a custom learnt bank of filters for each input. In order to prevent overfitting, the number of iterations used is imperative. The authors of the original DIP paper argue that the method could be considered to be used to solve many inverse problems (Ulyanov et al. 2020).

For PET, there have been a number of adaptations of DIP. (Gong et al. 2019) used a U-Net with relatively high count/low motion brain scans (Weng and Zhu 2021). In (Hashimoto et al. 2021) DIP is extended to Four Dimensional (4D) dynamic PET. To do this, multiple output branches are grafted onto the NN, one for each dynamic time point. (Hashimoto et al. 2019) uses the original or static PET acquisition as input to the NN, rather than noise. (Yang and Huang 2022) represents a more recent extension, where multiple NNs are used simultaneously.

This work seeks to extend or simplify previous work, in order to denoise 4D dynamic PET data.

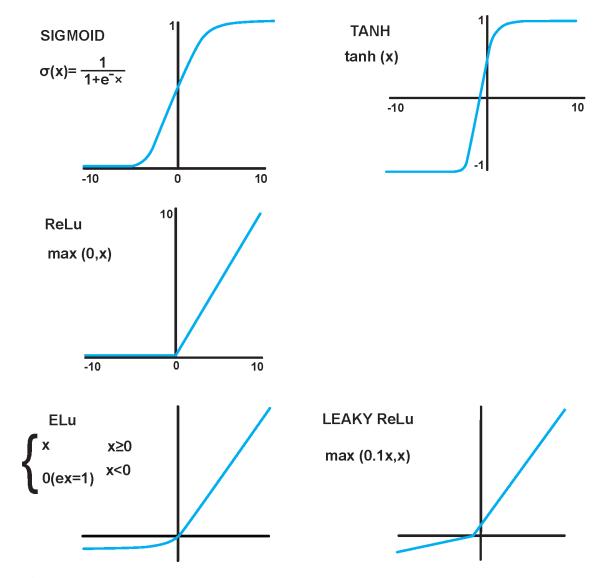


**Figure B.1:** Graphical representation of dropout. On the left of this figure a NN which is experiencing a large degree of dropout can be seen. Here, nodes which have been dropped are represented with a cross through it and their outgoing/incoming connections are remove. On the right of this figure the same network can be seen without dropout and will all nodes active and fully connected.

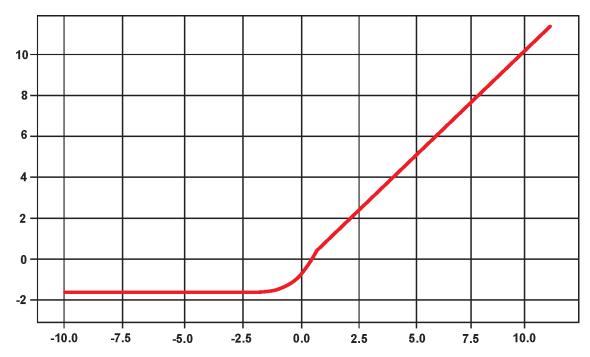
### **B.3** Methods

### **B.3.1** Network Design and Execution

Firstly, we reduced Graphics Processing Unit (GPU) memory requirements. Because DIP requires training at inference, the full amount of memory is needed every time the method is used. To aid in clinical adoption, a hard limit of 8.0 GB GPU memory was imposed. Secondly, a more robust stopping criteria is required. The stopping criteria should be flexible so as to allow the method to correct a wider variety of inputs. One method would be, to look at a window of previous loss function values, and exit, when the gradient of this drops below a tolerance. To aid in achieving this, as well as to address weaknesses of the original DIP, regularisation must be added, so as to stop the NN fitting to noise, similar to (Liu et al. 2019). Finally, because the output from this method is to be used in a further kinetic model fitting, it may provide improved results to have a metric of the uncertainty in the denoised images. We used dropout as in (Gal and Ghahramani 2016) to approximate (more expensive) Bayesian inference. A graphical representation of dropout can be seen in figure B.1. Furthermore, this work uses PET data with a FOV of the lung and liver, whereas most previous work uses a FOV of the head.

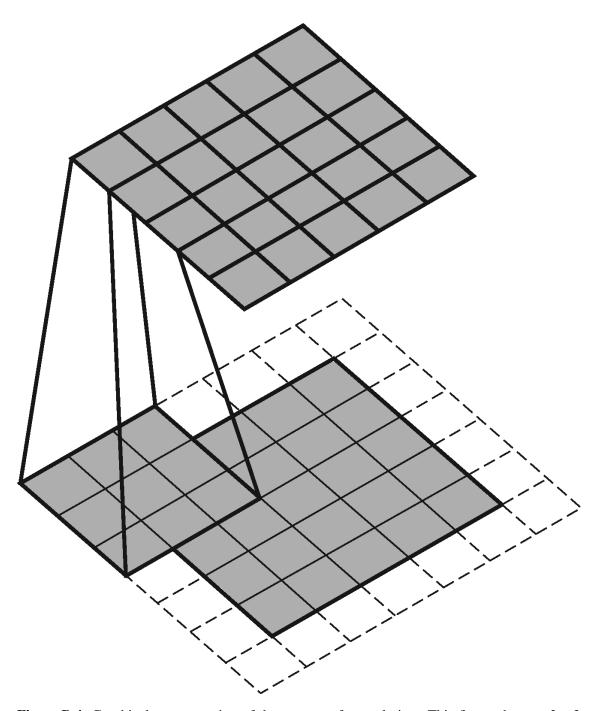


**Figure B.2:** Graphical representation of a number of activation functions. In the top left of this figure can be seen the sigmoid function. In the top right of this figure the tanh activation function can be seen. In the middle left of this figure the Rectified Linear Unit (ReLU) activation can be seen. In the bottom right of this figure the Exponential Linear Unit (ELU) activation can be seen. In the bottom right of this figure the leaky ReLU can be seen.



**Figure B.3:** Graphical representation of the Scaled Exponential Linear Unit (SELU) activation function. This figure shows a scaled version of the ELU activation function showed in figure B.2. The same advantages and disadvantages as with ELU are apparent here. However, as an additional advantage if the input data to this activation are normalised then the output of the activation will also be normalised.

The NN used was a modified U-Net (Weng and Zhu 2021), with seven down/upsampling stages. Each down/upsampling stage consisted of two convolutional layers (with two, four, eight, 16, 32, 64, or 128 channels, depending on depth). Followed by, either a split strided convolution and maxpooling layer (with the result concatenated), or a trilinear upsampling layer. Edge padding, group normalisation (Wu and He 2020), MISH activation (Misra 2019), and spatial dropout, were used with every convolutional layer. An example of some activation functions can be seen in figure B.2 and a self-regularising activation function, which was considered for this work, can be seen in figure B.3. A graphical representation of the process of convolution can be seen in figure B.4. Data was edge padded to the nearest power of two, and the input data had Gaussian noise summed to it. Both input and label data were standardised. Mean Squared Error (MSE) and TV were used as the loss function. AdamW (Loshchilov and Hutter 2019) was used as optimiser. Training continued for all methods until the gradient of the loss function, over a window of previous results, reduced below a threshold. Parameters were tuned using a grid search.



**Figure B.4:** Graphical representation of the process of convolution. This figure shows a  $3 \times 3$  convolution kernel operating on a padded image at the bottom of the figure to create the new image at the top of the figure. For a Convolutional Neural Network (CNN) each value in the convolution kernel is determined using weights and biases, like in a fully connected NN. However, rather than there being a weight and bias for every element in the input data there is only one for the kernel. Therefore, the number of parameters is significantly reduced and if the NN is fully convolutional then the input size does not have to be fixed.

Two training regimes were explored, one where each time point was treated independently, and another, where the model weights were saved and then independently updated on each time point (the mean of the new model's weights was taken for the next iteration).

As an aside, in order to justify further the selection of activation function, here are the advantages and disadvantages of the activations seen in figure B.2. The sigmoid function compresses the range of the input to be between zero and one. However, values which are much outside of the range of zero to one can have the contrast between them significantly reduced. For instance values with a magnitude of difference could become almost indistinguishable once transformed. Furthermore, because this activation is bounded both above and below the sigmoid function suffers with the vanish gradient problem (where the gradient becomes so small that values are updated very slowly or not at all). However, the sigmoid function does not suffer with the exploading gradient problem (where the gradient becomes so large that optimisation is unstable). The ReLU activation is very quick to compute. However, this activation function suffers when values become less than zero, there is no gradient here and as such the values cannot change, this is the dead ReLU problem. The ELU activation is similar to the ReLU activation, in contrast this activation does not suffer from the dead ReLU problem. However, because of the exponent this activation function is slow to compute. For leaky ReLU rather than setting all negative values to zero the negative region of the activation function is set to a scaled version of the input. This means that leaky ReLU does not suffer from the dead ReLU problem.

### **B.3.2** PET Acquisition Simulation and Image Reconstruction

A series of dynamic scans, following the clinical Dynamic Whole Body (DWB)-PET protocol, were generated using the XCAT phantom (Segars et al. 2010). Patient-derived kinetic parameters were assigned to 64 tissues, three tumours of 1.0 cm diameter in the left lung, and three tumours of 2.5 cm, 2.0 cm, and 1.0 cm diameter in the liver. An input function for Fluorine-18 Fludeoxyglucose ([<sup>18</sup>F]-FDG), taken from (Långsjö et al. 2004), was used to simulate Time Activity Curves (TACs) to create dynamic images.

PET acquisitions were simulated (and reconstructed) using Software for Tomographic Image Reconstruction (STIR) (Thielemans et al. 2012) through Synergistic Im-

age Reconstruction Framework (SIRF) (Ovtchinnikov et al. 2020). Non-Time of Flight (Non-TOF) sinogram data were simulated, using resolution modelling (using a 6.0 cm Full Width at Half Maximum (FWHM) Gaussian filter). Randoms and scatter were not included. Poisson noise was added.

Finally, all data sets were reconstructed using 10 iterations with 17 subsets of Ordered Subset Expectation Maximisation (OSEM) (Hudson and Larkin 1994).

### **B.3.3** Kinetic Modelling

Indirect Patlak estimation was used to generate  $K_i$  and intercept images (Patlak et al. 1983). Volume of Interests (VOIs) were defined for the three lesions and two background regions (lung and liver), and the mean values of  $K_i$  and  $V_d$  were calculated in each VOI. The uncertainties of the parameters were estimated as follows. Normally distributed noise was added to the dynamic images, with standard deviation given by the DIP-uncertainty, and Patlak analysis was performed. The procedure was repeated for 10 noise realisations, and the standard deviation of the  $K_i$  and  $V_d$  parameters were calculated.

### **B.3.4** Evaluation

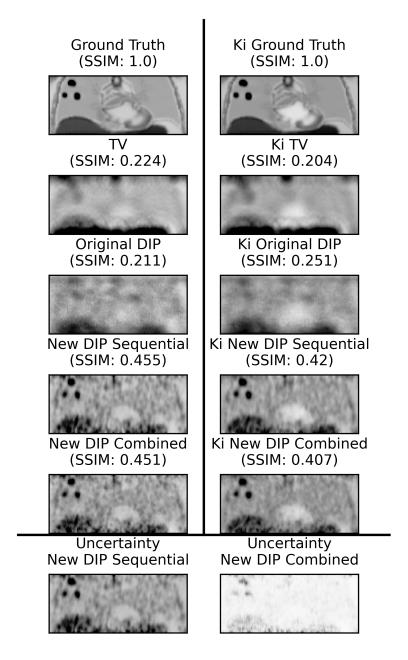
In addition to the denoising performed above in section B.3.1, data were also denoised using TV, and the DIP method presented in (Gong et al. 2019).

Comparisons used included: A visual analysis, SSIM to the ground truth (Wang and Bovik 2009),  $K_i$  values, a TAC through a lesion, a profile over a lesion, and Standard Uptake Value (SUV)<sub>max</sub> and SUV<sub>peak</sub> (defined following European Association of Nuclear Medicine (EANM) guidelines (Boellaard et al. 2015)).

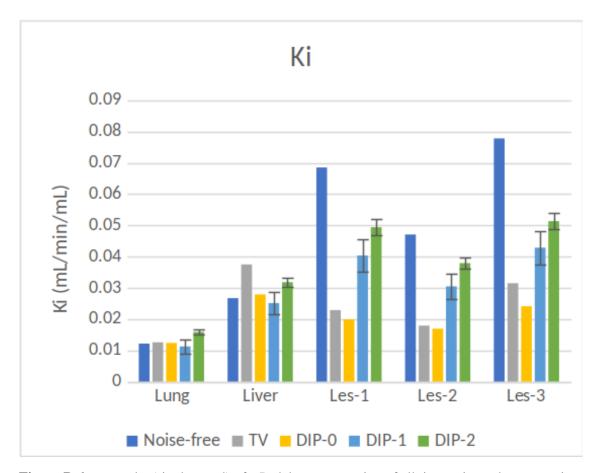
### **B.4** Results

A visual comparison of the reconstructed images (see figure B.5), shows that both of the new DIP methods perform comparably, if not for a slight reduction in noise in the combined case. Whereas, the TV and original DIP implementations appear to have struggled with over smoothing, reducing the contrast of the lesions, and introducing some edge artefacts. The uncertainty of the combined method can be seen reduced compared to the sequential method.

A comparison of  $K_i$  values across multiple lesions (see figure B.6), shows that the



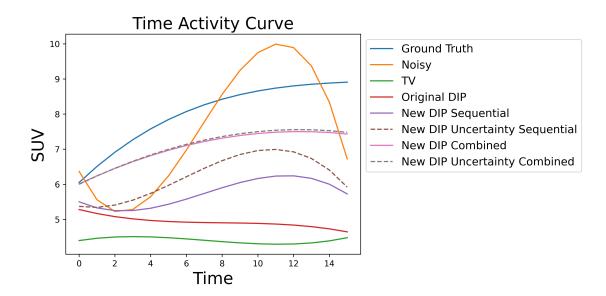
**Figure B.5:** First column contains, a visual analysis between the ground truth and denoised results (taken for the last time point, plus SSIM to the ground truth), and the second column contains, the  $K_i$  results (all voxels in a coronal view) of a Patlak reconstruction of all time points (plus SSIM to the ground truth), for the ground truth, and data denoised using, TV, the implementation of DIP from (Gong et al. 2019), and our new implementation of DIP, trained sequentially and combined (taken for the lung FOV). Last row contains, uncertainty volumes, for the data denoised using our new implementation of DIP, trained sequentially and combined (taken for the last time point of the lung FOV). Colour map ranges are consistent for all images in each section.



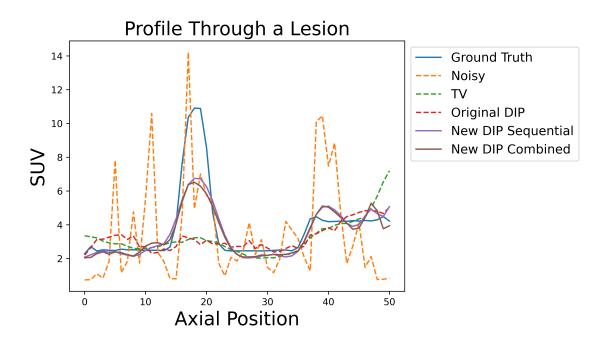
**Figure B.6:**  $K_i$  results (single voxel) of a Patlak reconstruction of all time points, plus uncertainty where applicable, for the ground truth, and data denoised using, TV, the implementation of DIP from (Gong et al. 2019), and our new implementation of DIP, trained sequentially and combined.

**Table B.1:** Comparison of  $SUV_{max}$  and  $SUV_{peak}$ , for the ground truth, the original noisy data, and this data denoised using, TV, the implementation of DIP from (Gong et al. 2019), and our new implementation of DIP, trained sequentially and combined (taken for the last time point of the lung FOV).

SUV	Max	Peak
<b>Ground Truth</b>	12.3	9.53
Noisy	21.5	5.99
TV	3.28	6.24
Original DIP	3.90	6.93
New DIP Sequential	9.19	8.02
New DIP Combined	9.27	8.21



**Figure B.7:** A TAC through a lesion, fit as a third order polynomial regression, with weighting using uncertainty (where available), for the ground truth, the original noisy data, and this data denoised using, TV, the implementation of DIP from (Gong et al. 2019), and our new implementation of DIP, trained sequentially and combined, both with and without uncertainty (taken for the lung FOV).



**Figure B.8:** A profile through a lesion, in the Superior Inferior (SI) direction, for the ground truth, the original noisy data, and this data denoised using, TV, the implementation of DIP from (Gong et al. 2019), and our new implementation of DIP, trained sequentially and combined (taken for the last time point of the lung FOV).

new DIP combined method most often estimates the greatest magnitude of  $K_i$  value, which is usually closest to the ground truth. The new DIP sequential is slightly less accurate (also with greater uncertainty), however, is more accurate than the TV and original DIP implementations (which consistently significantly underestimate  $K_i$ ).

The overall shape of the TAC (see figure B.7) for the new DIP combined method appears, most similar to the ground truth, however, with a slight reduction in quantification. The sequential method is less accurate, but still more so than both the original DIP and TV methods. There is significant variation in the noise TAC, somewhat masked by the regression, however, its shape is still least like the ground truth. Adding uncertainty appears to have improved the TAC of the new DIP sequential method, however, the uncertainty of the combined method is less, and as such, the inclusion of uncertainty has not affected results significantly.

The peak of the profile (see figure B.8) for both new DIP methods is comparable, and greater than both the original DIP and TV methods. The peak of the noise profile is greater than all other methods, including the ground truth, however, this is not necessarily beneficial, as can be seen by the rest of the profile not closely following the ground truth (it undulates unpredictably). The profile for both new DIP methods are significantly smoother, and more closely follow the ground truth.

SUV (and SSIM) results confirm the above (see table B.1).

### **B.5** Discussion and Conclusions

Evaluation indicated that the new DIP method, particularly when trained combined, provided images with less noise and more quantitative accuracy than other methods. The combined method had lower uncertainty.

Results presented here were obtained on a single bed position. Initial evaluation on a bed position, centred on the liver, indicated that parameter fine-tuning, depending on the distribution and count level, will be beneficial. Evaluation with patient data will follow.

The uncertainty estimates produced by the NN need to be validated by comparison with results obtained from repeated noise realisations.

In the future, research will focus on the application of the method to domains other than dynamic PET, where 4D data exists, such as motion correction.

## **Appendix C**

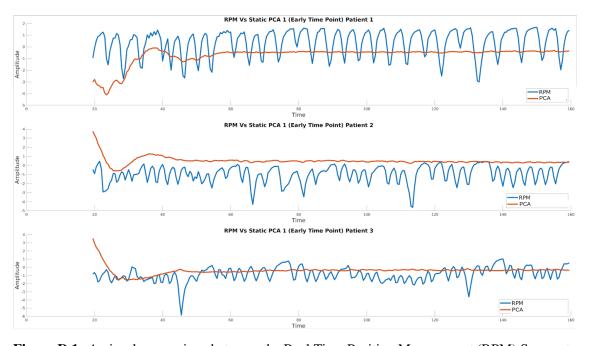
# PCA Data Driven Surrogate Signal Extraction Methods for Dynamic PET Full Overview



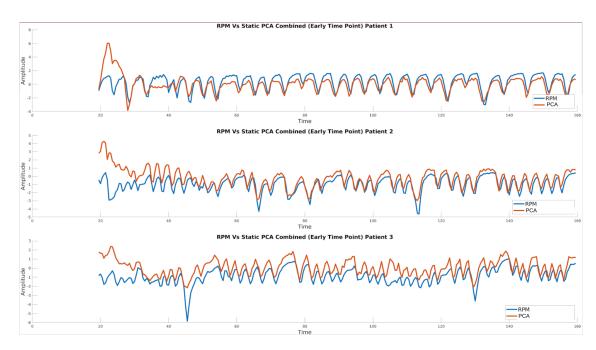
**Figure C.1:** A diagram showing an overview of the possible ways in which the method could be executed.

## **Appendix D**

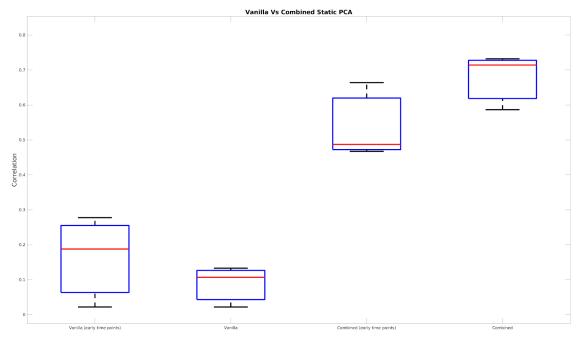
## PCA Data Driven Surrogate Signal Extraction Methods for Dynamic PET Results With and Without Pre- and



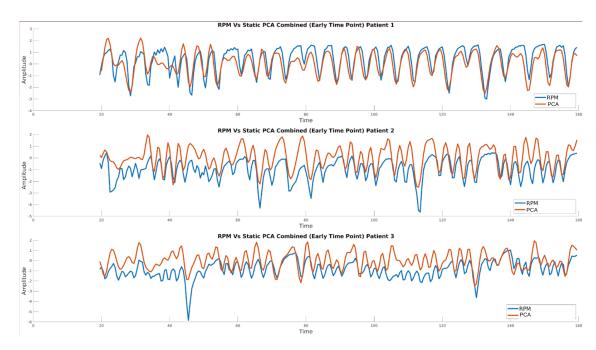
**Figure D.1:** A visual comparison between the Real Time Position Management (RPM) Surrogate Signal (SS) and the SS from the static Principal Component Analysis (PCA) method using only the first Principal Component (PC) for three patients between 20 s and 160 s. The three patients shown here are the ones on which parameters for the method were optimised.



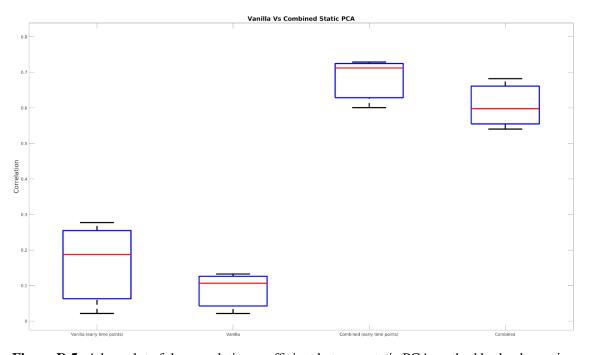
**Figure D.2:** A visual comparison between the RPM SS and the SS from the static PCA method by combining the first 20 PCs for three patients between 20 s and 160 s. The three patients shown here are the ones on which parameters for the method were optimised. Here no pre- or post-processing is used.



**Figure D.3:** A box plot of the correlation coefficient between static PCA method both when using only the first PC and by combining the first 20 PCs for all patients between 20 s and 160 s and for the entire acquisition. The parameters used here were optimised for and frozen based on the results shown in other diagrams on three patients. Here no preor post-processing is used.



**Figure D.4:** A visual comparison between the RPM SS and the SS from the static PCA method by combining the first 20 PCs for three patients between 20 s and 160 s. The three patients shown here are the ones on which parameters for the method were optimised. Here pre- and post-processing are used.



**Figure D.5:** A box plot of the correlation coefficient between static PCA method both when using only the first PC and by combining the first 20 PCs for all patients between 20 s and 160 s and for the entire acquisition. The parameters used here were optimised for and frozen based on the results shown in other diagrams on three patients. Here pre- and post-processing are used.

## **Post-Processing**

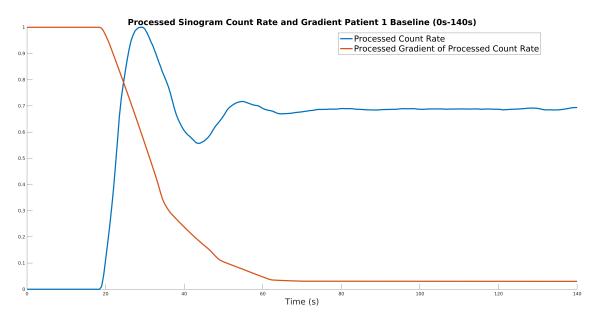
Results presented here are only for the static PCA case. For results presented here parameters were selected on a subset of the data (specifically three randomly selected scans) before being applied across the entire data set. A plot of the SS for static PCA using only one PC for three patients at early time points can be seen in figure D.1. It can be observed in this example that only using the static PCA method and one PC does not give satisfactory results at early time points. A plot of the SS for static PCA using a combination of the first 20 PCs for three patients at early time points can be seen in figure D.2, here no pre- or post-processing is used. Here for all patients from very early in the acquisition it can be seen that the method gives comparable results to the RPM. A plot of the SS for static PCA using a combination of the first 20 PCs for three patients at early time points can be seen in figure D.4, here pre- and post-processing are used. The parameters for the processing were optimised solely to improve the correlation coefficient for early time points, as such it can be seen that some inter-window variation has been removed (from parallel compression).

A box plot of the correlation coefficient of the SS for static PCA using only one PC compared to the RPM and the correlation coefficient of the SS for static PCA using a combination of the first 20 PCs compared to the RPM for all patients at both early time points and all time points and can be seen in figure D.3, here no pre- or post-processing is used. Here the improvement by incorporating multiple PCs is most apparent. A box plot of the correlation coefficient of the SS for static PCA using only one PC compared to the RPM and the correlation coefficient of the SS for static PCA using a combination of the first 20 PCs compared to the RPM for all patients at both early time points and all time points and can be seen in figure D.5, here pre- and post-processing are used. The parameters for the processing were optimised solely to improve the correlation coefficient for early time points, as such it can be seen that the correlation coefficient drops from no processing to processing for the entire acquisition, it can be assumed this effect is similar to minimising variance at the expense of bias or vice versa.

### Appendix E

# PCA Data Driven Surrogate Signal Extraction Methods for Dynamic PET Parallel Compression

This appendix contains a detailed description of what is presented briefly in section 5.4.2.1. This method attempts to limit the dynamic range of a signal proportionally, in time, to how much that signal is expected to be affected by artefacts. The signal is duplicated



**Figure E.1:** This figure shows an example of the processed count rate and processed gradient of the processed count rate, between 0 s and 140 s, which are used as part of the parallel compression post-processing step. These have been taken for the first acquisition of patient one. Here the processed count rate is normalised between zero and one by subtracting by the minimum and dividing by the maximum value of the signal.

### Algorithm 11: Parallel Compression

```
Data: timeSeriesSinograms, respiratorySignal, windowSize
  Result: respiratorySignal
1 for index in length of timeSeriesSinograms - windowSize do
      currentRespiratorySignal = respiratorySignal for data between index and
3
       index + windowSize
4
      currentRespiratorySignal = currentRespiratorySignal - mean of
5
       currentRespiratorySignal
                                            currentRespiratorySignal
      currentRespiratorySignal = -
6
                                  standard deviation of currentRespiratorySignal
7
      windowSignal = fill with Not a Numbers (NaNs) to index
8
      windowSignal append currentRespiratorySignal
      windowSignal append NaNs to length of timeSeriesSinograms
10
11
      respiratorySignals append windowSignal
12
13
14 end
15
  lowDynamicRangeSignal = mean of signals ignoring NaNs
17
18 weighting = get weighting from timeSeriesSinograms
19
20 lowDynamicRangeSignal = lowDynamicRangeSignal \times weighting
21 respiratorySignal = respiratorySignal \times (1 - weighting)
23 respiratorySignal = respiratorySignal + lowDynamicRangeSignal
```

into two channels, one of these channels is then split temporally into a series of small moving windows. These windows are 20 s in size and overlap by their length minus one (the stride size is one). The values in each window are then standardised independently, the windows are averaged back together, before being combined with the unadulterated channel following the weighting seen in figure E.1. An example of the algorithm can be seen in the algorithm 28.

The channels are combined following the gradient of the count rate. An example of the signals involved can be seen in figure E.1, and an example of the algorithm can be seen in the algorithm 28.

• Firstly, the count rate over time is found by summing all elements of each sinogram.

```
Algorithm 12: Extract Parallel Compression Weighting
```

```
Data: timeSeriesSinograms
  Result: processedGradientOfProcessedCountRate
1 for sinogram in timeSeriesSinograms do
     countRate append sum of sinogram
3 end
5 processedCountRate = moving average of countRate
7 gradientOfProcessedCountRate = gradient of processedCountRate
9 processedGradientOfProcessedCountRate = absolute of
   processedGradientOfProcessedCountRate
11 processedGradientOfProcessedCountRate = reverse of
   processedGradientOfProcessedCountRate
12 currentMaximumValue = 0.0
14 for value in processedGradientOfProcessedCountRate do
      if value >currentMaximumValue then
15
         currentMaximumValue = value
16
      else
17
         value = NaN
18
      end
19
20 end
21
22 processedGradientOfProcessedCountRate = reverse of
   processedGradientOfProcessedCountRate
23 linear interpolate NaNs in processedGradientOfProcessedCountRate
25 processedGradientOfProcessedCountRate = centred moving average of
   processedGradientOfProcessedCountRate
26
27 processedGradientOfProcessedCountRate =
   processedGradientOfProcessedCountRate - minimum of
   processedGradientOfProcessedCountRate
28 processedGradientOfProcessedCountRate =
          processedGradientOfProcessedCountRate
   maximum of processedGradientOfProcessedCountRate
```

A moving average smoothing is then applied to this processed count rate signal with a size of 5 s.

- Secondly, the gradient of the processed count rate is taken. The absolute of this signal is used, this is because we are interested in the change in gradient (but not the direction). Next the processed gradient of the processed count rate signal is iterated over from its last value to its first. If the value decreases, it is replaced by interpolation, using the previous value and the next value that is larger than the previous value. Finally a (centred) moving average smoothing is applied to the signal with a size of 20 s, and the signal is normalised between zero and one by subtracting by the minimum and dividing by the maximum value of the signal.
- Thirdly, the low and high dynamic range respiratory signals are summed, using the magnitude of the processed gradient of the processed count rate signal as a weighting. At time intervals where the magnitude of the processed count rate signal is greater, more of the compressed/standardised signal is summed, compared to the non-compressed/non-standardised signal. The total weighting is always equal to one, to maintain the scale. To achieve this, the low dynamic range respiratory signal is multiplied directly by the processed gradient of the processed count rate signal, and the high dynamic range respiratory signal is multiplied by one minus the processed count rate signal (before both respiratory signals are summed together).

### Appendix F

# A Bayesian Neural Network-Based Method for the Extraction of a Metabolite Corrected Arterial Input Function from Dynamic PBR28 PET

### F.1 Abstract

In Positron Emission Tomography (PET), arterial sampling and metabolite correction are prerequisites for the gold-standard measurement of values like the Volume of Distribution ( $V_T$ ), often necessary for the full quantification of radioligand binding. However, the invasiveness and technical demands of these procedures limit their application in both research and clinical PET studies. Machine learning approaches have been explored to predict  $V_T$  from PET images, but their integration in clinical routine is limited by their lack of transparency or thorough evaluation. Here we propose a Bayesian Neural Network (NN) to estimate the Arterial Input Function (AIF), while also outputting its prediction uncertainty, 1) directly from the entire dynamic PET images (Neural Network Auto-Encoder Input Function (NN-AEIF)), 2) from an Image Derived Input Function (IDIF) (Neural Network Image Derived Input Function (NN-IDIF)) and, as a sensitivity measure, 3) from the un-corrected plasma curve (Neural Network Arterial Input Function (NN-AIF)). All methods, applied on [ $^{11}$ C]-Peripheral Benzodiazepine Receptor ([ $^{11}$ C]-PBR28) PET data, were compared to the metabolite-corrected AIF in terms of  $V_T$ , and the

prediction uncertainty was assessed in terms of Normalised Coefficient of Variance (nCV). Overall, both NN-AEIF and NN-AIF were able to accurately predict  $V_T$ , outperforming the other methods, with NN-AEIF showing the lowest nCV (Whitehead et al. 2023a; Brusaferri et al. 2025).

### **F.2** Introduction

The  $V_T$  estimated with an AIF is utilised for quantification of many PET tracers, including [ $^{11}$ C]-PBR28. This, however, requires the concurrent measurement of the concentrations of unchanged radioligand in arterial plasma. Although insertion of an arterial catheter rarely results in clinically relevant adverse events, it is an invasive and laborious procedure.

IDIF represents a promising alternative to arterial sampling (Zanotti-Fregonara et al. 2011). However, its applicability in clinical research is hampered by several factors including the inaccuracy in the estimation of both shape and amplitude of the input function. Moreover, IDIF does not allow for radio-metabolites quantification (Sari et al. 2018). The application of machine learning is expected to improve the accuracy of predicting the AIF from PET images (Kuttner et al. 2020; Ferrante et al. 2022). While these methods have shown promising results, the vast majority of these approaches have been developed for PET tracers that do not produce radio-metabolites. Furthermore, even if the developed model shows sufficient prediction accuracy for unseen data, its applicability in the clinical setting remains questionable because of a lack of transparency or thorough evaluation (Salahuddin et al. 2022). Bayesian networks offer the significant advantage of making probabilistic predictions based on available evidence. Specifically, a Bayesian network would output uncertainty estimates in addition to the model prediction. For this reason, they have the potential to overcome the key barrier to the responsible adoption of Artificial Intelligence (AI) in clinical practice (Prabhudesai et al. 2023).

Here, we propose a Bayesian NN-based method for predicting a metabolite corrected AIF, while allowing for the estimation of uncertainty of the model's output. Specifically for the auto-encoder, although also present in the other networks, we try to enforce the low dimensional representation of the input data as disentangled and continuous. Furthermore, the network does not predict a single signal for each input. ather, it predicts a probability

density function of potential signals, which allows for the estimation of uncertainty of the model's output.

### F.3 Methods

### **F.3.1** Data Acquisition and Processing

Dynamic [ $^{11}$ C]-PBR28 PET/Magnetic Resonance (MR) images from 52 individuals (Age 55  $\pm$  16 years. Sex 27 Male, 25 Female. Genotype 32 High Affinity Binders (HABs), 20 Mixed Affinity Binders (MABs). Clinical population 12 Healthy Controls, 40 Chronic Pain patients. Injected Dose  $14.16\pm1.3$  Millicuries (mCis)) were acquired on a Siemens Biograph mMR whole-body tomograph for a time-period of 0-90 minutes post-injection. Data were pooled for multiple protocols (approved by the Partners Healthcare/Mass General Brigham Institutional Review Board) and reconstructed as in (Brusaferri et al. 2022a). All subjects had a radial artery catheter placed during the scan. Uncorrected plasma curves from blood samples were interpolated and metabolite-corrected to obtain the AIF. To further validate the proposed method, IDIF was calculated by segmenting the arterial carotid siphons using intensity thresholding of early dynamic PET frames. Data were split using ten-fold cross-validation, ensuring maximum within-variance and minimum between-variance in the training and testing sets.

All subjects received a radial artery catheter at the time of the scan. Uncorrected plasma curves from the blood samples were interpolated and metabolite-corrected to obtain the AIF. Uncorrected plasma curves were obtained from raw blood samples using linear fitting. In 16 subjects, arterial blood processing was performed using a HyperSep C18 solid extraction cartridge to separation of radio-metabolites. In 36 subjects, High Performance Liquid Chromatography (HPLC) for separation of radio-metabolites from parent radiotracer was used instead. Hill-fitted parent fractions were applied to the raw data and a radio-metabolite-corrected AIF was obtained for each subject. Previous cross-validation confirmed reliability of both pipelines allowing for the combination of both data-sets to increase statistical power of the study (Brusaferri et al. 2022a).

For validation purposes, IDIF was also calculated by segmenting the arterial carotid siphons via intensity thresholding of early dynamic PET frames. Data were split following ten-fold cross validation. The split was specifically designated to maximise the within-

variance while minimising the between-variance in training and testing sets.

### F.3.2 Neural Network Design

The method is comprised of three independent NNs.  $NN_1$  seeks to reduce the dimentionality of the input data (due to computational requirements) and extract the most relevant features.  $NN_2$  aims to extract a non-metabolite corrected signal from the low-dimensional representation output by the first network.  $NN_3$  metabolite corrects and reshapes/rescales the non-metabolite corrected signal.

All models use a novel activation function, defined as  $PSoftplus = \log(e^x + \alpha^2)$ , where x is the output from the previous layer and  $\alpha$  is a learnt parameter which is initialised as  $\alpha = 0$ . This activation function was designed as a fully differentiable drop-in replacement for PReLU (Ciuparu et al. 2020). The initial value of  $\alpha$  is selected such that the activation is initially linear (thus making the model easier to train at early iterations) and becomes non-linear as training progresses.

In order to try to enforce disentanglement and continuity of this layer, a regularisation term to enforce the orthogonality of its output was used. To promote stability of the optimisation, a regularisation term which compares the area under the curve of the prediction and true signals was used in each NN. Common to all NNs, each model was optimised using AdaBelief (Zhuang et al. 2020) with the warm-up proportion equivalent to one-tenth of the total number of epochs, after which it does not decrease. Weight decay was used to regularise against large weights. Each model was trained initially using Root Mean Square Error (RMSE), where the mean output of the model is penalised against the expected value and the standard deviation output of the model is encouraged to be close to 1. Subsequently, a second training regime was used where the negative log likelhood function is used to fine-tune the parameters. Gradient accumulation was used to allow for only one data-point to be loaded onto the Graphics Processing Unit (GPU) at any one time and for any arbitrary batch-size to be used. The batch size starts at two and increases to the size of total number of data points, where the batch size doubles as the current epoch number quadruples.

### F.3.2.1 Neural Network One Autoencoder

This network features three blocks, the downsampling block, the latent layer, and the upsampling block. The first block comprises three convolutions. The latent block is flanked on either side by two convolutions, with a variational latent layer in the middle. The upsampling block consists of a transposed convolution and two standard convolutions. Here, two downsampling and two upsampling blocks are used. The number of filters doubled or halved at each block respectively.

The input to the network is the dynamic PET images. Both the mean and standard deviation of the latent layer and the final layer are output from the model and passed onto  $NN_2$ . Both input and target data were standardised separately, based on parameters obtained from the training set. Each time frame was treated as an independent training example.

The target images were smoothed using a Gaussian filter with a Full Width at Half Maximum (FWHM) equal to three times the voxel size in each dimension. All images were padded such that the size of each dimension was equal to the next larger power of two. Both input and target images were standardised separately, based on parameters obtained from the training set.

### F.3.2.2 Neural Network Two Signal Extractor

This network consists of two blocks, the downsampling block and the fully connected block. The downsampling block follows the same structure as in NN<sub>1</sub>. The fully connected block consists solely of one fully connected layer. All time frames were used simultaneously, where the same convolutions are applied independently on each time frame before global average pooling and flattening. After flattening the clinical features were concatenated with the flattened output. Here, four downsampling blocks and eight fully connected blocks were used. The number of filters doubled and the number of units halved at each block respectively.

Each element of the clinical features was either standardised or encoded using one hot encoding depending on the nature of the feature. Because the autoencoder outputs both a mean and a standard deviation, values were sampled at the input of the network using the same *reparametrisation-trick* as for a variational autoencoder.

### F.3.2.3 Neural Network Three Metabolite Correction and Reshaping

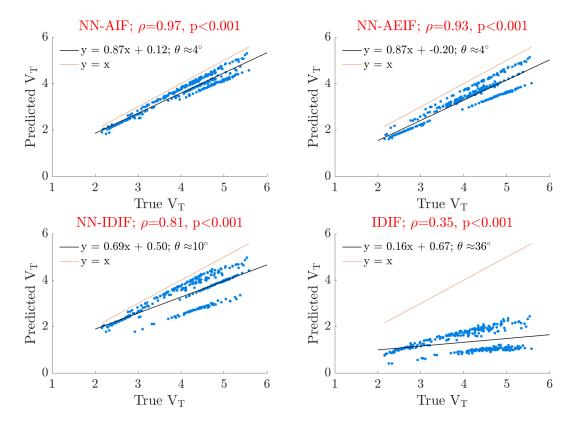
This network contains solely fully connected layers. If the network is to metabolite-correct a signal (e.g. from AIF or IDIF), it takes that signal as input together with the clinical/demographic features (age, sex, genotype, injected dose, clinical population). If the network is instead to correct a signal obtained with NN<sub>2</sub>, both the mean and the standard deviation of the uncorrected signal are input to NN<sub>3</sub>, in addition to the latent layer from NN<sub>2</sub>.

### F.3.3 Evaluation

The model was sampled 32 times resulting in multiple realisations of the estimated signal. Then,  $V_T \in \mathbb{R}^{r \times s \times b}$  were computed via the Logan graphical method, where r is the number of Region of Interests (ROIs) (r=69), s is the number of subjects (s=5) and b is the number of model samples (b=32). For both AIF and IDIF, b=1. For the NN-based methods,  $V_T$  were computed for each model realisation and then used to calculate the mean  $V_T$  and its standard deviation. Moreover, the nCV was defined as nCV = std( $Pred\ V_T$ )/ $True\ V_T$  with nCV  $\in \mathbb{R}^{r \times s}$ , where  $V_T$  is computed from the measured AIF. For each candidate signal, correlation analyses were performed on the  $V_T$  values (computed for all the ROIs and for all the subjects in the test-set) to the ones obtained with the ground truth signal (**TRUE-AIF**, see section F.3.1). To measure the accuracy of the prediction, the angle between the regression- and the identity-line was also computed, defined as  $\theta=45-\arctan(m)*180/\pi$ , with m being the slope of the regression-line. Furthermore, nCVs were averaged across ROIs for each subject of the test-set and compared via a paired t-test for each of the three NN-based methods.

### F.3.3.1 Candidate Signals

- **IDIF** generated as in section F.3.1.
- NN-IDIF metabolite-corrected input function obtained from IDIF input to NN<sub>3</sub>.
- NN-AIF metabolite-corrected AIF obtained from (uncorrected) arterial plasma input to NN<sub>3</sub>.
- NN-AEIF metabolite-corrected input function obtained from dynamic PET im-

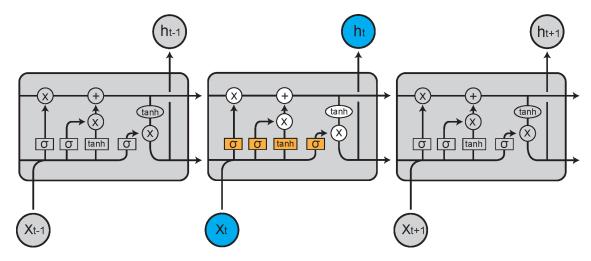


**Figure F.1:** Predicted  $V_T \in \mathbb{R}^{r \times s}$  in the test-set subjects, with r = 69 and s = 5 estimated with the four candidate signals, correlated to the True  $V_T$  (obtained with TRUE-AIF). Please note that for the NN-based methods, the displayed  $V_T$ s were averaged over all realisations.

ages input to  $NN_{1,2,3}$ .

### F.4 Results

Figure F.1 reports the correlation analyses between the predicted  $V_T$  values (obtained by the four candidate methods) and the true  $V_T$  values. For all methods, predicted  $V_T$  values positively correlate with true  $V_T$  values, with NN-AIF and IDIF showing the highest and the lowest Pearson correlation coefficient ( $\rho$ ), as well as the smallest and largest angular distance to the identity line, respectively  $\rho=0.97$  &  $\theta\approx 4^\circ$  vs  $\rho=0.35$  &  $\theta\approx 36^\circ$ . Overall, NN-AEIF outperformed NN-IDIF in terms of Pearson correlation coefficient and angular distance  $\rho=0.93$  &  $\theta\approx 4^\circ$  vs  $\rho=0.81$  &  $\theta\approx 10^\circ$ . With regard to the variance analysis, NN-AEIF outperforms both NN-AIF and NN-IDIF, showing the lowest nCV, while NN-AIF and NN-IDIF do not differ in terms of nCV values (p>0.05).



**Figure F.2:** Graphical representation of an unrolled Long Short Term Memory (LSTM). This figure shows from the left to the right the process of feeding a time series to an LSTM. This figure also shows not only the output at a given time point but also how information flows from the past to the future. In the centre of this figure an LSTM node can be seen transforming  $X_t$  to  $h_t$ . This transformation takes place taking into account information passed from the transformation of  $x_{t-1}$  to  $h_{t-1}$  etc. The flow of information is shown within the grey box. There are numerous gates which allow the LSTM to selectively remember or forget information by controlling how this flows to the transformation of  $X_{t+1}$  to  $h_{t+1}$ .

### F.5 Discussion and Conclusions

This work presents an innovative Bayesian NN-based approach for estimating the AIF from dynamic PET images and clinical variables. This approach shares similarities with previous methods developed for PET tracers that do not produce radio-metabolites, such as [18F]Fluorodeoxyglucose (FDG). In this study, additional efforts were devoted to address [11C]-PBR28 radio-metabolite correction. One of the main advantages of the proposed method is that it provides a measure of confidence in the generated signal for unseen data. Additionally, the method's modular design allows each part to be used independently. For example, in this work, metabolite correction was applied to a signal generated by a more traditional method (IDIF).

The four candidate signals were compared to the gold standard TRUE-AIF, obtained from arterial blood sampling and metabolite correction. Overall, NN-AEIF demonstrated comparable performance in terms of correlation and bias to NN-AIF, with the lowest variance of the estimated  $V_T$ , as measured by the nCV. This improved performance can potentially be explained by the larger amount of input data and the consequently

more complex model with additional parameters. Interestingly, the NN-IDIF method was able to improve on the traditional IDIF approach, as evidenced by a higher correlation coefficient and a lower angular distance from the identity line.

The proposed approach has some limitations, including the small training size, which hindered the assessment of the prediction accuracy within subsets of clinical populations in the test-set (i.e., patients vs healthy controls). In the future, the accuracy of the model could be improved through the inclusion of an attention layer either before or after the latent layer of the auto-encoder, validated through the use of an ablation study. As well as the replacement of the fully connected layers with an LSTM or transformer based approach. A graphical representation of the structure of an LSTM can be seen in figure F.2.

## **Bibliography**

- [1] A. Afshar-Oromieh, A. Malcher, M. Eder, M. Eisenhut, H. G. Linhart, B. A. Hadaschik, T. Holland-Letz, F. L. Giesel, C. Kratochwil, S. Haufe, U. Haberkorn and C. M. Zechmann. 'Pet imaging with a [68ga]gallium-labelled psma ligand for the diagnosis of prostate cancer: Biodistribution in humans and first evaluation of tumour lesions'. In: *European Journal of Nuclear Medicine and Molecular Imaging* 40.4 (Nov. 2013), pp. 486–495. ISSN: 16197089. DOI: 10.1007/s00259-012-2298-2.
- [2] Adeyemi Akintonde. 'Surrogate driven respiratory motion model derived from CBCT projection data'. In: *Doctoral thesis, UCL (University College London)*. (Apr. 2021).
- [3] Robert E. Alvarez and Albert MacOvski. 'Energy-selective reconstructions in X-ray computerised tomography'. In: *Physics in Medicine and Biology* 21.5 (1976), pp. 733–744. ISSN: 00319155. DOI: 10.1088/0031-9155/21/5/002.
- [4] Wenjia Bai and Michael Brady. 'Regularized B-spline deformable registration for respiratory motion correction in PET images'. In: *Physics in Medicine and Biology* 54.9 (2009), pp. 2719–2736. ISSN: 00319155. DOI: 10.1088/0031-9155/54/9/008.
- [5] Wenjia Bai and Michael Brady. 'Motion correction and attenuation correction for respiratory gated PET images'. In: *IEEE Transactions on Medical Imaging* 30.2 (Feb. 2011), pp. 351–365. ISSN: 02780062. DOI: 10.1109/TMI.2010. 2078514.
- [6] D L Bailey, P E Townsend D.W. and Valk and M N Maisey. 'Positron Emission Tomography'. In: *Basic Science, Springer* (2006).

- [7] Dale L. Bailey. 'Transmission scanning in emission tomography'. In: *European Journal of Nuclear Medicine* 25.7 (1998), pp. 774–787. ISSN: 03406997. DOI: 10.1007/s002590050282.
- [8] Eric Bardinet, Laurent D. Cohen and Nicholas Ayache. 'Tracking and motion analysis of the left ventricle with deformable superquadrics'. In: *Medical Image Analysis* 1.2 (1996), pp. 129–149. ISSN: 13618415. DOI: 10.1016/S1361-8415(96)80009-0.
- [9] Ottavia Bertolli. 'Data-Driven methods for respiratory signal detection in Positron Emission Tomography'. PhD thesis. University College London, Apr. 2018.
- [10] Ottavia Bertolli, Simon Arridge, Charles W. Stearns, Scott D. Wollenweber, Brian F. Hutton and Kris Thielemans. 'Data Driven Respiratory signal detection in PET taking advantage of time-of-flight Data'. In: 2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop, NSS/MIC/RTSD 2016. Vol. 2017-Janua. Institute of Electrical and Electronics Engineers Inc., Oct. 2017. ISBN: 9781509016426. DOI: 10. 1109/NSSMIC.2016.8069426.
- [11] Ottavia Bertolli, Simon Arridge, Scott D. Wollenweber, Charles W. Stearns, Brian F. Hutton and Kris Thielemans. 'Sign determination methods for the respiratory signal in data-driven PET gating'. In: *Physics in Medicine and Biology* 62.8 (Mar. 2017), pp. 3204–3220. ISSN: 13616560. DOI: 10.1088/1361-6560/aa6052.
- [12] V. Bettinardi, E. De Bernardi, L. Presotto and M. C. Gilardi. 'Motion-tracking hardware and advanced applications in PET and PET/CT'. In: *PET Clinics* 8.1 (Jan. 2013), pp. 11–28. ISSN: 18799809. DOI: 10.1016/j.cpet.2012.09.008.
- [13] V. Bettinardi, L. Presotto, E. Rapisarda, M. Picchio, L. Gianolli and M. C. Gilardi. 'Physical Performance of the new hybrid PETCT Discovery-690'. In: *Medical Physics* 38.10 (2011), pp. 5394–5411. ISSN: 00942405. DOI: 10.1118/1.3635220.

- [14] Thomas Beyer, Gerald Antoch, Todd Blodgett, Lutz F. Freudenberg, Tim Akhurst and Stephan Mueller. 'Dual-modality PET/CT imaging: The effect of respiratory motion on combined image quality in clinical oncology'. In: *European Journal of Nuclear Medicine and Molecular Imaging* 30.4 (Apr. 2003), pp. 588–596. ISSN: 16197070. DOI: 10.1007/s00259-002-1097-6.
- [15] Ander Biguri, Fotis Kotasidis, Alexander C. Whitehead, Irene Burger, Brian F. Hutton and Kris Thielemans. 'Systematic Evaluation of the Impact of Involuntary Motion in Whole Body Dynamic PET'. In: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference Record, NSS/MIC 2021 and 28th International Symposium on Room-Temperature Semiconductor Detectors, RTSD 2022. Institute of Electrical and Electronics Engineers Inc., 2021. ISBN: 9781665421133. DOI: 10.1109/NSS/MIC44867.2021.9875689.
- Ronald Boellaard, Roberto Delgado-Bolton, Wim J.G. Oyen, Francesco Giammarile, Klaus Tatsch, Wolfgang Eschner, Fred J. Verzijlbergen, Sally F. Barrington, Lucy C. Pike, Wolfgang A. Weber, Sigrid Stroobants, Dominique Delbeke, Kevin J. Donohoe, Scott Holbrook, Michael M. Graham, Giorgio Testanera, Otto S. Hoekstra, Josee Zijlstra, Eric Visser, Corneline J. Hoekstra, Jan Pruim, Antoon Willemsen, Bertjan Arends, Jörg Kotzerke, Andreas Bockisch, Thomas Beyer, Arturo Chiti and Bernd J. Krause. 'FDG PET/CT: EANM procedure guidelines for tumour imaging: version 2.0'. In: *European Journal of Nuclear Medicine and Molecular Imaging* 42.2 (2015), pp. 328–354. ISSN: 16197089. DOI: 10.1007/s00259-014-2961-x.
- [17] Alexandre Bousse, Ottavia Bertolli, David Atkinson, Simon Arridge, Sébastien Ourselin, Brian F. Hutton and Kris Thielemans. 'Maximum-likelihood joint image reconstruction and motion estimation with misaligned attenuation in TOF-PET/CT'. In: *Physics in Medicine and Biology* 61.3 (Feb. 2016), pp. L11–L19. ISSN: 13616560. DOI: 10.1088/0031–9155/61/3/L11.
- [18] Alexandre Bousse, Ottavia Bertolli, David Atkinson, Simon Arridge, Sébastien Ourselin, Brian F. Hutton and Kris Thielemans. 'Maximum-likelihood joint image reconstruction/motion estimation in attenuation-corrected respiratory gated

- PET/CT using a single attenuation map'. In: *IEEE Transactions on Medical Imaging* 35.1 (Jan. 2016), pp. 217–228. ISSN: 1558254X. DOI: 10.1109/TMI. 2015.2464156.
- [19] Ludovica Brusaferri. 'Improving quantification in non-TOF 3D PET / MR by incorporating photon energy information'. PhD thesis. Mar. 2020, p. 201.
- [20] Ludovica Brusaferri, Zeynab Alshelh, Daniel Martins, Minhae Kim, Akila Weerasekera, Hope Housman, Erin J. Morrissey, Paulina C. Knight, Kelly A. Castro-Blanco, Daniel S. Albrecht, Chieh En Tseng, Nicole R. Zürcher, Eva Maria Ratai, Oluwaseun Akeju, Meena M. Makary, Ciprian Catana, Nathaniel D. Mercaldo, Nouchine Hadjikhani, Mattia Veronese, Federico Turkheimer, Bruce R. Rosen, Jacob M. Hooker and Marco L. Loggia. 'The pandemic brain: Neuroinflammation in non-infected individuals during the COVID-19 pandemic'. In: Brain, behavior, and immunity 102 (May 2022), pp. 89–97. ISSN: 1090-2139. DOI: 10.1016/J.BBI.2022.02.018.
- [21] Ludovica Brusaferri, Alexandre Bousse, Elise C. Emond, Richard Brown, Yu-Jung Tsai, David Atkinson, Sebastien Ourselin, Charles C. Watson, Brian F. Hutton, Simon Arridge and Kris Thielemans. 'Joint Activity and Attenuation Reconstruction From Multiple Energy Window Data With Photopeak Scatter Re-Estimation in Non-TOF 3-D PET'. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 4.4 (Mar. 2020), pp. 410–421. ISSN: 2469-7311. DOI: 10.1109/trpms.2020.2978449.
- [22] Ludovica Brusaferri, Elise C. Emond, Alexandre Bousse, Robert Twyman, Alexander C. Whitehead, David Atkinson, Sebastien Ourselin, Brian F. Hutton, Simon Arridge and Kris Thielemans. 'Detection Efficiency Modeling and Joint Activity and Attenuation Reconstruction in Non-TOF 3-D PET from Multiple-Energy Window Data'. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 6.1 (2022), pp. 87–97. ISSN: 24697311. DOI: 10.1109/TRPMS.2021.3064239.
- [23] Ludovica Brusaferri, Alexander C. Whitehead, Lucia Maccioni, Zaynab Alshelh, Mattia Veronese, Federico Turkheimer, Nicola Toschi, Matteo Ferrante, Marianna Inglese, Ciprian Catana, Enrico Grisan and Marco L. Loggia. 'Advancing Gener-

- alisable Neural Network-Based PET Quantification: A Multicenter [11C]PBR28 study'. In: 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (EMBC 2025). 2025.
- [24] Philippe P. Bruyant, Michael A. King and P. Hendrik Pretorius. 'Correction of the respiratory motion of the heart by tracking of the center of mass of thresholded projections: A simulation study using the dynamic MCAT phantom'. In: *IEEE Transactions on Nuclear Science* 49 I.5 (Oct. 2002), pp. 2159–2166. ISSN: 00189499. DOI: 10.1109/TNS.2002.803678.
- [25] Ralph A. Bundschuh, Axel Martínez-Moeller, Markus Essler, María José Martínez, Stephan G. Nekolla, Sibylle I. Ziegler and Markus Schwaiger. 'Postac-quisition detection of tumor motion in the lung and upper abdomen using list-mode PET data: A feasibility study'. In: *Journal of Nuclear Medicine* 48.5 (May 2007), pp. 758–763. ISSN: 01615505. DOI: 10.2967/jnumed.106.035279.
- [26] Florian Büther, Iris Ernst, James Hamill, Hans T. Eich, Otmar Schober, Michael Schäfers and Klaus P. Schäfers. 'External radioactive markers for PET data-driven respiratory gating in positron emission tomography.' In: *European journal of nuclear medicine and molecular imaging* 40.4 (Dec. 2013), pp. 602–614. ISSN: 16197089. DOI: 10.1007/s00259-012-2313-7.
- [27] Florian Büther, Judson Jones, Robert Seifert, Lars Stegger, Paul Schleyer and Michael Schäfers. 'Clinical evaluation of a data-driven respiratory gating algorithm for whole-body PET with continuous bed motion'. In: *Journal of Nuclear Medicine* 61.10 (Oct. 2020), pp. 1520–1527. ISSN: 2159662X. DOI: 10.2967/jnumed.119.235770.
- [28] Paulo R.R.V. Caribé, M. Koole, Yves D'Asseler, Timothy W. Deller, K. Van Laere and S. Vandenberghe. 'NEMA NU 2–2007 performance characteristics of GE Signa integrated PET/MR for different PET isotopes'. In: *EJNMMI Physics* 6.1 (Dec. 2019), pp. 1–13. ISSN: 21977364. DOI: 10.1186/s40658-019-0247-x.
- [29] Jonathan P.J. Carney, David W. Townsend, Vitaliy Rappoport and Bernard Bendriem. 'Method for transforming CT images for attenuation correction

- in PET/CT imaging'. In: *Medical Physics* 33.4 (2006), pp. 976–983. ISSN: 00942405. DOI: 10.1118/1.2174132.
- [30] J. W. Cates and C. S. Levin. 'Evaluation of a TOF-PET detector design that achieves less than or equal to 100 ps coincidence time resolution'. In: 2017 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2017 Conference Proceedings (2018). DOI: 10.1109/NSSMIC.2017.8532998.
- [31] Chung Chan, Xiao Jin, Edward K. Fung, Mika Naganawa, Tim Mulnix, Richard E. Carson and Chi Liu. 'Event-by-event respiratory motion correction for PET with 3D internal-1D external motion correlation'. In: *Medical Physics* 40.11 (2013), p. 112507. ISSN: 00942405. DOI: 10.1118/1.4826165.
- [32] Chung Chan, John Onofrey, Yiqiang Jian, Mary Germino, Xenophon Papademetris, Richard E. Carson and Chi Liu. 'Non-Rigid Event-by-Event Continuous Respiratory Motion Compensated List-Mode Reconstruction for PET'. In: *IEEE Transactions on Medical Imaging* 37.2 (Feb. 2018), pp. 504–515. ISSN: 1558254X. DOI: 10.1109/TMI.2017.2761756.
- [33] R. W. Chan and P. B. Goldsmith. 'Psychovisually-based image quality evaluator for JPEG images'. In: *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics* 2 (2000), pp. 1541–1546. ISSN: 08843627. DOI: 10.1109/ICSMC.2000.886075.
- [34] Simon R. Cherry, Terry Jones, Joel S. Karp, Jinyi Qi, William W. Moses and Ramsey D. Badawi. 'Total-body PET: Maximizing sensitivity to create new opportunities for clinical research and patient care'. In: *Journal of Nuclear Medicine* 59.1 (Jan. 2018), pp. 3–12. ISSN: 2159662X. DOI: 10.2967/jnumed.116.184028.
- [35] Andrei Ciuparu, Adriana Nagy-Dăbâcan and Raul C. Mureşan. 'Soft++, a multi-parametric non-saturating non-linearity that improves convergence in deep neural architectures'. In: *Neurocomputing* 384 (Apr. 2020), pp. 376–388. ISSN: 0925-2312. DOI: 10.1016/J.NEUCOM.2019.12.014.
- [36] Maurizio Conti and Lars Eriksson. 'Physics of pure and non-pure positron emitters for PET: A review and a discussion'. In: *EJNMMI Physics* 3.1 (Dec. 2016). ISSN: 21977364. DOI: 10.1186/s40658-016-0144-5.

- [37] Mohammad Dawood, Florian Büther, Xiaoyi Jiang and Klaus P. Schäfers. 'Respiratory motion correction in 3-D PET data with advanced optical flow algorithms'. In: *IEEE Transactions on Medical Imaging* 27.8 (Aug. 2008), pp. 1164–1175. ISSN: 02780062, DOI: 10.1109/TMI.2008.918321.
- [38] Mohammad Dawood, Thomas Kösters, Michael Fieseler, Florian Büther, Xiaoyi Jiang, Frank Wübbeling and Klaus P. Schäfers. 'Motion Correction in Respiratory Gated Cardiac PET/CT Using Multi-scale Optical Flow'. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 5242 LNCS.PART 2 (2008), pp. 155–162. ISSN: 1611-3349. DOI: 10.1007/978-3-540-85990-1{\\_}19.
- [39] Mohammad Dawood, Norbert Lang, Xiaoyi Jiang and Klaus P. Schäfers. 'Lung motion correction on respiratory gated 3-D PET/CT images'. In: *IEEE Transactions on Medical Imaging* 25.4 (Apr. 2006), pp. 476–485. ISSN: 02780062. DOI: 10.1109/TMI.2006.870892.
- [40] Michel Defrise, Ahmadreza Rezaei and Johan Nuyts. 'Time-of-flight PET data determine the attenuation sinogram up to a constant'. In: *Physics in Medicine and Biology* 57.4 (Feb. 2012), pp. 885–899. ISSN: 00319155. DOI: 10.1088/0031-9155/57/4/885.
- [41] Dominique Delbeke, R. Edward Coleman, Milton J. Guiberteau, Manuel L. Brown, Henry D. Royal, Barry A. Siegel, David W. Townsend, Lincoln L. Berland, J. Anthony Parker, Karl Hubner, Michael G. Stabin, George Zubal, Marc Kachelriess, Valerie Cronin and Scott Holbrook. 'Procedure guideline for tumor imaging with 18F-FDG PET/CT 1.0'. In: *Journal of Nuclear Medicine* 47.5 (2006), pp. 885–895. ISSN: 01615505.
- [42] A. P. Dempster, N. M. Laird and D. B. Rubin. 'Maximum Likelihood from Incomplete Data Via the EM Algorithm'. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 39.1 (1977), pp. 1–22. DOI: 10.1111/j. 2517-6161.1977.tb01600.x.

- [43] Jean Duchon. 'Splines minimizing rotation-invariant semi-norms in Sobolev spaces'. In: Springer, Berlin, Heidelberg, 1977, pp. 85–100. DOI: 10.1007/bfb0086566.
- [44] N. Efthimiou, A. C. Whitehead, M. Stockhoff, C. Thyssen, S. J. Archibald and S. Vandenberghe. 'Preliminary investigation of the impact of Axial Ring Splitting on Image Quality for the Cost Reduction of Total-Body PET'. In: 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2019. Institute of Electrical and Electronics Engineers Inc., Oct. 2019. ISBN: 9781728141640. DOI: 10.1109/NSS/MIC42101.2019.9059650.
- [45] Nikos Efthimiou, Elise Emond, Palak Wadhwa, Christopher Cawthorne, Charalampos Tsoumpas and Kris Thielemans. 'Implementation and validation of time-of-flight PET image reconstruction module for listmode and sinogram projection data in the STIR library'. In: *Physics in Medicine and Biology* 64.3 (Jan. 2019), p. 035004. ISSN: 13616560. DOI: 10.1088/1361-6560/aaf9b9.
- [46] Elise C. Emond, Alexandre Bousse, Maria MacHado, Joanna Porter, Kjell Erlandsson, Ashley M. Groves, Brian F. Hutton and Kris Thielemans. 'Respiratory Motion Correction in Dynamic PET with a Single Attenuation Map'. In: 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2019 (Oct. 2019). DOI: 10.1109/NSS/MIC42101.2019.9059813.
- [47] Elise C. Emond, Alexandre Bousse, Maria Machado, Joanna Porter, Ashley M. Groves, Brian F. Hutton and Kris Thielemans. 'Effect of attenuation mismatches in time of flight PET reconstruction'. In: *Physics in Medicine and Biology* 65.8 (Apr. 2020), p. 085009. ISSN: 13616560. DOI: 10.1088/1361-6560/ab7a6f.
- [48] Elise Claudine Emond. 'Improving Quantification of Dynamic PET / CT Biomarkers in the Fibrotic Lung for the Evaluation of Disease Progression and Treatment Effectiveness'. PhD thesis. Apr. 2020. ISBN: 3000030921.
- [49] Yusuf E. Erdi, Sadek A. Nehmeh, Tinsu Pan, Alexander Pevsner, Kenneth E. Rosenzweig, Gikas Mageras, Ellen D. Yorke, Heiko Schoder, Wendy Hsiao, Olivia D. Squire, Phil Vernon, Jonathan B. Ashman, Hassan Mostafavi, Steven M. Larson and John L. Humm. 'The CT motion quantitation of lung lesions and its

- impact on PET-measured SUVs'. In: *Journal of Nuclear Medicine* 45.8 (2004), pp. 1287–1292. ISSN: 01615505.
- [50] Hadi Fayad, Holger Schmidt, Christian Wuerslin and Dimitris Visvikis. 'Reconstruction-Incorporated Respiratory Motion Correction in Clinical Simultaneous PET/MR Imaging for Oncology Applications'. In: *Journal of Nuclear Medicine* 56.6 (June 2015), pp. 884–889. ISSN: 0161-5505. DOI: 10.2967/JNUMED.114.153007.
- [51] Hadi J. Fayad, Frederic Lamare, Catherine Cheze Le Rest, Valentino Bettinardi and Dimitris Visvikis. 'Generation of 4-Dimensional CT Images Based on 4-Dimensional PET-Derived Motion Fields'. In: *Journal of Nuclear Medicine* 54.4 (Apr. 2013), pp. 631–638. ISSN: 0161-5505. DOI: 10.2967/JNUMED.112.110809.
- [52] Tao Feng, Jizhe Wang, Youjun Sun, Wentao Zhu, Yun Dong and Hongdi Li. 'Self-gating: An adaptive center-of-mass approach for respiratory gating in PET'. In: *IEEE Transactions on Medical Imaging* 37.5 (May 2018), pp. 1140–1148. ISSN: 1558254X. DOI: 10.1109/TMI.2017.2783739.
- [53] Matteo Ferrante, Marianna Inglese, Ludovica Brusaferri, Alexander Whitehead, Marco Loggia and Nicola Toschi. 'Physically Informed Neural Network for Non-Invasive Arterial Input Function Estimation In Dynamic PET Imaging'. In: June 2022. DOI: 10.1038/s42254-021-00314-5.
- [54] Matteo Ferrante, Marianna Inglese, Ludovica Brusaferri, Alexander C. Whitehead, Lucia Maccioni, Federico E. Turkheimer, Maria A. Nettis, Valeria Mondelli, Oliver Howes, Marco L. Loggia, Mattia Veronese and Nicola Toschi. 'Physically informed deep neural networks for metabolite-corrected plasma input function estimation in dynamic PET imaging'. In: *Computer Methods and Programs in Biomedicine* 256 (Nov. 2024), p. 108375. ISSN: 0169-2607. DOI: 10.1016/J. CMPB.2024.108375.
- [55] Martin A. Fischler and Robert C. Bolles. 'Random sample consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Carto-

- graphy'. In: *Communications of the ACM* 24.6 (June 1981), pp. 381–395. ISSN: 15577317. DOI: 10.1145/358669.358692.
- [56] R. Fletcher. Practical Methods of Optimization. Chichester, West Sussex England: John Wiley & Sons, Ltd, May 2000. ISBN: 9781118723203. DOI: 10.1002/ 9781118723203.
- [57] Murray F. Freeman and John W. Tukey. 'Transformations Related to the Angular and the Square Root'. In: *The Annals of Mathematical Statistics* 21.4 (1950), pp. 607–611. ISSN: 0003-4851. DOI: 10.1214/aoms/1177729756.
- [58] Niccolo Fuin, Stefano Pedemonte, Onofrio A. Catalano, David Izquierdo-Garcia, Andrea Soricelli, Marco Salvatore, Keith Heberlein, Jacob M. Hooker, Koen Van Leemput and Ciprian Catana. 'PET/MRI in the presence of metal implants: Completion of the attenuation map from PET emission data'. In: *Journal of Nuclear Medicine* 58.5 (May 2017), pp. 840–845. ISSN: 2159662X. DOI: 10.2967/jnumed.116.183343.
- [59] Sebastian Fürst, Robert Grimm, Inki Hong, Michael Souvatzoglou, Michael E. Casey, Markus Schwaiger, Stephan G. Nekolla and Sibylle I. Ziegler. 'Motion correction strategies for integrated PET/MR'. In: *Journal of Nuclear Medicine* 56.2 (Feb. 2015), pp. 261–269. ISSN: 2159662X. DOI: 10.2967/jnumed.114. 146787.
- [60] Yarin Gal and Zoubin Ghahramani. 'Dropout as a Bayesian approximation: Representing model uncertainty in deep learning'. In: *33rd International Conference on Machine Learning, ICML 2016* 3 (June 2016), pp. 1651–1660. DOI: 10.48550/arxiv.1506.02142.
- [61] Gerhard W. Goerres, Cyrill Burger, Ehab Kamel, Burkhardt Seifert, Achim H. Kaim, Alfred Buck, Tobias C. Buehler and Gustav K. Von Schulthess. 'Respiration-induced attenuation artifact at PET/CT: Technical considerations'. In: *Radiology* 226.3 (Mar. 2003), pp. 906–910. ISSN: 00338419. DOI: 10.1148/radiol.2263011732.

- [62] Kuang Gong, Ciprian Catana, Jinyi Qi and Quanzheng Li. 'PET Image Reconstruction Using Deep Image Prior'. In: *IEEE Transactions on Medical Imaging* 38.7 (July 2019), pp. 1655–1665. ISSN: 1558254X. DOI: 10.1109/TMI.2018. 2888491.
- [63] Alexander M. Grant, Timothy W. Deller, Mohammad Mehdi Khalighi, Sri Harsha Maramraju, Gaspar Delso and Craig S. Levin. 'NEMA NU 2-2012 performance studies for the SiPM-based ToF-PET component of the GE SIGNA PET/MR system'. In: *Medical Physics* 43.5 (May 2016), pp. 2334–2343. ISSN: 00942405. DOI: 10.1118/1.4945416.
- [64] Willem Grootjans, Lioe Fee de Geus-Oei, Antoi P.W. Meeuwis, Charlotte S. van der Vos, Martin Gotthardt, Wim J.G. Oyen and Eric P. Visser. 'Amplitude-based optimal respiratory gating in positron emission tomography in patients with primary lung cancer'. In: *European Radiology* 24.12 (Nov. 2014), pp. 3242–3250. ISSN: 14321084. DOI: 10.1007/s00330-014-3362-z.
- [65] Olivier Guivarc'h, Alexandre Turzo, Dimitris Visvikis and Yves Bizais. 'Synchronization of pulmonary scintigraphy by respiratory flow and by impedance plethysmography'. In: *Medical Imaging 2004: Image Processing*. Ed. by J. Michael Fitzpatrick and Milan Sonka. Vol. 5370. SPIE, May 2004, p. 1166. DOI: 10.1117/12.536329.
- [66] Fumio Hashimoto, Hiroyuki Ohba, Kibo Ote, Akihiro Kakimoto, Hideo Tsukada and Yasuomi Ouchi. '4D deep image prior: Dynamic PET image denoising using an unsupervised four-dimensional branch convolutional neural network'. In: *Physics in Medicine and Biology* 66.1 (Jan. 2021), p. 015006. ISSN: 13616560. DOI: 10.1088/1361-6560/abcd1a.
- [67] Fumio Hashimoto, Hiroyuki Ohba, Kibo Ote, Atsushi Teramoto and Hideo Tsukada. 'Dynamic PET Image Denoising Using Deep Convolutional Neural Networks Without Prior Training Datasets'. In: *IEEE Access* 7 (2019), pp. 96594–96603. ISSN: 21693536. DOI: 10.1109/ACCESS.2019.2929230.

- [68] D. L.G. Hill, P. G. Batchelor, M. Holden and D. J. Hawkes. 'Medical image registration'. In: *Physics in Medicine and Biology* 46.3 (Mar. 2001), R1–R45. ISSN: 00319155. DOI: 10.1088/0031-9155/46/3/201.
- [69] Jeremy D.P. Hoisak, Katharina E. Sixel, Romeo Tirona, Patrick C.F. Cheung and Jean Philippe Pignol. 'Correlation of lung tumor motion with external surrogate indicators of respiration'. In: *International Journal of Radiation Oncology Biology Physics* 60.4 (2004), pp. 1298–1306. ISSN: 03603016. DOI: 10.1016/j.ijrobp.2004.07.681.
- [70] David F.C. Hsu, Ezgi Ilan, William T. Peterson, Jorge Uribe, Mark Lubberink and Craig S. Levin. 'Studies of a next-generation silicon-photomultiplier-based time-of-flight PET/CT system'. In: *Journal of Nuclear Medicine* 58.9 (2017), pp. 1511–1518. ISSN: 2159662X. DOI: 10.2967/jnumed.117.189514.
- [71] H. Malcolm Hudson and Richard S. Larkin. 'Accelerated Image Reconstruction Using Ordered Subsets of Projection Data'. In: *IEEE Transactions on Medical Imaging* 13.4 (1994), pp. 601–609. ISSN: 1558254X. DOI: 10.1109/42.363108.
- [72] B. S. Ishkhanov. *The atomic nucleus*. Vol. 67. 1. 2012, pp. 1–24. DOI: 10.3103/S0027134912010092.
- [73] Roey Izhaki. *Mixing audio: concepts, practices, and tools*. Vol. 49. 12. Focal Press, 2012, pp. 49–6905. ISBN: 978-0-240-52068-1. DOI: 10.5860/choice.49-6905.
- [74] M. W. Jacobson and J. A. Fessler. 'Joint estimation of image and deformation parameters in motion-corrected PET'. In: *IEEE Nuclear Science Symposium Conference Record*. Vol. 5. 2003, pp. 3290–3294. DOI: 10.1109/nssmic.2003. 1352599.
- [75] Xiao Jin, Chung Chan, Tim Mulnix, al -, P J Markiewicz, K Thielemans, J M Schott, L Livieratos, L Stegger, P M Bloomfield, K Schafers, D L Bailey and P G Camici. 'Rigid-body transformation of list-mode projection data for respiratory motion correction in cardiac PET'. In: *Physics in Medicine & Biology* 50.14 (July 2005), p. 3313. ISSN: 0031-9155. DOI: 10.1088/0031-9155/50/14/008.

- [76] Kevin M. Johnson, Sean B. Fain, Mark L. Schiebler and Scott Nagle. 'Optimized 3D ultrashort echo time pulmonary MRI'. In: *Magnetic Resonance in Medicine* 70.5 (2013), pp. 1241–1250. ISSN: 15222594. DOI: 10.1002/mrm.24570.
- [77] Valen E. Johnson. 'A Note on Stopping Rules in EM-ML Reconstructions of ECT Images'. In: *IEEE Transactions on Medical Imaging* 13.3 (1994), pp. 569–571.
   ISSN: 1558254X. DOI: 10.1109/42.310891.
- [78] F Jolivet, A Rezaei, G Schramm, K Schafers, M Fieseler, F Boada and J Nuyts. 'Joint estimation of activity, attenuation and motion in respiratory-gated time-of-flight PET'. In: 2022 IEEE Nuclear Science Symposium and Medical Imaging Conference Record, NSS/MIC 2022 and 29th International Symposium on Room-Temperature Semiconductor Detectors, RTSD 2022. Institute of Electrical and Electronics Engineers Inc., 2022.
- [79] Alark Joshi, Dustin Scheinost, Hirohito Okuda, Dominique Belhachemi, Isabella Murphy, Lawrence H. Staib and Xenophon Papademetris. 'Unified framework for development, deployment and robust testing of neuroimaging algorithms'. In: *Neuroinformatics* 9.1 (Mar. 2011), pp. 69–84. ISSN: 15392791. DOI: 10.1007/S12021-010-9092-8/FIGURES/12.
- [80] Faraz Kalantari, Tianfang Li, Mingwu Jin and Jing Wang. 'Respiratory motion correction in 4D-PET by simultaneous motion estimation and image reconstruction (SMEIR)'. In: *Physics in Medicine and Biology* 61.15 (July 2016), pp. 5639–5661. ISSN: 13616560. DOI: 10.1088/0031-9155/61/15/5639.
- [81] Faraz Kalantari and Jing Wang. 'Attenuation correction in 4D-PET using a single-phase attenuation map and rigidity-adaptive deformable registration:' in: *Medical Physics* 44.2 (Feb. 2017), pp. 522–532. ISSN: 00942405. DOI: 10.1002/mp. 12063.
- [82] Adam L. Kesner, Paul J. Schleyer, Florian Büther, Martin A. Walter, Klaus P. Schäfers and Phillip J. Koo. 'On transcending the impasse of respiratory motion correction applications in routine clinical imaging a consideration of a fully automated data driven motion control framework'. In: *EJNMMI Physics* 1.1 (Dec. 2014), pp. 1–11. ISSN: 21977364. DOI: 10.1186/2197-7364-1-8.

- [83] Adam Leon Kesner, Ralph Alexander Bundschuh, Nicole Christine Detorie, Magnus Dahlbom, Sibylle Ilsa Ziegler, Johannes Czernin and Daniel Hillel Silverman. 'Respiratory gated PET derived in a fully automated manner from raw PET data'. In: *IEEE Transactions on Nuclear Science* 56.3 (June 2009), pp. 677–686. ISSN: 00189499. DOI: 10.1109/TNS.2009.2016341.
- [84] Adam Leon Kesner and Claudia Kuntner. 'A new fast and fully automated software based algorithm for extracting respiratory signal from raw PET data and its comparison to other methods'. In: *Medical Physics* 37.10 (Sept. 2010), pp. 5550–5559. ISSN: 00942405. DOI: 10.1118/1.3483784.
- [85] Magdy M. Khalil. *Basic science of PET imaging*. Springer International Publishing, Jan. 2016, pp. 1–619. ISBN: 9783319400709. DOI: 10.1007/978-3-319-40070-9.
- [86] Hak Soo Kim, Jae Cheol Lee and Kyu Tae Park. 'Motion estimation method using multiple linear regression model'. In: *Visual Communications and Image Processing* '97. Ed. by Jan Biemond and Edward J. Delp III. Vol. 3024. SPIE, Jan. 1997, pp. 600–607. DOI: 10.1117/12.263272.
- [87] Diederik P. Kingma and Jimmy Lei Ba. 'Adam: A Method for Stochastic Optimization'. In: 3rd International Conference on Learning Representations, ICLR 2015 Conference Track Proceedings (Dec. 2014).
- [88] Yoshiyuki Kitamura, Shingo Baba, Takuro Isoda, Yasuhiro Maruoka, Satoshi Kawanami, Kazuhiko Himuro, Masayuki Sasaki and Hiroshi Honda. 'The Efficiency of Respiratory-gated 18F-FDG PET/CT in Lung Adenocarcinoma: Amplitude-gating Versus Phase-gating Methods.' In: *Asia Oceania journal of nuclear medicine & biology* 5.1 (2017), pp. 30–36. ISSN: 2322-5718. DOI: 10.22038/aojnmb.2016.7747.
- [89] G. J. Klein, B. W. Reutter, E. H. Botvinick, T. F. Budinger and R. H. Huesman. 'Fine-scale motion detection using intrinsic list mode PET information'. In: *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis*. 2001, pp. 71–78. DOI: 10.1109/mmbia.2001.991701.

- [90] B Kröse, B Krose, P van der Smagt and P Smagt. *An introduction to neural networks*. 1993.
- [91] Samuel Kuttner, Kristoffer Knutsen Wickstrm, Gustav Kalda, S. Esmaeil Dorraji, Montserrat Martin-Armas, Ana Oteiza, Robert Jenssen, Kristin Fenton, Rune Sundset and Jan Axelsson. 'Machine learning derived input-function in a dynamic 18F-FDG PET study of mice'. In: *Biomedical physics & engineering express* 6.1 (2020). ISSN: 2057-1976. DOI: 10.1088/2057-1976/AB6496.
- [92] Andre Z. Kyme and Roger R. Fulton. 'Motion estimation and correction in SPECT, PET and CT'. In: *Physics in Medicine & Biology* 66.18 (Sept. 2021), 18TR02. ISSN: 0031-9155. DOI: 10.1088/1361-6560/AC093B.
- [93] F. Lamare, A. Bousse, K. Thielemans, C. Liu, T. Merlin, H. Fayad and D. Visvikis. *PET respiratory motion correction: Quo vadis?* Feb. 2022. DOI: 10.1088/1361-6560/ac43fc.
- [94] F. Lamare, M. J. Ledesma Carbayo, T. Cresson, G. Kontaxakis, A. Santos,
  C. Cheze Le Rest, A. J. Reader and D. Visvikis. 'List-mode-based reconstruction
  for respiratory motion correction in PET using non-rigid body transformations'.
  In: *Physics in Medicine & Biology* 52.17 (Aug. 2007), p. 5187. ISSN: 0031-9155.
  DOI: 10.1088/0031-9155/52/17/006.
- [95] Adriaan A. Lammertsma. 'Forward to the past: The case for quantitative PET imaging'. In: *Journal of Nuclear Medicine* 58.7 (2017), pp. 1019–1024. ISSN: 2159662X. DOI: 10.2967/jnumed.116.188029.
- [96] K. Lange and R. Carson. 'EM reconstruction algorithms for emission and transmission tomography'. In: *Journal of Computer Assisted Tomography* 8.2 (1984), pp. 306–316. ISSN: 03638715.
- [97] Jaakko W. Långsjö, Elina Salmi, Kaike K. Kaisti, Sargo Aalto, Susanna Hinkka, Riku Aantaa, Vesa Oikonen, Tapio Viljanen, Timo Kurki, Martti Silvanto and Harry Scheinin. 'Effects of Subanesthetic Ketamine on Regional Cerebral Glucose Metabolism in Humans'. In: *Anesthesiology* 100.5 (May 2004), pp. 1065–1071. ISSN: 00033022. DOI: 10.1097/00000542-200405000-00006. URL: https://pubmed.ncbi.nlm.nih.gov/15114201/.

- [98] P. Lecoq. 'Pushing the Limits in Time-of-Flight PET Imaging'. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 1.6 (2017), pp. 473–485. ISSN: 24697311. DOI: 10.1109/TRPMS.2017.2756674.
- [99] Craig S. Levin and Edward J. Hoffman. 'Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution'. In: *Physics in Medicine and Biology* 44.3 (Jan. 1999), pp. 781–799. ISSN: 00319155. DOI: 10.1088/0031-9155/44/3/019.
- [100] Tianfang Li, Brian Thorndyke, Eduard Schreibmann, Yong Yang and Lei Xing. 'Model-based image reconstruction for four-dimensional PET'. In: *Medical Physics* 33.5 (2006), pp. 1288–1298. ISSN: 00942405. DOI: 10.1118/1.2192581.
- [101] Chi Liu, Adam Alessio, Larry Pierce, Kris Thielemans, Scott Wollenweber, Alexander Ganin and Paul Kinahan. 'Quiescent period respiratory gating for PET/CT'.
   In: *Medical Physics* 37.9 (2010), pp. 5037–5043. ISSN: 00942405. DOI: 10. 1118/1.3480508.
- [102] Chi Liu, Adam M. Alessio and Paul E. Kinahan. 'Respiratory motion correction for quantitative PET/CT using all detected events with internal-external motion correlation'. In: *Medical Physics* 38.5 (2011), pp. 2715–2723. ISSN: 00942405. DOI: 10.1118/1.3582692.
- [103] Jiaming Liu, Yu Sun, Xiaojian Xu and Ulugbek S. Kamilov. 'Image Restoration Using Total Variation Regularized Deep Image Prior'. In: *ICASSP*, *IEEE International Conference on Acoustics, Speech and Signal Processing Proceedings* 2019-May (May 2019), pp. 7715–7719. ISSN: 15206149. DOI: 10.1109/ICASSP.2019.8682856.
- [104] J. Llacer, E. Veklerov, L. R. Baxter, S. T. Grafton, L. K. Griffeth, R. A. Hawkins, C. K. Hoh, J. C. Mazziotta, E. J. Hoffman and C. E. Metz. 'Results of a clinical receiver operating characteristic study comparing filtered backprojection and maximum likelihood estimator images in FDG PET studies'. In: *Journal of Nuclear Medicine* 34.7 (1993), pp. 1198–1203. ISSN: 01615505.

- [105] Ilya Loshchilov and Frank Hutter. 'Decoupled weight decay regularization'. In: Conference on Learning Representations. Nov. 2019. DOI: 10.48550/arxiv. 1711.05101.
- [106] Daniel A. Low, Parag J. Parikh, Wei Lu, James F. Dempsey, Sasha H. Wahab, James P. Hubenschmidt, Michelle M. Nystrom, Maureen Handoko and Jeffrey D. Bradley. 'Novel breathing motion model for radiotherapy'. In: *International Journal of Radiation Oncology Biology Physics* 63.3 (Nov. 2005), pp. 921–929. ISSN: 03603016. DOI: 10.1016/j.ijrobp.2005.03.070.
- [107] Yihuan Lu, Kathryn Fontaine, Tim Mulnix, John A. Onofrey, Silin Ren, Vladimir Panin, Judson Jones, Michael E. Casey, Robert Barnett, Peter Kench, Roger Fulton, Richard E. Carson and Chi Liu. 'Respiratory motion compensation for PET/CT with motion information derived from matched attenuation-corrected gated PET data'. In: *Journal of Nuclear Medicine* 59.9 (Sept. 2018), pp. 1480–1486. ISSN: 2159662X. DOI: 10.2967/jnumed.117.203000.
- [108] Richard Manber, Kris Thielemans, Brian F. Hutton, Anna Barnes, Sébastien Ourselin, Simon Arridge, Celia O'Meara, Simon Wan and David Atkinson. 'Practical PET respiratory motion correction in clinical PET/MR'. In: *Journal of Nuclear Medicine* 56.6 (June 2015), pp. 890–896. ISSN: 2159662X. DOI: 10.2967/jnumed.114.151779.
- [109] Richard Manber, Kris Thielemans, Brian F. Hutton, Simon Wan, Francesco Fraioli, Anna Barnes, Sébastien Ourselin, Simon Arridge and David Atkinson. 'Clinical impact of respiratory motion correction in simultaneous PET/MR, using a joint PET/MR predictive motion model'. In: *Journal of Nuclear Medicine* 59.9 (Sept. 2018), pp. 1467–1473. ISSN: 2159662X. DOI: 10.2967/JNUMED.117. 191460,.
- [110] Richard Manber, Kris Thielemans, Brian F. Hutton, Simon Wan, Jamie McClelland, Anna Barnes, Simon Arridge, Sébastien Ourselin and David Atkinson. 'Joint PET-MR respiratory motion models for clinical PET motion correction'. In: *Physics in Medicine and Biology* 61.17 (Sept. 2016), pp. 6515–6530. ISSN: 13616560. DOI: 10.1088/0031-9155/61/17/6515.

- [111] Ravindra Manjeshwar, Xiaodong Tao, Evren Asma and Kris Thielemans. 'Motion compensated image reconstruction of respiratory gated PET/CT'. In: 2006 3rd IEEE International Symposium on Biomedical Imaging: From Nano to Macro-Proceedings. Vol. 2006. 2. IEEE, 2006, pp. 674–677. ISBN: 0780395778. DOI: 10.1109/isbi.2006.1625006.
- [112] David Mattes, David R. Haynor, Hubert Vesselle, Thomas K. Lewellen and William Eubank. 'PET-CT image registration in the chest using free-form deformations'. In: *IEEE Transactions on Medical Imaging* 22.1 (2003), pp. 120–128. ISSN: 02780062. DOI: 10.1109/TMI.2003.809072.
- [113] J. R. McClelland, D. J. Hawkes, T. Schaeffter and A. P. King. 'Respiratory motion models: A review'. In: *Medical Image Analysis* 17.1 (Jan. 2013), pp. 19–42. ISSN: 13618415. DOI: 10.1016/j.media.2012.09.005.
- [114] Jamie R. McClelland, Jane M. Blackall, Ségolène Tarte, Adam C. Chandler, Simon Hughes, Shahreen Ahmad, David B. Landau and David J. Hawkes. 'A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy'. In: *Medical Physics* 33.9 (2006), pp. 3348–3358. ISSN: 00942405. DOI: 10.1118/1.2222079.
- [115] Jamie R. McClelland, Benjamin A.S. Champion and David J. Hawkes. 'Combining image registration, respiratory motion modelling, and motion compensated image reconstruction'. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 8545 LNCS (2014), pp. 103–113. ISSN: 16113349. DOI: 10.1007/978-3-319-08554-8{\\_}11.
- [116] Jamie R. McClelland, Marc Modat, Simon Arridge, Helen Grimes, Derek D'Souza, David Thomas, Dylan O. Connell, Daniel A. Low, Evangelia Kaza, David J. Collins, Martin O. Leach and David J. Hawkes. 'A generalized framework unifying image registration and respiratory motion models and incorporating image reconstruction, for partial image data or full images'. In: *Physics in Medicine and Biology* 62.11 (June 2017), pp. 4273–4292. ISSN: 13616560. DOI: 10.1088/1361-6560/aa6070.

- [117] Qing Le Meng, Rui Yang, Run Ze Wu, Lei Xu, Hao Liu, Gang Yang, Yun Dong, Feng Wang, Zhengguo Chen and Hongbing Jiang. 'Evaluation of a respiratory motion-corrected image reconstruction algorithm in 2-[18F]FDG and [68Ga]Ga-DOTA-NOC PET/CT: impacts on image quality and tumor quantification'. In: *Quantitative Imaging in Medicine and Surgery* 13.1 (Jan. 2023), pp. 370–383. ISSN: 22234306. DOI: 10.21037/QIMS-22-557/COIF.
- [118] G. Mettivier, R. Salvati, M. Conti and P. Russo. 'The effect of count statistics on the convergence value in OSEM reconstruction in PET and TOF PET'. In: *IEEE Nuclear Science Symposium Conference Record*. Institute of Electrical and Electronics Engineers Inc., 2011, pp. 2400–2406. ISBN: 9781467301183. DOI: 10.1109/NSSMIC.2011.6152654.
- [119] Fred A. Mettler and Milton J. Guiberteau. *Essentials of nuclear medicine imaging:*Expert consult. Vol. 31. Elsevier, Jan. 2012, pp. 1–475. ISBN: 9781455701049.

  DOI: 10.1016/C2016-0-00043-8.
- [120] Alan Miranda, Steven Staelens, Sigrid Stroobants and Jeroen Verhaeghe. 'Markerless rat head motion tracking using structured light for brain PET imaging of unrestrained awake small animals'. In: *Physics in Medicine and Biology* 62.5 (Mar. 2017), pp. 1744–1758. ISSN: 13616560. DOI: 10.1088/1361-6560/aa5a46.
- [121] Diganta Misra. 'Mish: A Self Regularized Non-Monotonic Activation Function'.

  In: British Machine Vision Conference. Aug. 2019. DOI: 10.48550/arxiv.

  1908.08681. URL: http://arxiv.org/abs/1908.08681.
- [122] Marc Modat, David M. Cash, Pankaj Daga, Gavin P. Winston, John S. Duncan and Sébastien Ourselin. 'Global image registration using a symmetric block-matching approach'. In: *Journal of Medical Imaging* 1.2 (2014), p. 024003. ISSN: 2329-4302. DOI: 10.1117/1.jmi.1.2.024003.
- [123] Marc Modat, Pankaj Daga, M. Jorge Cardoso, Sebastien Ourselin, Gerard R. Ridgway and John Ashburner. 'Parametric non-rigid registration using a stationary velocity field'. In: *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis* (2012), pp. 145–150. DOI: 10.1109/MMBIA.2012.6164745.

- [124] Marc Modat, Gerard R. Ridgway, Zeike A. Taylor, Manja Lehmann, Josephine Barnes, David J. Hawkes, Nick C. Fox and Sébastien Ourselin. 'Fast free-form deformation using graphics processing units'. In: *Computer Methods and Programs in Biomedicine* 98.3 (2010), pp. 278–284. ISSN: 01692607. DOI: 10.1016/j.cmpb.2009.09.002.
- [125] Jan Modersitzki. Fair: Flexible Algorithms for Image Registration. Society for Industrial and Applied Mathematics, Jan. 2009. ISBN: 978-0-89871-690-0. DOI: 10.1137/1.9780898718843.
- [126] A. Michael Morey and Dan J. Kadrmas. 'Effect of varying number of OSEM subsets on PET lesion detectability'. In: *Journal of Nuclear Medicine Technology* 41.4 (2013), pp. 268–273. ISSN: 00914916. DOI: 10.2967/jnmt.113.131904.
- [127] Mark Muzi, Finbarr O'Sullivan, David A. Mankoff, Robert K. Doot, Larry A. Pierce, Brenda F. Kurland, Hannah M. Linden and Paul E. Kinahan. 'Quantitative assessment of dynamic PET imaging data in cancer imaging'. In: *Magnetic Resonance Imaging* 30.9 (Nov. 2012), pp. 1203–1215. ISSN: 0730725X. DOI: 10.1016/j.mri.2012.05.008.
- [128] In Jae Myung. 'Tutorial on maximum likelihood estimation'. In: *Journal of Mathematical Psychology* 47.1 (Feb. 2003), pp. 90–100. ISSN: 00222496. DOI: 10. 1016/S0022-2496(02)00028-7.
- [129] Sadek A. Nehmeh and Yusuf E. Erdi. 'Respiratory Motion in Positron Emission Tomography/Computed Tomography: A Review'. In: *Seminars in Nuclear Medicine* 38.3 (May 2008), pp. 167–176. ISSN: 00012998. DOI: 10.1053/j.semnuclmed.2008.01.002.
- [130] Sadek A. Nehmeh, Yusuf E. Erdi, Gustavo S.P. Meirelles, Olivia Squire, Steven M. Larson, John L. Humm and Heiko Schöder. 'Deep-inspiration breath-hold PET/CT of the thorax'. In: *Journal of Nuclear Medicine* 48.1 (2007), pp. 22–26. ISSN: 01615505.
- [131] NHS England. *Diagnostic Imaging Dataset Annual Statistical Release 2019/20*. Tech. rep. 2020, pp. 1–29.

- [132] Koen Nieman, Oliver Gaemperli, Patrizio Lancellotti and Sven Plein. *Advanced cardiac imaging*. Elsevier Inc., Sept. 2015, pp. 1–781. ISBN: 9781782422945. DOI: 10.1016/C2013-0-16517-8.
- [133] P. J. Noonan, J. Howard, W. A. Hallett and R. N. Gunn. 'Repurposing the Microsoft Kinect for Windows v2 for external head motion tracking for brain PET'. In: *Physics in Medicine and Biology* 60.22 (Nov. 2015), pp. 8753–8766. ISSN: 13616560. DOI: 10.1088/0031-9155/60/22/8753.
- [134] Philip J. Noonan, Jon Howard, Deborah Tout, Ian Armstrong, Heather A. Williams, Tim F. Cootes, William A. Hallett and Rainer Hinz. 'Accurate markerless respiratory tracking for gated whole body PET using the Microsoft Kinect'. In: *IEEE Nuclear Science Symposium Conference Record.* 2012, pp. 3973–3974. ISBN: 9781467320306. DOI: 10.1109/NSSMIC.2012.6551910.
- [135] Johan Nuyts, Girish Bal, Frank Kehren, Matthias Fenchel, Christian Michel and Charles Watson. 'Completion of a truncated attenuation image from the attenuated PET emission data'. In: *IEEE Transactions on Medical Imaging* 32.2 (2013), pp. 237–246. ISSN: 02780062. DOI: 10.1109/TMI.2012.2220376,.
- [136] Johan Nuyts, Ahmadreza Rezaei and Michel Defrise. 'ML-reconstruction for TOF-PET with simultaneous estimation of the attenuation factors'. In: *IEEE Nuclear Science Symposium Conference Record* (2012), pp. 2147–2149. ISSN: 10957863. DOI: 10.1109/NSSMIC.2012.6551491.
- [137] Matthew J. Nyflot, Tzu Cheng Lee, Adam M. Alessio, Scott D. Wollenweber, Charles W. Stearns, Stephen R. Bowen and Paul E. Kinahan. 'Impact of CT attenuation correction method on quantitative respiratory-correlated (4D) PET/CT imaging'. In: *Medical Physics* 42.1 (Jan. 2015), pp. 110–120. ISSN: 00942405. DOI: 10.1118/1.4903282.
- [138] H. Nyquist. 'Certain Topics in Telegraph Transmission Theory'. In: *Transactions of the American Institute of Electrical Engineers* 47.2 (1928), pp. 617–644. ISSN: 00963860. DOI: 10.1109/T-AIEE.1928.5055024.

- [139] Se An Oh, Ji Woon Yea, Sung Kyu Kim and Jae Won Park. 'Optimal Gating Window for Respiratory-Gated Radiotherapy with Real-Time Position Management and Respiration Guiding System for Liver Cancer Treatment'. In: *Scientific Reports* 9.1 (Mar. 2019), pp. 1–6. ISSN: 20452322. DOI: 10.1038/s41598-019-40858-2.
- [140] Francisco P.M. Oliveira and João Manuel R.S. Tavares. 'Medical image registration: A review'. In: *Computer Methods in Biomechanics and Biomedical Engineering* 17.2 (Jan. 2014), pp. 73–93. ISSN: 10255842. DOI: 10.1080/10255842. 2012.670855.
- [141] Evgueni Ovtchinnikov, Richard Brown, Christoph Kolbitsch, Edoardo Pasca, Casper da Costa-Luis, Ashley G. Gillman, Benjamin A. Thomas, Nikos Efthimiou, Johannes Mayer, Palak Wadhwa, Matthias J. Ehrhardt, Sam Ellis, Jakob S. Jørgensen, Julian Matthews, Claudia Prieto, Andrew J. Reader, Charalampos Tsoumpas, Martin Turner, David Atkinson and Kris Thielemans. 'SIRF: Synergistic Image Reconstruction Framework'. In: *Computer Physics Communications* 249 (July 2020), pp. 1–3. ISSN: 00104655. DOI: 10.1016/j.cpc.2019.107087.
- [142] Evgueni Ovtchinnikov, Richard Brown, Johannes Mayer, Edoardo Pasca, Casper da Costa-Luis, David Atkinson, Christoph Kolbitsch, Nikos Efthimiou, Matthew Strugari, Ashley Gillman, Sam David Porter, Ander Biguri, Daniel Deidda, Alexander Whitehead, Evangelos Papoutsellis, Gemma Fardell, Ben Thomas, Francesca Leek, Matthias Ehrhardt and Kris Thielemans. 'SIRF Synergistic Image Reconstruction Framework'. In: (Jan. 2023). DOI: 10.5281/ZENOD0.7551178. URL: https://zenodo.org/record/7551178.
- [143] Evgueni Ovtchinnikov, Richard Brown, Kris Thielemans, Edoardo Pasca, Casper da Costa-Luis, Benjamin A. Thomas, David Atkinson, Johannes Mayer, Ashley Gillman, Christoph Kolbitsch, Matthias J. Ehrhardt and Alexander Whitehead. 'SIRF Virtual Machine'. In: (May 2021). DOI: 10.5281/ZENOD0.7050371.
- [144] Tinsu Pan, Samuel A. Einstein, Srinivas Cheenu Kappadath, Kira S. Grogg, Cristina Lois Gomez, Adam M. Alessio, William C. Hunter, Georges El Fakhri,

- Paul E. Kinahan and Osama R. Mawlawi. 'Performance evaluation of the 5-Ring GE Discovery MI PET/CT system using the national electrical manufacturers association NU 2-2012 Standard'. In: *Medical Physics* 46.7 (2019), pp. 3025–3033. ISSN: 00942405. DOI: 10.1002/mp.13576.
- [145] Tinsu Pan, Ting Yim Lee, Eike Rietzel and George T.Y. Chen. '4D-CT imaging of a volume influenced by respiratory motion on multi-slice CT'. In: *Medical Physics* 31.2 (2004), pp. 333–340. ISSN: 00942405. DOI: 10.1118/1.1639993.
- [146] Tinsu Pan, Osama Mawlawi, Dershan Luo, Hui H. Liu, Pai Chun M. Chi, Martha V. Mar, Gregory Gladish, Mylene Truong, Jeremy Erasmus, Zhongxing Liao and H. A. Macapinlac. 'Attenuation correction of PET cardiac data with low-dose average CT in PET/CT'. In: *Medical Physics* 33.10 (Oct. 2006), pp. 3931–3938.
  ISSN: 2473-4209. DOI: 10.1118/1.2349843.
- [147] V. Y. Panin, M. Aykac and M. E. Casey. 'Simultaneous reconstruction of emission activity and attenuation coefficient distribution from TOF data, acquired with external transmission source'. In: *Physics in Medicine and Biology* 58.11 (2013), pp. 3649–3669. ISSN: 00319155. DOI: 10.1088/0031-9155/58/11/3649,.
- [148] Edoardo Pasca, Ben Thomas, Casper da Costa-Luis, Richard Brown, Ander Biguri, Evgueni Ovtchinnikov, Ashley Gillman, Christoph Kolbitsch, Johannes Mayer, Alexander Whitehead, Georg Schramm, David Atkinson and Kris Thielemans. 'SIRF-SuperBuild'. In: (Aug. 2022). DOI: 10.5281/ZENOD0.7022530.
- [149] C. S. Patlak, R. G. Blasberg and J. D. Fenstermacher. 'Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data'. In: *Journal of Cerebral Blood Flow and Metabolism* 3.1 (1983), pp. 1–7. ISSN: 0271678X. DOI: 10.1038/jcbfm.1983.1. URL: https://pubmed.ncbi.nlm.nih.gov/6822610/.
- [150] Karl Pearson. 'On lines and planes of closest fit to systems of points in space'. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (Nov. 1901), pp. 559–572. ISSN: 1941-5982. DOI: 10.1080/14786440109462720.

- [151] I. Polycarpou, C. Tsoumpas and P. K. Marsden. 'Analysis and comparison of two methods for motion correction in PET imaging'. In: *Medical Physics* 39.10 (Oct. 2012), pp. 6474–6483. ISSN: 00942405. DOI: 10.1118/1.4754586.
- [152] Stefanie Pösse, Florian Büther, Dirk Mannweiler, Inki Hong, Judson Jones, Michael Schäfers and Klaus Peter Schäfers. 'Comparison of two elastic motion correction approaches for whole-body PET/CT: motion deblurring vs gate-to-gate motion correction'. In: *EJNMMI Physics* 7.1 (Dec. 2020), pp. 1–14. ISSN: 21977364. DOI: 10.1186/S40658-020-0285-4/TABLES/4.
- [153] Snehal Prabhudesai, Jeremiah Hauth, Dingkun Guo, Arvind Rao, Nikola Banovic and Xun Huan. 'Lowering the computational barrier: Partially Bayesian neural networks for transparency in medical imaging AI'. In: *Frontiers in Computer Science* 5 (Feb. 2023), p. 1071174. ISSN: 26249898. DOI: 10.3389/FCOMP. 2023.1071174/BIBTEX.
- [154] Feng Qiao, Tinsu Pan, John W. Clark and Osama R. Mawlawi. 'A motion-incorporated reconstruction method for gated PET studies'. In: *Physics in Medicine and Biology* 51.15 (July 2006), pp. 3769–3783. ISSN: 00319155. DOI: 10.1088/0031-9155/51/15/012.
- [155] Silin Ren, Xiao Jin, Chung Chan, Yiqiang Jian, Tim Mulnix, Chi Liu and Richard E. Carson. 'Data-driven event-by-event respiratory motion correction using TOF PET list-mode centroid of distribution'. In: *Physics in Medicine and Biology* 62.12 (June 2017), pp. 4741–4755. ISSN: 13616560. DOI: 10.1088/1361-6560/aa700c.
- [156] Ahmadreza Rezaei, Michel Defrise, Girish Bal, Christian Michel, Maurizio Conti, Charles Watson and Johan Nuyts. 'Simultaneous reconstruction of activity and attenuation in time-of-flight PET'. In: *IEEE Transactions on Medical Imaging* 31.12 (2012), pp. 2224–2233. ISSN: 02780062. DOI: 10.1109/TMI.2012. 2212719.
- [157] Ahmadreza Rezaei, Christophe M. Deroose, Thomas Vahle, Fernando Boada and Johan Nuyts. 'Joint reconstruction of activity and attenuation in time-of-flight

- PET: A quantitative analysis'. In: *Journal of Nuclear Medicine* 59.10 (2018), pp. 1630–1635. ISSN: 2159662X. DOI: 10.2967/jnumed.117.204156.
- [158] Ahmadreza Rezaei, Christian Michel, Michael E. Casey and Johan Nuyts. 'Simultaneous reconstruction of the activity image and registration of the CT image in TOF-PET'. In: *Physics in Medicine and Biology* 61.4 (2016), pp. 1852–1874. ISSN: 13616560. DOI: 10.1088/0031-9155/61/4/1852.
- [159] Steve Ross. 'Q. Clear'. In: GE Healthcare, White Paper (2014), pp. 1–9.
- [160] D. Rueckert. 'Nonrigid registration using free-form deformations: Application to breast mr images'. In: *IEEE Transactions on Medical Imaging* 18.8 (1999), pp. 712–721. ISSN: 02780062. DOI: 10.1109/42.796284.
- [161] Zohaib Salahuddin, Henry C. Woodruff, Avishek Chatterjee and Philippe Lambin.
  'Transparency of deep neural networks for medical image analysis: A review of interpretability methods'. In: *Computers in Biology and Medicine* 140 (Jan. 2022), p. 105111. ISSN: 0010-4825. DOI: 10.1016/J.COMPBIOMED.2021.105111.
- [162] André Salomon, Andreas Goedicke, Bernd Schweizer, Til Aach and Volkmar Schulz. 'Simultaneous reconstruction of activity and attenuation for PET/MR'. In: *IEEE Transactions on Medical Imaging* 30.3 (Mar. 2011), pp. 804–813. ISSN: 02780062. DOI: 10.1109/TMI.2010.2095464.
- [163] Hasan Sari, Kjell Erlandsson, Lisbeth Marner, Ian Law, Henrik B.W. Larsson, Kris Thielemans, Sébastien Ourselin, Simon Arridge, David Atkinson and Brian F. Hutton. 'Non-invasive kinetic modelling of PET tracers with radiometabolites using a constrained simultaneous estimation method: evaluation with 11C-SB201745'. In: *EJNMMI Research* 8 (2018), p. 58. ISSN: 2191219X. DOI: 10.1186/S13550-018-0412-6.
- [164] Abraham Savitzky and Marcel J.E. Golay. 'Smoothing and Differentiation of Data by Simplified Least Squares Procedures'. In: *Analytical Chemistry* 36.8 (July 1964), pp. 1627–1639. ISSN: 15206882. DOI: 10.1021/ac60214a047.

- [165] Orazio Schillaci and Nicoletta Urbano. *Digital PET/CT: a new intriguing chance for clinical nuclear medicine and personalized molecular imaging*. June 2019. DOI: 10.1007/s00259-019-04300-z.
- [166] P. J. Schleyer, K. Thielemans and P. K. Marsden. 'Extracting a respiratory signal from raw dynamic PET data that contain tracer kinetics'. In: *Physics in Medicine and Biology* 59.15 (2014), pp. 4345–4356. ISSN: 13616560. DOI: 10.1088/0031-9155/59/15/4345.
- [167] Paul Schleyer, Inki Hong, Judson Jones, James Hamill, Vladimir Panin and Sebastian Fuerst. 'Data-Driven Respiratory Gating Whole Body PET Using Continuous Bed Motion'. In: 2018 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2018 Proceedings. Institute of Electrical and Electronics Engineers Inc., Nov. 2018. ISBN: 9781538684948. DOI: 10.1109/NSSMIC.2018.8824401.
- [168] Paul J. Schleyer, Michael J. O'Doherty, Sally F. Barrington and Paul K. Marsden. 'Retrospective data-driven respiratory gating for PET/CT'. In: *Physics in Medicine and Biology* 54.7 (2009), pp. 1935–1950. ISSN: 00319155. DOI: 10.1088/0031-9155/54/7/005.
- [169] Paul J. Schleyer, Michael J. O'Doherty and Paul K. Marsden. 'Extension of a data-driven gating technique to 3D, whole body PET studies'. In: *Physics in Medicine and Biology* 56.13 (2011), pp. 3953–3965. ISSN: 00319155. DOI: 10.1088/0031-9155/56/13/013.
- [170] W. P. Segars, G. Sturgeon, S. Mendonca, Jason Grimes and B. M.W. Tsui. '4D XCAT phantom for multimodality imaging research'. In: *Medical Physics* 37.9 (Aug. 2010), pp. 4902–4915. ISSN: 00942405. DOI: 10.1118/1.3480985.
- [171] Andrew P. Selwyn, Roger M. Allan, Antonio L'Abbate, Peter Horlock, Paolo Camici, John Clark, Harold A. O'Brien and Peter M. Grant. 'Relation between regional myocardial uptake of rubidium-82 and perfusion: Absolute reduction of cation uptake in ischemia'. In: *The American Journal of Cardiology* 50.1 (1982), pp. 112–121. ISSN: 00029149. DOI: 10.1016/0002-9149(82)90016-9.

- [172] Ahmed H. Shahin, An Zhao, Alexander C. Whitehead, Daniel C. Alexander, Joseph Jacob and David Barber. 'CenTime: Event-conditional modelling of censoring in survival analysis'. In: *Medical Image Analysis* 91 (Jan. 2024), p. 103016. ISSN: 13618423. DOI: 10.1016/J.MEDIA.2023.103016.
- [173] Claude E. Shannon. 'Communication in the Presence of Noise'. In: *Proceedings of the IRE* 37.1 (1949), pp. 10–21. ISSN: 00968390. DOI: 10.1109/JRPROC. 1949.232969.
- [174] L. A. Shepp and Y. Vardi. 'Maximum Likelihood Reconstruction for Emission Tomography'. In: *IEEE Transactions on Medical Imaging* 1.2 (1982), pp. 113–122. ISSN: 1558254X. DOI: 10.1109/TMI.1982.4307558.
- [175] K. Shibuya, E. Yoshida, F. Nishikido, T. Suzuki, T. Tsuda, N. Inadama, T. Yamaya and H. Murayama. 'Limit of spatial resolution in FDG-PET due to annihilation Photon Non-Collinearity'. In: *IFMBE Proceedings*. Vol. 14. 1. Springer Verlag, 2007, pp. 1667–1671. ISBN: 9783540368397. DOI: 10.1007/978-3-540-36841-0{\\_}411.
- [176] Jonathan Sigfridsson, Elin Lindström, Victor Iyer, Maria Holstensson, Irina Velikyan, Anders Sundin and Mark Lubberink. 'Prospective data-driven respiratory gating of [68Ga]Ga-DOTATOC PET/CT'. In: *EJNMMI Research* 11.1 (Mar. 2021), pp. 1–9. ISSN: 2191219X. DOI: 10.1186/S13550-021-00775-W/FIGURES/4.
- [177] Evan Silverstein and Michael Snyder. 'Comparative analysis of respiratory motion tracking using Microsoft Kinect v2 sensor'. In: *Journal of Applied Clinical Medical Physics* 19.3 (May 2018), pp. 193–204. ISSN: 15269914. DOI: 10.1002/acm2.12318.
- [178] Virginia Ch Spanoudaki and Craig S. Levin. 'Photo-detectors for time of flight positron emission tomography (ToF-PET)'. In: *Sensors (Switzerland)* 10.11 (Nov. 2010), pp. 10484–10505. ISSN: 14248220. DOI: 10.3390/s101110484.
- [179] Waheeda Sureshbabu and Osama Mawlawi. 'PET/CT imaging artifacts'. In: *Journal of Nuclear Medicine Technology* 33.3 (2005), pp. 156–161. ISSN: 00914916.

- [180] Suleman Surti. 'Update on time-of-flight PET imaging'. In: *Journal of Nuclear Medicine* 56.1 (2015), pp. 98–105. ISSN: 01615505. DOI: 10.2967/jnumed. 114.145029.
- [181] Andrew M. Taylor, Permi Jhooti, Frank Wiesmann, Jennifer Keegan, David N. Firmin and Dudley J. Pennell. 'MR navigator-echo monitoring of temporal changes in diaphragm position: Implications for MR coronary angiography'. In: *Journal of Magnetic Resonance Imaging* 7.4 (July 1997), pp. 629–636. ISSN: 10531807. DOI: 10.1002/jmri.1880070404.
- [182] M. M. Ter-Pogossian, Nizar A. Mullani, David C. Ficke, Joanne Markham and Donald L. Snyder. 'Photon time-of-flight-assisted positron emission tomography'.
   In: *Journal of Computer Assisted Tomography* 5.2 (Apr. 1981), pp. 227–239.
   ISSN: 15323145. DOI: 10.1097/00004728-198104000-00014.
- [183] Kris Thielemans, Shailendra Rathore, Fredrik Engbrant and Pasha Razifar. 'Device-less gating for PET/CT using PCA'. In: *IEEE Nuclear Science Symposium Conference Record*. IEEE, 2011, pp. 3904–3910. ISBN: 9781467301183. DOI: 10.1109/NSSMIC.2011.6153742.
- [184] Kris Thielemans, Paul Schleyer, Paul K. Marsden, Ravindra M. Manjeshwar, Scott D. Wollenweber and Alexander Ganin. 'Comparison of different methods for data-driven respiratory gating of PET data'. In: *IEEE Nuclear Science Symposium Conference Record*. IEEE, Oct. 2013, pp. 1–4. ISBN: 9781479905348. DOI: 10.1109/NSSMIC.2013.6829055.
- [185] Kris Thielemans, Charalampos Tsoumpas, Sanida Mustafovic, Tobias Beisel, Pablo Aguiar, Nikolaos Dikaios and Matthew W. Jacobson. 'STIR: Software for tomographic image reconstruction release 2'. In: *Physics in Medicine and Biology* 57.4 (Feb. 2012), pp. 867–883. ISSN: 00319155. DOI: 10.1088/0031-9155/57/4/867.
- [186] Nicholas J. Tustison, Brian B. Avants and James C. Gee. 'Directly manipulated free-form deformation image registration'. In: *IEEE Transactions on Image Processing* 18.3 (2009), pp. 624–635. ISSN: 10577149. DOI: 10.1109/TIP.2008. 2010072.

- [187] Dmitry Ulyanov, Andrea Vedaldi and Victor Lempitsky. 'Deep Image Prior'. In: *International Journal of Computer Vision* 128.7 (July 2020), pp. 1867–1888. ISSN: 15731405. DOI: 10.1007/s11263-020-01303-4.
- [188] Joyce Van Sluis, Johan De Jong, Jenny Schaar, Walter Noordzij, Paul Van Snick, Rudi Dierckx, Ronald Borra, Antoon Willemsen and Ronald Boellaard. 'Performance characteristics of the digital biograph vision PET/CT system'. In: *Journal of Nuclear Medicine* 60.7 (2019), pp. 1031–1036. ISSN: 2159662X. DOI: 10. 2967/jnumed.118.215418.
- [189] N. Venkatanath, D. Praneeth, B. H. Maruthi Chandrasekhar, Sumohana S. Channappayya and Swarup S. Medasani. 'Blind image quality evaluation using perception based features'. In: 2015 21st National Conference on Communications, NCC 2015. Institute of Electrical and Electronics Engineers Inc., Apr. 2015. ISBN: 9781479966196. DOI: 10.1109/NCC.2015.7084843.
- [190] Tom Vercauteren, Xavier Pennec, Aymeric Perchant and Nicholas Ayache. 'Diffeomorphic demons: efficient non-parametric image registration.' In: *NeuroImage* 45.1 Suppl (Mar. 2009), S61–S72. ISSN: 10959572. DOI: 10.1016/j.neuroimage.2008.10.040.
- [191] Joseph Stephen Verducci, Xiaotong Shen and John Lafferty. *Prediction and Discovery*. Ed. by Joseph Stephen Verducci, Xiaotong Shen and John Lafferty. Vol. 443. Contemporary Mathematics. Providence, Rhode Island: American Mathematical Society, 2007. ISBN: 978-0-8218-4195-2. DOI: 10.1090/CONM/443.
- [192] Christopher Voscopoulos, Jordan Brayanov, Diane Ladd, Michael Lalli, Alexander Panasyuk and Jenny Freeman. 'Evaluation of a novel noninvasive respiration monitor providing continuous measurement of minute ventilation in ambulatory subjects in a variety of clinical scenarios'. In: *Anesthesia and Analgesia* 117.1 (July 2013), pp. 91–100. ISSN: 00032999. DOI: 10.1213/ANE.0b013e3182918098.
- [193] M D Walker, K Su, S D Wollenweber, R Johnsen and D R McGowan. 'Automatic Classification of Data-Driven Respiratory Waveforms Using AI'. In: *European*

- *Journal of Nuclear Medicine and Molecular Imaging*. Vol. 47. SUPPL 1. Dec. 2020, S485–S485.
- [194] Matthew D. Walker, Kevin M. Bradley and Daniel R. McGowan. 'Evaluation of principal component analysis-based data-driven respiratory gating for positron emission tomography'. In: *British Journal of Radiology* 91.1085 (2018). ISSN: 1748880X. DOI: 10.1259/BJR.20170793,.
- [195] Matthew D. Walker, Andrew J. Morgan, Kevin M. Bradley and Daniel R. McGowan. 'Evaluation of data-driven respiratory gating waveforms for clinical PET imaging'. In: *EJNMMI Research* 9 (2019). ISSN: 2191219X. DOI: 10.1186/s13550-018-0470-9.
- [196] Zhou Wang and Alan C. Bovik. 'Mean squared error: Lot it or leave it? A new look at signal fidelity measures'. In: *IEEE Signal Processing Magazine* 26.1 (2009), pp. 98–117. ISSN: 10535888. DOI: 10.1109/MSP.2008.930649.
- [197] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh and Eero P. Simoncelli. 'Image quality assessment: From error visibility to structural similarity'. In: *IEEE Transactions on Image Processing* 13.4 (2004), pp. 600–612. ISSN: 10577149. DOI: 10.1109/TIP.2003.819861.
- [198] Joachim Weiss. *Handbook of ion chromatography*. Vol. 1-3. wiley, Dec. 2016, pp. 1–1522. ISBN: 9783527651610. DOI: 10.1002/9783527651610.
- [199] Weihao Weng and Xin Zhu. 'INet: Convolutional Networks for Biomedical Image Segmentation'. In: *IEEE Access* 9 (May 2021), pp. 16591–16603. ISSN: 21693536. DOI: 10.1109/ACCESS.2021.3053408.
- [200] A. C. Whitehead, L. Brusaferri, L. Maccioni, M. Ferrante, M. Inglese, Z. Alshelh, M. Veronese, N. Toschi, J. Gilman, K. Thielemans and M. L. Loggia. 'A Bayesian Neural Network-Based Method For The Extraction Of A Metabolite Corrected Arterial Input Function From Dynamic [11 C]PBR28 PET'. In: 2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTSD). Institute of Electrical and Electronics Engineers (IEEE), Dec. 2023, pp. 1–1. DOI: 10.1109/NSSMICRTSD49126.2023.10338687.

- [201] A. C. Whitehead, A. H. Shahin, A. Zhao, D. C. Alexander, J. Jacob and D. Barber. 'Neural Network Based Methods for the Survival Analysis of Idiopathic Pulmonary Fibrosis Patients from a Baseline CT Acquisition'. In: 2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTSD). Institute of Electrical and Electronics Engineers (IEEE), Dec. 2023, pp. 1–1. DOI: 10.1109/NSSMICRTSD49126.2023.10337922.
- [202] Alexander C. Whitehead, Ander Biguri, Nikos Efthimiou, Kuan Hao Su, Scott W. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'PET/CT Respiratory Motion Correction with a Single Attenuation Map Using NAC Derived Deformation Fields'. In: 2020 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2020. IEEE, Aug. 2020, pp. 1–3. ISBN: 9781728176932. DOI: 10.1109/NSS/MIC42677. 2020.9507890.
- [203] Alexander C. Whitehead, Ander Biguri, Kuan Hao Su, Scott D. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'Comparison of Motion Correction Methods Incorporating Motion Modelling for PET/CT Using a Single Breath Hold Attenuation Map'. In: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference Record, NSS/MIC 2021 and 28th International Symposium on Room-Temperature Semiconductor Detectors, RTSD 2022. Institute of Electrical and Electronics Engineers Inc., Sept. 2021, pp. 1–4. ISBN: 9781665421133. DOI: 10.1109/NSS/MIC44867.2021. 9875560.
- [204] Alexander C. Whitehead and Ludovica Brusaferri. 'Mortality Maps: Inherently Interpretable Event-Conditional Survival Analysis'. In: *arXiv* (2025).
- [205] Alexander C. Whitehead, Elise C. Emond, Nikos Efthimiou, Adeyemi Akintonde, Scott W. Wollenweber, Charles W. Stearns, Brian F. Hutton, Jamie R. McClelland and Kris Thielemans. 'Impact of Time-of-Flight on Respiratory Motion Modelling using Non-Attenuation-Corrected PET'. In: 2019 IEEE Nuclear Science

- Symposium and Medical Imaging Conference, NSS/MIC 2019. IEEE, Oct. 2019. ISBN: 9781728141640. DOI: 10.1109/NSS/MIC42101.2019.9059812.
- [206] Alexander C. Whitehead, Kjell Erlandsson, Ander Biguri, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Pseudo-Bayesian DIP Denoising as a Preprocessing Step for Kinetic Modelling in Dynamic PET'. In: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845.2022.10399318.
- [207] Alexander C. Whitehead, Kuan Hao Su, Elise C. Emond, Ander Biguri, Ludovica Brusaferri, Maria Machado, Joanna C. Porter, Helen Garthwaite, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Data driven surrogate signal extraction for dynamic PET using selective PCA: time windows versus the combination of components'. In: *Physics in Medicine & Biology* 69.17 (Aug. 2024), p. 175008. ISSN: 0031-9155. DOI: 10.1088/1361-6560/AD5EF1.
- [208] Alexander C. Whitehead, Kuan Hao Su, Elise C. Emond, Ander Biguri, Maria MacHado, Joanna C. Porter, Helen Garthwaite, Scott D. Wollenweber, Jamie R. McClelland and Kris Thielemans. 'Data Driven Surrogate Signal Extraction for Dynamic PET Using Selective PCA'. In: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845. 2022.10399196.
- [209] Alexander C. Whitehead, Kuan Hao Su, Scott D. Wollenweber, Jamie R. Mc-Clelland and Kris Thielemans. 'PET/CT Motion Correction Exploiting Motion Models Fit on Coarsely Gated Data Applied to Finely Gated Data'. In: 2022 IEEE NSS/MIC RTSD IEEE Nuclear Science Symposium, Medical Imaging Conference and Room Temperature Semiconductor Detector Conference. Institute of Electrical and Electronics Engineers Inc., 2022. ISBN: 9781665488723. DOI: 10.1109/NSS/MIC44845.2022.10399154.

- [210] Alexander Charles Whitehead. 'Motion Signal Extraction Framework for the Microsoft Kinect Camera: Point Cloud Registration and its Application as a Motion Correction Metric in PET/CT'. PhD thesis. University of Hull, 2018.
- [211] E. T. Whittaker. 'On the Functions which are represented by the Expansions of the Interpolation-Theory'. In: *Proceedings of the Royal Society of Edinburgh* 35 (Jan. 1915), pp. 181–194. ISSN: 0370-1646. DOI: 10.1017/S0370164600017806.
- [212] Yuxin Wu and Kaiming He. 'Group Normalization'. In: *International Journal of Computer Vision* 128.3 (Mar. 2020), pp. 742–755. ISSN: 15731405. DOI: 10.1007/s11263-019-01198-w.
- [213] Junyi Xia and R. Alfredo Siochi. 'A real-time respiratory motion monitoring system using KINECT: Proof of concept'. In: *Medical Physics* 39.5 (May 2012), pp. 2682–2685. ISSN: 00942405. DOI: 10.1118/1.4704644.
- [214] Cheng Hsun Yang and Hsuan Ming Huang. 'Simultaneous Denoising of Dynamic PET Images Based on Deep Image Prior'. In: *Journal of Digital Imaging* 35.4 (Mar. 2022), pp. 834–845. ISSN: 1618727X. DOI: 10.1007/s10278-022-00606-x.
- [215] I. N.Kwon Yeo and Richard A. Johnson. 'A new family of power transformations to improve normality or symmetry'. In: *Biometrika* 87.4 (Dec. 2000), pp. 954–959. ISSN: 00063444. DOI: 10.1093/biomet/87.4.954.
- [216] Yunhan Yu, Chung Chan, Tianyu Ma, Yaqiang Liu, Jean Dominique Gallezot, Mika Naganawa, Olivia J. Kelada, Mary Germino, Albert J. Sinusas, Richard E. Carson and Chi Liu. 'Event-by-event continuous respiratory motion correction for dynamic PET imaging'. In: *Journal of Nuclear Medicine* 57.7 (July 2016), pp. 1084–1090. ISSN: 2159662X. DOI: 10.2967/jnumed.115.167676.
- [217] Paolo Zanotti-Fregonara, Kewei Chen, Jeih San Liow, Masahiro Fujita and Robert B. Innis. 'Image-derived input function for brain PET studies: many challenges and few opportunities'. In: *Journal of cerebral blood flow and metabolism : official journal of the International Society of Cerebral Blood Flow and Metabolism* 31.10 (Oct. 2011), pp. 1986–1998. ISSN: 1559-7016. DOI: 10.1038/JCBFM. 2011.107.

- [218] Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar Tatikonda, Nicha Dvornek, Xenophon Papademetris and James S. Duncan. 'AdaBelief Optimizer: Adapting Stepsizes by the Belief in Observed Gradients'. In: Advances in Neural Information Processing Systems 2020-December (Oct. 2020). ISSN: 10495258.
- [219] H Zimmermann, S Zuehlsdorff, R Umathum, S Volz, W Semmler and M Bock. 'Use of Fiducial Markers for Motion Corrected MRI'. In: *Magnetic Resonance in Medicine* 11 (2003), pp. 2003–2003.