



Improved PET Data Quantification in Simultaneous PET/MR Neuroimaging

Citation

Chonde, Daniel B. 2015. Improved PET Data Quantification in Simultaneous PET/MR Neuroimaging. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:17467219>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Improved PET Data Quantification in Simultaneous PET/MR Neuroimaging

A dissertation presented

By

Daniel Burje Chonde

To

The Program in Biophysics

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

In the subject of

Biophysics

Harvard University

Cambridge, Massachusetts

September 2014

© 2014 Daniel Burje Chonde

All rights reserved.

Improved PET Data Quantification in Simultaneous PET/MR Neuroimaging

Abstract

Recently, systems that integrate positron emission tomography and magnetic resonance imaging (PET/MR) have become available for clinical use. This new technology, which combines the high spatial resolution and superior soft-tissue contrast of MR with the picomolar sensitivity, quantitative capabilities, and wide array of tracers of PET, has the potential to benefit patients and provide insights that were previously unattainable in standalone systems. Simultaneous measurement of PET and MR parameters provides complementary information, allowing for a more complete assessment of disease, as well as cross validation and calibration of MR and PET measurements and techniques.

To take full advantage of such a multi-modal system, accurate quantification of the PET data is necessary. Due to the low spatial resolution of PET – which can be further reduced by external factors like patient motion – and the inherent lack of anatomic detail, accurate quantification can be challenging. The simultaneously acquired MR information provides an opportunity to optimize PET quantification and analysis. In order to fully realize the benefits provided by the simultaneously acquired MR data, the MR data cannot be treated as discrete sequences, but as the continuous flow of information. This is due to differences in the time required for data collection to generate PET and MR images.

This work describes the development and optimization of a pipeline for the reconstruction and analysis of PET data in a brain-dedicated prototype PET/MR system, the

BrainPET (Siemens Healthcare). First, the performance of the BrainPET system was optimized for neurological imaging. MR-hardware interference and characteristics of the PET camera were quantified and a method for multimodal alignment was developed. To simplify and streamline the reconstruction and quantification process, a platform was designed which utilizes the functionality of a number of specialized brain imaging analysis software packages in an automated fashion. Second, MR-based methods addressing specific challenges to PET quantification were addressed. Simultaneously acquired structural MR data was used to correct the PET data for attenuation and partial volume effects. The use of MR data for motion correction was addressed and a unified algorithm which derives motion estimates from the PET data when MR data is unavailable was presented. Finally, the value of the optimized PET processing for neurological studies was evaluated in three instances: first an upper limit on the physiologic noise introduced by MR imaging on cerebral metabolism was estimated using PET and found to be minimal; next the benefit of MR-based motion correction and partial volume effect correction were estimated in a patient study; and lastly, a method to derive the PET radiotracer input function from the PET data using multiple MR sequences was presented.

“Build a better mousetrap, and the world will beat a path to your door”

—“Ralph Waldo Emerson”

Acknowledgments

This dissertation would not have been possible without the support of a great number of people to whom I would like to express my sincerest gratitude.

First of all I would like to thank my family and friends for always being there when I needed them.

I would like to thank my committee—Drs. Jacob Hooker, Joshua Roffman, and Bruce Rosen—for their help, guidance, and support throughout this project. Additionally I would like to thank Michele Jakoulov and Jim Hogle of the Biophysics Department for their invaluable help throughout my graduate career. I owe additional recognition to the Department, Harvard University, and the Ashford Family for financial support throughout my graduate education.

I would like to thank the scientists and staff of Siemens Molecular Imaging Inc at Rockford for all their expertise regarding the technical aspects of this thesis, the BrainPET in general, and for their help in the development of the automated blood sampler. This gratitude extends to: Johannes Breuer, Holli Irwin, James Corbeil, Nan Zhang, Ron Grazioso, Matthew Martin, and Matthias Schmand. I find myself especially lucky to have had the opportunity to work alongside Christian J Michel and Larry Byars, both of whom I consider to be the greatest scientists I will ever have the pleasure of working with. They will continue to be role models that I strive to emulate throughout my career.

I would like to thank the enormous faculty and staff of the Martinos Center, whose names could fill multiple pages, for creating the collaborative environment that shaped this work in tremendous ways. I would especially like to thank Douglas Greve for his help with FreeSurfer, priceless insights on motion estimation, and being a sounding board for ideas.

I owe many thanks to the members of the Catana Group: Spencer Bowen, for having the answer to whatever question I might have; Shirley Hsu, for always being there when I needed help and more than happy to provide it; David Izquierdo-Garcia, for his dedication to the production of tools to be shared with the larger scientific community, and for his insuppressible curiosity, which has led to amazing and fun projects; and Kevin Chen, for his willingness to lend support and his continued interest in advancing our projects. Working with them truly made this experience a pleasure.

Without the help of Grae Arabasz, none of this work would have been possible. He has been an integral part of this undertaking from the very beginning and has been a great support both in the lab and outside of it.

I owe a great debt of gratitude to Ciprian Catana, who accepted me as his student without a second thought, spent countless nights and weekends in the lab with me, and has been an outstanding PhD advisor, a true friend, and, in the sincerest sense, family. Working as his student has been one of the crowning achievements in my life. I am proud to have been his first graduate student.

Finally, I would like to thank the collaborators and users of Masamune, for making use of my work.

Contents

Chapter 1 Integrated PET/MR Imaging	1
1.1 Positron Emission Tomography.....	3
1.1.1 Physical Principles.....	3
1.1.2 Image Reconstruction and Correction Factors	6
1.1.3 Quantitative PET Imaging.....	20
1.1.4 Factors That Affect Spatial Resolution.....	30
1.1.5 Application to Neurological Imaging.....	36
1.2 Magnetic Resonance Imaging	38
1.2.1 Physical Principles.....	38
1.2.2 Application to Neurological Imaging.....	41
1.3 Integrated PET/MR Scanners.....	42
1.3.1 History of PET/MR.....	43
1.3.2 Applications to Neurological Imaging.....	48
Chapter 2 Assessment and Optimization of the PET/MR Scanner	52
2.1 Siemens BrainPET Prototype	52
2.1.1 PET Data Format	54
2.2 PET Reconstruction Optimization.....	59
2.2.1 Regional Spread Function Model.....	59
2.2.2 OP-OSEM Optimization.....	63
2.2.3 Dynamic Imaging Accuracy	77
2.3 MR Gradient and RF Effects on BrainPET Performance	78
2.3.1 Assessment of Detector-level MR Interference	79
2.3.2 Assessment of System-level MR Interference	82
2.4 Multi-modal Alignment	85
2.4.1 Spatial Coregistration.....	85
2.4.2 Temporal Correlation.....	87
2.5 BrainPET Automated Data Processing and Image Reconstruction	90
2.5.1 Detection of Malfunctioning Detector Blocks	91
2.5.2 Automated Bolus Arrival Time Determination	95
2.5.3 BrainPET Image Reconstruction.....	100
2.6 Integrated Tools for Advanced Automated PET Analysis.....	104
2.6.1 Graphical User Interface for PET Reconstruction and PET/MR Analysis ...	104
2.6.2 Region of Interest Labeling	105
2.6.3 Intensity Normalization	112

2.6.4 Surface-based Smoothing	114
2.6.5 Image Reconstruction and Analysis Workflow	117
2.6.6 Blood-based PET Radiotracer Arterial Input Function (AIF) Estimation	119
Chapter 3 MR-based Methods for Addressing Specific Challenges to PET Quantification in	
PET/MR	122
3.1 Attenuation Correction	122
3.1.1 Head Attenuation Map Estimation	123
3.1.2 MR-Hardware Attenuation Map Estimation.....	131
3.2 Partial Volume Effects Correction	132
3.2.1 Applying a PSF to an ROI	133
3.2.2 Geometric Transfer Matrix	135
3.2.3 Region-based Voxel-Wise Correction	136
3.3 Head Motion Compensation	137
3.3.1 Rigid-body Motion Estimation	138
3.3.2 Rigid-body Motion Correction	155
Chapter 4 Assess the Added Value of Improved PET Data Quantification for Neurological	
Applications	174
4.1 Analysis of MR-derived Interference Effects on FDG-PET Brain Quantification in	
Human Studies	175
4.1.1 Materials and Methods.....	177
4.1.2 Results.....	183
4.1.3 Discussion.....	191
4.1.4 Conclusion.....	197
4.2 Impact of PET Data Optimization on NNC112 Binding Potential Estimation	198
4.2.1 Materials and Methods.....	199
4.2.2 Results.....	201
4.2.3 Discussion.....	206
4.3 MR-guided Image-derived PET Radiotracer Input Function Estimation	208
4.3.1 Materials and Methods.....	210
4.3.2 Results.....	220
4.3.3 Discussion.....	225
Chapter 5 Conclusion.....	228
Bibliography.....	230
Appendix.....	244
BrainPET Quality Control Experiment Protocols.....	244

Masamune System State Executable Commands.....	245
Useful Commands for PET data processing	245

Figures

Figure 1 Positron Emission and Annihilation	4
Figure 2 Example of a Sinogram	6
Figure 3 PET Image Reconstruction	9
Figure 4 Attenuation Effects on a Uniform Phantom	11
Figure 5 The Two-Tissue Compartment Model	25
Figure 6 Methods of PET Analysis.....	29
Figure 7 Partial Volume Effect	34
Figure 8 Multiparametric PET/MR Imaging of Glioblastoma	50
Figure 9 Time-of-Flight MR Angiography.....	51
Figure 10 The Siemens BrainPET	52
Figure 11 BrainPET Crystal Orientation and Sinograms	57
Figure 12 Sinogram Gap Filling Due to Motion Correction	58
Figure 13 Reconstruction with Resolution Modeling	60
Figure 14 BrainPET Point Spread Function	62
Figure 15 Gaussian and Radial Blur Images	65
Figure 16 Determination of Optimal Radial Motion Blur	66
Figure 17 PET Reconstruction Effective Radius Estimation	68
Figure 18 Reconstruction Optimization Graphical Layout.....	69
Figure 19 Derenzo Profile Models	71
Figure 20 Effective Scanner Radius Estimation	72
Figure 21 Effect of Iteration Number of Image Quality.....	73
Figure 22 Iteration Number Optimization	74
Figure 23 Percent Bias and Noise Versus Iteration Number	76
Figure 24 Dynamic Reconstruction Effects.....	78
Figure 25 BrainPET Photopeak Spectra During MR Imaging.....	82
Figure 26 PET Count Rate Fluctuation	84
Figure 27 Spatial Offset of MR and PET	86
Figure 28 Multi-modal Temporal Alignment	89

Figure 29 PET Block Design and SD Regions	92
Figure 30 Normal DCmap.....	93
Figure 31 Faulty APD's Effect on the DCmap.....	94
Figure 32 DCmap of Defect in Middle-Edge of Block	95
Figure 33 Automated Bolus Arrival Time.....	97
Figure 34 Nonuniform injection	98
Figure 35 Nonuniform injection fit with 4-part model.....	99
Figure 36 Uniform injection fit with 4-part model.	99
Figure 37 Comparison of 4- and 3-part BAT models.	100
Figure 38 PET/MR Graphical User Interface.....	104
Figure 39 FreeSurfer-derived ROIs	107
Figure 40 SPM-derived ROIs	110
Figure 41 Mean Gray Matter Contribution to Brodmann Atlas	111
Figure 42 ROI-based Intensity Normalization.....	113
Figure 43 ROI-driven FDG-Intensity Normalization Across Subjects	114
Figure 44 Automated Reconstruction and Analysis Workflow.....	119
Figure 45 Implemented Mu Map Options	123
Figure 46 PET-based Attenuation Correction GUI.....	124
Figure 47 Partial Volume Effect Correction Processing for Region of Interest.....	134
Figure 48 Rim Artifacts in RBV Images Without rPSF Modeling....	137
Figure 49 Effect of B_0 inhomogeneity on EPI registration	142
Figure 50 Effect of Alternative Field Mapping Scheme on Motion Estimates.....	144
Figure 51 Representative MR-based Motion Estimates for FDG-PET/MR Study.....	149
Figure 52 PET-based Motion Correction Algorithm	152
Figure 53 PET-based MC Estimates for Dynamic FDG-PET/MR Study	154
Figure 54 Motion Correction Algorithm	157
Figure 55 MC Reconstruction Effect.....	163
Figure 56 MC Reconstruction Simulation	164
Figure 57 Effects of different affine conventions.....	167

Figure 58 Averaging Motion Estimates.....	168
Figure 59 Percent volume overlap between spheres using proper and arithmetic averages	173
Figure 60 Design of paradigm, processing, and analysis for simultaneous FDG-PET/MR experiment to illicit MR-based physiologic effects on FDG-PET uptake	180
Figure 61 Creation of high-resolution ROIs for PET imaging	183
Figure 62 TACs	184
Figure 63 SPM8 Glass Brain of parametric image	186
Figure 64 Cortical FreeSurfer analysis of parametric image	188
Figure 65 SPM8 Glass Brain of static image.....	189
Figure 66 Cortical FreeSurfer analysis of static image.....	191
Figure 67 NNC112 Motion Profiles.	201
Figure 68 Effect of Motion Correction on TACs.....	202
Figure 69 Partial Volume Effect Correction for a representative NNC112 subject	202
Figure 70 Effects of Motion and Partial Volume Effect Correction on a Single Subject.....	203
Figure 71 Relative Change of BPnd Due to Corrections	204
Figure 72 BPnd by Motion Amplitude	205
Figure 73 BPnd Relative Change By Motion Amplitude	206
Figure 74 Intrasequence Motion for a Single Subject Derived from MR Estimates.....	207
Figure 75 TOF Morphological Filtering	211
Figure 76 Multi-step Seeded Region Growing of Arterial Mask....	213
Figure 77 Multi-MR-Sequence Derivation of Arterial Mask.....	220
Figure 78 Representative Image-Derived Input Function	221
Figure 79 Kinetic Parameter Estimation Using AIF and IDIF for a Representative Subject.....	222
Figure 80 Comparison of Late Time-Point Blood Sampling Using the Injection Catheter to Arterial Line Measurements.....	223
Figure 81 Representative TOF+MPRAGE PET IDIF Results for FDG, NNC112, and PBR28.....	224

Tables

Table 1 PET Radiotracers for Neuroimaging.....	37
Table 2 BrainPET System Specifications	54
Table 3 Parameters of Derenzo Profile.....	71
Table 4 Mean Multimodal Offsets	87
Table 5 Mean Offsets for MC Reconstruction Simulation	164

Equations

(1.1).....	10
(1.2).....	13
(1.3).....	15
(1.4).....	16
(1.5).....	17
(1.6).....	17
(1.7).....	18
(1.8).....	18
(1.9).....	18
(1.10).....	19
(1.11).....	22
(1.12).....	30
(1.13).....	32
(1.14).....	38
(1.15).....	38
(1.16).....	38
(1.17).....	38
(1.18).....	39
(1.19).....	39
(1.20).....	39
(1.21).....	40
(2.1).....	56
(2.2).....	61
(2.3).....	61
(2.4).....	61
(2.5).....	61
(2.6).....	70
(2.7).....	75

(2.8).....	75
(2.9).....	96
(2.10).....	98
(2.11).....	121
(2.12).....	121
(3.1).....	127
(3.2).....	127
(3.3).....	127
(3.4).....	127
(3.5).....	129
(3.6).....	135
(3.7).....	135
(3.8).....	136
(3.9).....	136
(3.10).....	137
(3.11).....	153
(3.12).....	153
(3.13).....	153
(3.14).....	153
(3.15).....	158
(3.16).....	161
(3.17).....	162
(3.18).....	162
(3.19).....	169
(3.20).....	169
(3.21).....	170
(4.1).....	200
(4.2).....	214
(4.3).....	216

Chapter 1 INTEGRATED PET/MR IMAGING

Tomography, derived from the Greek *tomos* (section) and *graph* (write) refers to a collection of methods that produce three dimensional images of the internal structure of “solid” objects through the use of penetrating waves. The nature of these imaging modalities makes them appealing for biological imaging as they range from minimally invasive to completely noninvasive and allow for the organism of interest to be imaged in its nascent state. Since the introduction of computed tomography (CT) by Godfrey Hounsfield in 1972, tomographic imaging has become standard in medicine and today includes magnetic resonance imaging (MRI), ultrasound, single photon emission tomography (SPECT), and positron emission tomography (PET). Tomographic imaging methods can be separated broadly into methods that resolve the structure of tissues (structural imaging) and methods that resolve the function of tissues (functional imaging). Some tomography devices, or tomographs as they are called, are capable of both structural and functional imaging, while others are only capable of structural or functional of imaging. The diversity of tomographic techniques, or modalities, that exist for medical imaging is due to the inherent strengths and weaknesses of the different modalities. Since no single modality is perfect for every situation, considerable attention has been invested

in the combination of multiple modalities in an effort to capitalize on their individual strengths and mitigate their weaknesses [1].

Tomography is routinely performed on every organ and tissue of the body; nevertheless, the brain is an organ of particular interest due to its involvement in many homeostatic functions, its broad range of biochemical and physiologic pathways, and its need for noninvasive diagnostic techniques. While structural imaging allows for noninvasive diagnosis of atrophy, neurodegenerative diseases, and detection of tumors, many cognitive/psychiatric disorders have been shown to have stereotyped functional patterns on imaging, making functional brain imaging useful in their diagnosis and treatment [2]. As more advanced and powerful scanners are becoming available, considerable attention has been directed at merging structural and functional information to better understand the mechanism of disease and provide a more complete characterization of the state of the brain for diagnosis [3].

This work is concerned with the application of devices that combine PET, a functional imaging modality, with MRI, a powerful modality capable of both functional and structural imaging, for neuroimaging and how the combination of MRI and PET can improve the quality and quantification of the PET data. This chapter reviews the basic concepts of PET imaging and quantification, with particular attention paid to corrections that need to be applied to the data for accurate quantification and effects that can degrade the reliability of the PET data. Pertinent concepts of MRI, the history of simultaneous PET/MR systems, and unique challenges in these multimodal systems are also addressed.

1.1 Positron Emission Tomography

This section provides an overview of PET with a focus on the phenomena and corrections that must be addressed for accurate estimation of the radiotracer distribution. Subsequently, the correlation between the time-varying radiotracer distribution along with its biochemical interpretation is discussed, followed by factors that can reduce the reliability of the inferences on the biochemical interpretation. Finally, examples of the application of PET to neuroscience and clinical imaging are presented.

1.1.1 Physical Principles

PET is a nuclear medicine technique that uses ionizing radiation to noninvasively image molecular aspects of a tissue *in vivo*. Specifically, PET cameras image the spatial distribution of radioactivity through the detection of coincident 511 keV photons (gamma rays) produced from the annihilation of positron-electron pairs. Endogenous or exogenous molecules are altered to contain a positron emitting radioisotope (e.g. ^{18}F , ^{11}C , ^{15}O) and are administered intravenously into the subject. The radioisotope undergoes radioactive decay via beta-plus emission wherein a proton is transformed into a neutron and a positron is released from its nucleus (along with a neutrino). The specific kinetic energy that the positron is released with is dependent on the radioisotope and is related to the isotope's energy. As the positron travels through the tissue it inelastically interacts with surrounding atoms, losing its kinetic energy, typically over the distance of a few millimeters. Once the positron has lost all of its kinetic energy and comes to rest, it can interact with a nearby electron. The particles pair and form a metastable positronium intermediate which quickly undergoes annihilation, converting the particles to a

pair of coincident photons. To conserve energy and momentum, each photon has an energy of 511 keV and the two travel in opposite (antiparallel) directions. A cartoon of the decay process can be found in Figure 1.




Figure 1 Positron Emission and Annihilation. A proton of the radionuclide decays through the emission of a positron. The positron travels through space, losing its kinetic energy until it interacts with an electron at which point the particles annihilate and two antiparallel photons are created. Particles not drawn to scale.

The three-dimensional distribution of the radioisotope is not directly measured in PET; rather, it is estimated from a series of projections that correspond to the paths of the annihilation photons. The annihilation photons are detected in photodetectors, dense materials and/or electronics that convert high energy photons into an electrical signal. If a pair of annihilation photons are stopped by a pair of photodetectors and the photons travel antiparallel from their point of annihilation, then the line between the photodetectors – referred to as the Line-of-Response (LOR) – contains all the possible locations in space where the annihilation could have taken place. If both photons are detected within a certain time interval, called the coincidence window, they are considered to have originated from the same annihilation event. These events, called coincidence event, are assigned to the corresponding LOR. For a ring of detectors, integrating the coincident counts in a set of parallel LORs over a

period of time yields a one-dimensional projection (histogram) along that direction. This direction can be represented as an angular offset from some arbitrary direction. The one-dimensional histograms for different angles can then be reorganized into a two-dimensional matrix called sinogram in which each position, or bin, represents an individual LOR. A sinogram of a digital phantom can be found in Figure 2. Mathematically this transformation of an object into its projections is known as the Radon transform, and in tomography the transformation of an object to its sinogram is known as a forward projection. For systems composed of multiple rings of detectors, three-dimensional sinograms can be generated. The organization of three-dimensional sinograms is reviewed elsewhere [4]. If coincidences are accepted between detectors in different rings, there will be a redundancy in sampling the field of view as points in space are covered multiple times by direct (in-plane) LORs and those from oblique detectors. With more LORs passing through a given point in space, the sensitivity of the scanner improves. Images can then be generated from these three-dimensional sinograms through a number of techniques, ranging from the purely analytical approaches to methods that attempt to iteratively determine the distribution of the radioisotope and are exhaustively reviewed elsewhere [5]. To simplify and speed up the reconstruction, rather than the sinogram bins directly correspond to the LORs with a 1-to-1 mapping, multiple LORs can be assigned to the same sinogram bin, e.g. neighboring LORs or those that are oblique with regards to the ring by a few crystals.

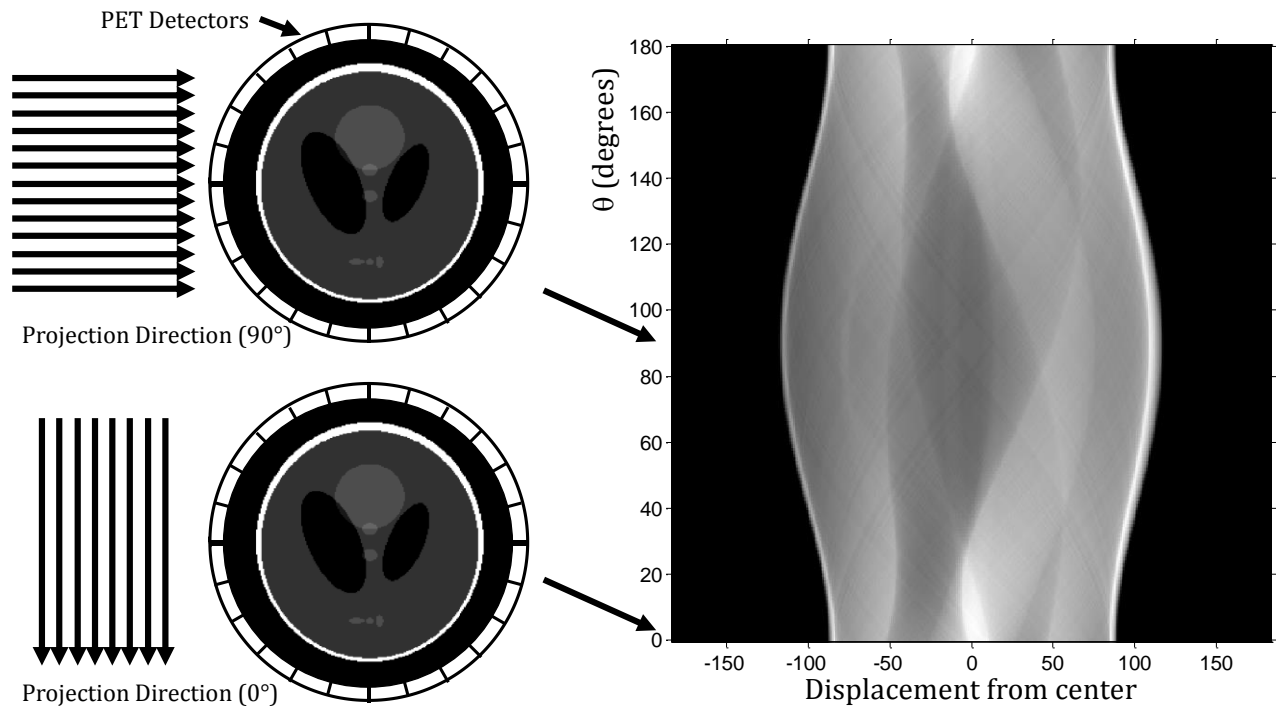


Figure 2 Example of a Sinogram. The PET crystals define lines-of-response where detected coincidence events are assigned. Parallel sets of LORs define a row in the sinogram where the angular offset of the projection defines the position of the row. Only projections between 0° to slightly less than 180° are considered in PET imaging.

Two of the defining characteristics of PET are its quantitative nature and its high sensitivity (in the picomolar range). However, in order to accurately estimate the distribution of the radioisotope, a number of corrections must be considered. For instance, photons can be absorbed or deflected by atoms, preventing them from being detected in the photodetectors. Similarly, if the course of the photon is altered, the coincidence event could be assigned to the incorrect LOR. Additionally, the detection efficiency for the photodetectors and associated hardware must be considered.

1.1.2 Image Reconstruction and Correction Factors

As previously mentioned, there are a number of techniques that can resolve the radiotracer distribution from emission sinograms. The simplest approaches, back-projection

and filtered back-projection (FBP), analytically invert the Radon transform directly to resolve the image. While FBP is fast and easy to implement, it also has some associated challenges. One challenge is the sampling requirements; to minimize artifacts, FBP requires a fine linear and angular sampling (i.e. that the sinogram be sampled over 180°) and, most importantly, that the data be complete which can present a problem (typically the presence of artifacts) if there are sizeable gaps between the photodetectors.

An alternative to FBP is to guess the distribution, compare the results to the measured emission data, and update the guess, iteratively, to approach the true image. These methods comprise the iterative reconstruction algorithms. Often the initial estimate is a blank or uniform image. The estimated image is then forward projected to produce a sinogram. The estimated and emission data are then compared and the difference is estimated through a *cost function*, like ordinary Poisson statistics (OP). Based on the cost function, the estimate can be updated via an update function and a new estimated image can be produced and the cycle can begin again. Eventually the estimated image will produce a difference that falls below a specified level and the process will stop. Iterative reconstruction algorithms are named based on their cost and update functions. Popular update functions include maximum likelihood expectation maximization (MLEM) and an adaptation that uses ordered subsets to improve speed, ordered subset MLEM (OSEM). More advanced reconstruction techniques like the maximum *a posteriori* (MAP) can utilize anatomic information from sources like MRI to resolve the activity distribution [6].

The sinogram corresponding to the coincident photons emitted from the radiotracer distribution is called the true sinogram as it contains all the true events. What is measured by

the PET camera is the emission sinogram that represents the photons absorbed by the photodetectors. In order to accurately resolve the true radiotracer distribution it is necessary to account for intrinsic and extrinsic factors that may affect the measured counts in each sinogram bin. An example of an intrinsic factor is the individual sensitivity of each photodetector, which will directly affect the sensitivities of all LORs that include it. An extrinsic factor would be the direct absorption or scattering of photons as they travel through the object to the photodetectors. The common corrections to the emission sinogram, along with how some of them are derived can be found in

Figure 3. The following subsections will discuss the major necessary corrections to the emission sinogram and reconstructed image.

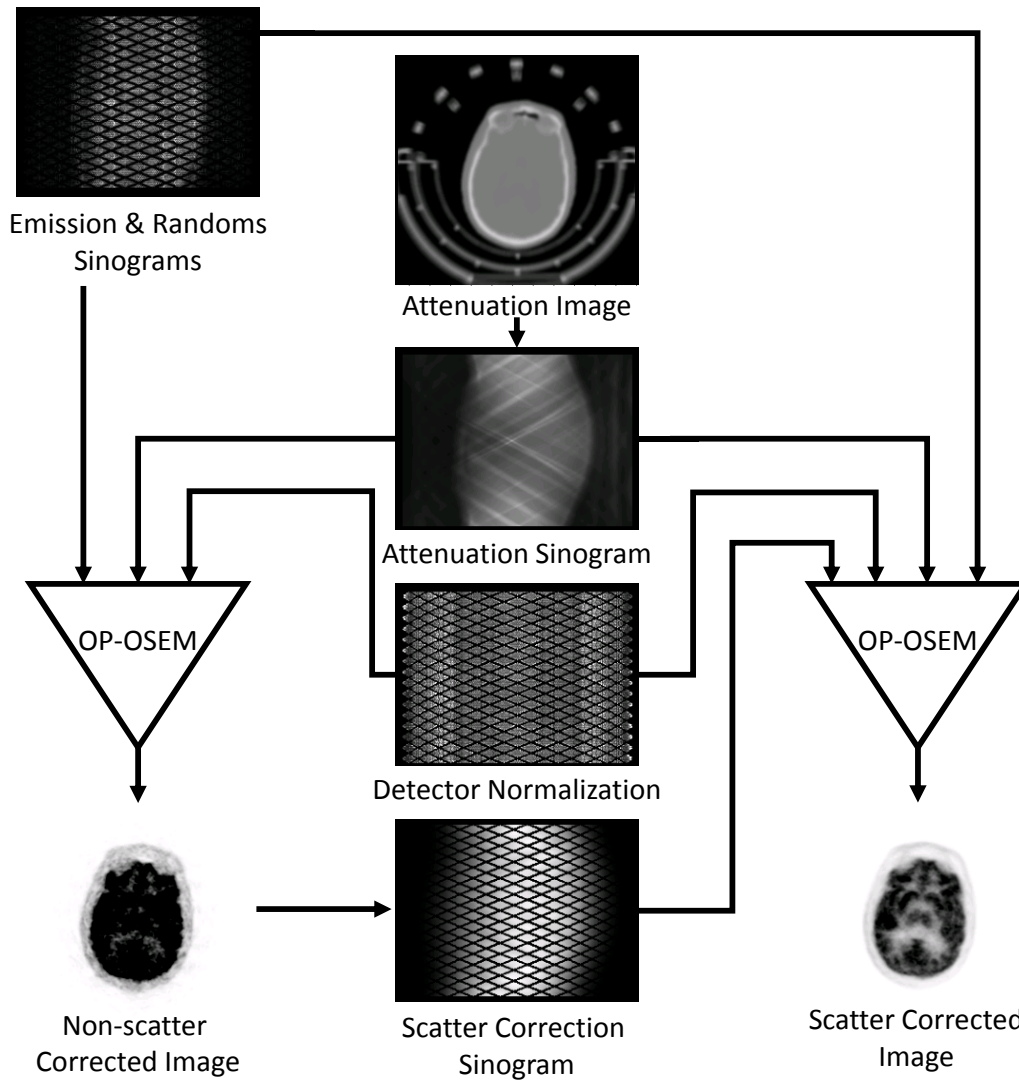


Figure 3 PET Image Reconstruction. Hardware effects, lost photons (attenuation) and misplaced photons (scatter) must all be accounted for in order to estimate an accurate estimation of the tracer distribution.

1.1.2.1 Attenuation Correction

After the coincident photons are emitted, there is a finite probability that, as they travel through space between their point of origin and the photodetectors, they will be absorbed or deflected away by some material, both of which result in a loss of counts in a LOR. The net result of the attenuation of the photons is the underestimation of the activity, which becomes more pronounced further from the surface. An example of the effect of attenuation on a

phantom of uniform activity with a uniform attenuation can be found in Figure 4. Attenuation effects can severely impact the values of the reconstructed PET image, and thus accurate correction is necessary.

If the attenuation properties of the object can be determined, then a more realistic radiotracer distribution can be resolved. The probability of both photons successfully passing through a LOR to the photodetectors, p_{coinc} , where the attenuating object exists along the length D of the LOR is given by

$$p_{coinc} = \int_{x_1}^{x_2} \exp(-\mu(x)dx) \quad (1.1)$$

where μ is the linear attenuation coefficient (in units of cm^{-1}) which may be spatially varying, and $|X_1 - X_2|$ is D . If an attenuation map (or mu map) is created, an image of the object through which the radiotracer is distributed where each voxel is the linear attenuation coefficient at that point in space, then a sinogram representing the effective attenuation through each LOR is merely the result of forward projecting the object and taking the negative exponential of the resulting values. The effects of attenuation on the emission sinogram can be removed by dividing the emission sinogram by the attenuation sinogram. In PET neuroimaging three main classes of attenuating tissue types exist for the 511 keV photons: bone (highest attenuation; $\mu=0.12-0.15$), soft tissue (e.g. brain, fat, etc; $\mu \approx 0.098$), and air (no attenuation) [7, 8].




Figure 4 Attenuation Effects on a Uniform Phantom

Before the advent of combined PET/CT systems, the attenuation correction for the PET data was determined through transmission (Tx) scans, whereby a radioactive rod source (e.g. $^{68}\text{Ge}/^{68}\text{Ga}$) was rotated around the body. In PET/CT systems, the time-consuming transmission scan (~ 3 minutes per field of view or 18 minutes to cover the entire body) is replaced with a CT scan which is performed faster (~ 35 seconds to cover the entire body), has lower noise, can be done post tracer injection (since the flux of CT photons is significantly higher than PET photons), and does not require the periodic replacement of decayed transmission sources or related equipment [9, 10]. Using CT images as a surrogate for Tx scans could be done because CT measures effects related to electron densities which also compose a majority of the PET signal attenuation. Despite its drawbacks, one of the benefits of the Tx scan is that only mono-energetic 511 keV photons are present. This is important as the attenuation coefficient is energy dependent and the Tx photons are equal in energy to the annihilation photons from the PET tracer. CT photons, on the other hand, have a much broader energy spectrum ranging from 40 keV to 140 keV. In order to determine the attenuation coefficient, it is necessary to choose a single effective energy to represent the CT spectrum [11]. There are three ways to generate the 511 keV attenuation coefficients from the CT data: 1) each voxel can be scaled using the

ratio of attenuation coefficients for a known substance at the two energies (e.g. water); 2) a threshold can be applied to the data to determine the tissue type (bone, soft tissue, cerebrospinal fluid) and a single value assigned to each tissue; or 3) a combined approach can be used [12-14]. The hybrid approach has been shown to be the most accurate method as it first separates bone from soft tissues before applying scaling factors, which have significantly different scaling factors due to the high concentrations of calcium and phosphorous in bone [15].

In contrast to CT, standard MR images cannot directly interrogate the probability of photon attenuation because MR images reflect proton density rather than electron density. In fact, the solid, crystalline structure of bone (the tissue with the highest LAC) leads to a short T2 which makes bone indistinguishable from air on most MR images. This presents a problem in integrated system PET/MR systems where space is limited and the addition of a rotating Tx source is not possible. An approach analogous to the Tx scan has been presented, where the attenuation map is constructed by placing an annulus-shaped phantom around the subject and bed; however, it must be filled each time it is used which is impractical, subjects the subject to a higher cumulative dose, and is limited to Time-of-flight PET scanners [16]. Alternatively, a number of schemes have been developed to construct attenuation maps from MR images. Early approaches relied on complicated sequential regional segmentation and applied a single value to the segmented tissue classes [17]. With the advent of ultra-short echo-time (UTE) MR imaging, which were developed to image tendon and bone tissues [18, 19], the principal attenuation classes could be derived directly from the MR images [20, 21]. Another set of methods relied on an atlas based approach where the MR image could be registered to a

template based on a collection of patients with corresponding attenuation information [22-24]. The template MR volume, and corresponding attenuation information, could be applied to the patient MR using non-rigid transformations. This method has the benefit of using Tx or CT data for the attenuation information which allows for a continuous distribution of values rather than a discrete-valued segmentation approach. Recently, more advanced approaches, which combine atlas-based information with UTE or similar MR imaging sequences have been developed [25].

1.1.2.2 Scatter Correction

While attenuation correction increases the signal in an attempt to account for lost coincident events due to absorption or scattering, scatter correction seeks to remove events that have been placed in improper LORs as a result of Compton scattering. Compton scattering is the predominant interaction for 511 keV photons with elements that have an atomic number of approximately 50 and below, which includes most of the elements that make up the body [26, 27]. Compton scattering describes the particle-like behavior of a photon wherein it collides with an outer shell electron, liberating the electron. Unlike the photoelectric effect where the photon is consumed in the interaction, in Compton scattering the photon's course is deflected. Since energy must be conserved, the scattered photon's energy is lowered by the amount lost to the electron. This loss can be related to the deflection angle of the electron and its resulting energy, E_{sc} ,

$$E_{sc} = \frac{E_0}{\left[1 + \frac{E_0}{511}(1 - \cos \theta)\right]} \quad (1.2)$$

where E_0 is the incident photon energy and θ is the deflection angle [28].

Due to the relatively poor energy resolution of LSO, the detection electronics typically accept photons with an energy range between 350 to 650 keV [29]. As such, photons that have undergone a single Compton scatter event and have altered their trajectory by as much as 45° still have sufficient energy to be considered an annihilation photon. Assuming one, if not both photons are scattered, the resulting LOR defined by the detection of the coincident photons may not pass remotely near where the initial radioactive decay occurred. In some 3D acquisition PET scanners, up to 40% of the detected events have undergone Compton scattering [30]. To accurately estimate the true radiotracer distribution, the contribution to these LORs due to scattered photons must be removed.

There are a number of techniques to estimate the scatter, including: empirical techniques where the tails of measured data in sinogram space are fit with functions; multiple energy window techniques where the scatter is estimated from portions of the photon spectra; convolution approaches; and simulations.

While scatter can be estimated in several ways in PET studies, the fact that the physics behind scatterings is so well understood makes the simulation based scatter approaches very robust and reasonably accurate to within less than 1% [31]. The modeling of multiple scatter events is complex and computationally intensive for practical application so only single scattered events are considered. Simplifying the system to only include single scatter events has been demonstrated as a reasonable approximation as it has been shown that 75-80% of scattered coincidences arise from single scattered events for axial lengths of 10 cm [32, 33]. A popular technique for modeling the scatter is the analytical single scatter simulation technique

(SSS) [30, 34, 35]. Here, R_s , The expected coincidence rate between a pair of detectors (A and B) due to a single scatter event arising at point S is given by an integral of the Klein-Nishina formula,

$$R_S = \int_{V_s} dV_s \left(\frac{\sigma_{AS} \sigma_{BS}}{4\pi R_{AS}^2 R_{BS}^2} \right) \frac{\mu d\sigma_c}{\sigma_c d\Omega} (I^A + I^B) \quad (1.3)$$

where

$$I^A = \varepsilon_{AS} \varepsilon'_{BS} e^{-\left(\int_s^A \mu ds + \int_s^B \mu' ds \right)} \int_S^A \lambda ds,$$

$$I^B = \varepsilon'_{AS} \varepsilon_{BS} e^{-\left(\int_s^A \mu' ds + \int_s^B \mu ds \right)} \int_S^B \lambda ds, \quad (1.3)$$

μ is the attenuation coefficient (evaluated at 511 keV for the unprimed values and E_{sc} at the primed values), R_{*s} are the distances from the detectors (A or B) to the scatter point, ε_{*s} the respective efficiencies, σ_{*s} the cross sections as evaluated from the scatter point, λ is the emitter intensity, σ_c is the Compton scatter cross section as calculated from the Klein-Nishina formula, and Ω is the solid angle. The theory and implementation of the SSS is well detailed elsewhere [30]. The SSS requires three pieces of information to estimate the scatter: the attenuation volume to derive μ , an estimate of the emission volume (typically corrected for normalization and attenuation) to derive λ , and the detector sensitivity (usually in the form of the normalization) to derive ε . To reduce the computational complexity, typically only a subset of the image is used in the form of sparsely separated points and the data is interpolated, requiring an additional scaling step to ensure a good fit between the measured and interpolated points.

The subtraction of scattered events, along with the subtraction of random coincidence events, recorded events from photons that do not arise from the same annihilation event, is performed in sinogram space. The specific equation that relates the true sinogram that represents the tracer distribution, T , to the recorded emission sinograms depends on the formulation of the algorithm; however, for the specifics of this work the equation is given by:

$$T = [(P - D) \times N - S] \times A \quad (1.4)$$

where P is the prompt sinogram (the emission sinogram), D is the delays (an estimate of the random coincidence events), N is the normalization, S is the scatter sinogram, and A is the attenuation sinogram.

1.1.2.3 Count-rate Corrections

In addition to the extrinsic corrections associated with objects in the field of view (discussed above), additional intrinsic corrections accounting for the performance of the scanner and properties of the radiotracer are necessary. An understanding of these factors is important as they can affect the shape of the dynamic PET signal (Section 1.1.3.3) and in turn quantitation, introduce artifacts in the data, and may vary between radionuclides. These corrections should be considered in the design of any PET experiment and the imaging limitations they can introduce. These corrections are discussed last as they are typically applied after the reconstruction of the PET image as global scaling factors.

1.1.2.3.1 Decay Correction

The most intuitive of the count-rate corrections is the decay correction. The signal detected in PET is a result of a radioactive decay, and therefore is related to the number of

radioisotopes present at the beginning of the measurement or frame, where the frame can be defined by the time that the measurement is started and the duration of the measurement. Stated another way, given a uniform volume of a molecule that contains a radioisotope, the signal will change over time as the radioisotopes decay away and successive frames will detect fewer and fewer photons while the concentration of the molecule will remain constant. In most scenarios it is the accumulation of the molecule, and not that radioisotope, that is of interest; however, the concentration of the molecule is inferred from the radioisotope concentration.

There are two forms of decay correction, *inter-* and *intraframe* decay correction.

Interframe correction relates an instantaneous estimate of the activity or a series of estimates at different instances in time to a common time point. The amount of a radioactive substance at a time t is given by

$$N(t) = N(t_0)e^{-\frac{\ln 2}{\tau_{1/2}}t} \quad (1.5)$$

where $N(t_0)$ is the amount at some earlier time and $\tau_{1/2}$ is the half-life of the isotope. Typically it takes time for the radioisotope to distribute throughout the body and the total activity in the body will be less than the injected activity. The lower total activity results in an underestimation of the concentration of the molecule for all frames following the injection as the distribution of the molecule at any point in time is relative to the total amount of radioisotope present at that time. To determine the measured activity distribution at a given time relative to a different moment in time, t_0 , the measured activity A can be multiplied by a correction factor C

$$A(t)_{t_0} = A(t)_t C = A(t)_t e^{\frac{\ln 2}{\tau_{1/2}}(t-t_0)} \quad (1.6)$$

Intraframe decay correction accounts decay over the duration of the frame. The PET image does not represent an instantaneous measurement, but rather the accumulated counts recorded over a period of time on the order of seconds to minutes. PET measures coincident photons that result from the annihilation of electrons-positron pairs, which in turn result from the radioactive decay of an atom. For a given concentration of an undecayed radioisotope, the probability of any individual atom undergoing decay is randomly distributed in time, and is statistically independent from all other atoms and for any non-overlapping time intervals. As such, the probability of a radioactive decay follows a Poisson statistic and the probability of n decay events in a short time interval Δt is given by

$$P(n; \Delta t) = e^{-\Gamma \Delta t} \frac{(\Gamma \Delta t)^n}{n!} \quad (1.7)$$

where Γ is the rate of decay in a unit time and is related to the half-life, $\tau_{1/2}$, and the total activity at the beginning of the time interval, $N(t)$, thusly,

$$\Gamma = N(t) \frac{\tau_{1/2}}{\ln 2} \quad (1.8)$$

Following Equation (1.7), to acquire a sufficient number of coincidence counts to generate an image the length of the frame may be on the order of minutes. For an image of duration Δt , the intraframe decay correction is given by,

$$R = \frac{\ln 2}{\tau_{1/2} \left(1 - e^{-\frac{\ln 2}{\tau_{1/2}} \Delta t} \right)} \Delta t \quad (1.9)$$

And the overall decay correction factor, RC, for a frame of duration Δt with start time t is given by

$$RC = \frac{\ln 2}{\tau_{1/2} \left(1 - e^{-\frac{\ln 2}{\tau_{1/2}} \Delta t} \right)} \Delta t e^{-\frac{\ln 2}{\tau_{1/2} (t-t_0)}}. \quad (1.10)$$

If a single decay correction factor is used for a frame, then the decay correction can be applied to either the delay and prompt sinograms or the reconstructed image with similar results.

1.1.2.3.2 *Deadtime Correction*

The conversion of the high energy photons to electronic signals is accomplished through a number of distinct hardware elements. Each element requires a finite amount of time for the processing of an individual event, the deadtime, with the signal integration electronics typically having the largest effect [36]. During the deadtime the system electronics ignore subsequent events arising at the photodetector level, even for photodetectors not involved in the initial event detection. During this time there is a probability that a true event may be lost because it occurs while the system is processing another event. Ignoring these events leads to nonlinearity in the detection of the PET camera at high count rates. This effect sets a limit on the maximum radiotracer that can be present in the PET field of view. In addition to the events being ignored in regards to counting, they can have an additional impact on when future events can be detected. Counting electronics can be categorized into those that have a paralyzable and nonparalyzable deadtime. Those with a paralyzable deadtime only ignore events that occur during the deadtime window set by the initializing event. Nonparalyzable deadtime, as is the case with PET scanners, will introduce deadtime according to any event that it detects. Therefore, if a second event occurs just before the end of the deadtime window, the total time before the camera will accept another event for counting will be approximately twice the

deadtime. The cumulative effect for nonparalizable systems is that the error at high count rates will be magnified as events detected during the deadtime will be ignored and the total deadtime will be larger than the characteristic deadtime.

1.1.2.3.3 *Pileup Correction*

At high count rates multiple annihilation events can be detected in the same photodetector in such a way that the system cannot distinguish them from a single event. At sufficiently high count rates, multiple photons that have undergone large angle Compton scatter events, and individually fall below the lower energy threshold of the system, can mistakenly be counted as a real event. This in turn can introduce artifacts in the image. Pileup effects are typically addressed by minimizing the activity used in the study, however hardware-based methods also exist [37, 38].

1.1.3 Quantitative PET Imaging

PET is a quantitative imaging modality. In a properly calibrated system, the signal in a given voxel of unit size is the concentration of disintegration events over the imaging duration. This signal can in turn be directly converted into the absolute concentration of the radioisotope in that voxel. One of the most appealing aspects of PET is its molecular specificity, making it possible to measure a diverse set of biological processes in the body. While PET has the potential to interrogate many biochemical pathways and systems, the PET images alone rarely are capable of providing this information directly; rather, PET only resolves the spatial distribution of a radioisotope or radioligand. In the case of fluorodeoxyglucose (FDG), this equates to the subtle but important difference between the accumulation of glucose molecules

in a tissue and the metabolism of glucose by that tissue. The biologically relevant effect, the metabolic rate of glucose, can be estimated from the dynamic PET data with the aid of a model that describes the biochemical and physiologic processes, referred to as a kinetic model, as it is used to estimate the kinetics of the system. Kinetic models typically require an estimation of the free radiotracer or input function which may be measured through serial sampling of the arterial blood, which adds a level of complexity to the study protocol. To simplify the protocol at the sacrifice of some of the quantitation, semi-quantitative methods, which are directly related to the tracer accumulation, may be used. Semi-quantitative analyses are routinely used in the clinic with PET imaging as they do not necessarily require dynamic imaging, may not require an input function, are typically shorter in duration than dynamic scans, and use information that is easy to measure in a clinical setting (e.g. height, weight, pre-injection blood-glucose levels).

1.1.3.1 Qualitative and Semi-quantitative Imaging

The most basic use of PET imaging for clinical applications involves qualitative visual interpretation of the images. In these cases the subject is injected and the tracer is allowed to circulate for a given time, known as the uptake time, accumulating in cells, and then a single image of the tracer localization is acquired. Here interpretation is based on relative intensity or relative change in a region as compared to surrounding structures over longitudinal scans. While the simplest form of analysis, visual interpretation has been shown to provide clinical utility. With regards to oncologic applications, the complete resolution of abnormal FDG uptake, which can be made through visual interpretation, was a strong predictor of

progression-free survival [39]. However, accurate visual interpretation requires reproducibility in the injected dose, pre-scanning procedure/environment/routine, and that the metabolism of the patient remains relatively constant over time, save for the regions of interest. Visual interpretation also fails when comparing different subjects as basal metabolic rates and blood flow may differ as well as other factors that may influence uptake. To help standardize across subjects and account for subtle changes across visits for a single subject, the Standard Uptake Value (SUV), a normalized semi-quantitative measure, was constructed. The SUV is given by

$$SUV = Q \times \frac{W}{Q_{inj}} \quad (1.11)$$

where Q is the measured activity of a voxel in an image, W is the patient weight, and Q_{inj} is the injected activity. SUVs have been shown to have a dependence on weight requiring corrections based on lean-body mass or body surface area for studies involving comparisons across diverse body types or in longitudinal studies where significant changes in body weight are observed [40-43]. Similarly, corrections based on circulating blood-glucose levels before the FDG injection have been used to correct FDG data [43-47]. Regardless, as both qualitative and semi-quantitative methods rely on a single time-point measurement, the contribution of tracer delivery cannot be disentangled from tracer uptake (i.e. tissues with a high tracer uptake but low delivery may appear similar to those with high delivery but low uptake).

1.1.3.2 Kinetic Modeling

The spatial distribution of the radiotracer in the body can depend on a number of factors, such as: infusion rate, delivery to the tissues, binding to a receptor, transport into and out of the cell, metabolism, etc. There is a temporal component to the tracer accumulation,

which can provide a measurement of the biologically relevant quantity and information about the kinetics and rates of the system. The kinetics of the system and related rates are estimated by representing the biological process of interest by a mathematical model called a kinetic model.

An early formalism for kinetic modeling was first presented by Sokoloff *et al.* [48] in the context of a glucose analogue [¹⁴C]Deoxyglucose (CDG) and can be used to determine kinetic parameters that describe a series of biochemical processes. The model is defined by a set of compartments, volumes where a tracer quickly distributes itself uniformly (i.e. there are minimal concentration gradients), and can have physical interpretations (e.g. the vasculature) or not (e.g. phosphorylated or unphosphorylated compound). The radiotracer concentration in each compartment is a function in time that is dependent on concentrations in the other compartments. This model can be further abstracted as a set of mathematical equations where the kinetic parameters are the rate constants at which a compound moves from compartment to compartment and the concentrations in the compartments can be represented through sets of ordinary differential equations related to mass balances. The model is a simplification of the exact biology; however, it has good predicative power. An example of a kinetic model is shown in Figure 5. The object of interest is the radionuclide rather than the tracer itself as the tracer may undergo chemical reactions or may not, as is the case with receptor-ligand binding models. For example, glucose is transported from the blood into the tissues via glucose transporters (GLUT1 and GLUT3 in the brain). Glucose can then be converted to glucose-6-phosphate (G6P) by the action of hexokinase. From this step G6P can be further metabolized to provide energy and produce a number of biomolecules necessary for survival. Similarly, glucose-6-phosphate

can be converted back to glucose through the action of glucose-6-phosphatase. Finally, glucose can also leave the cell through glucose transporters. Each of these steps can be represented by a rate constant. 2-deoxy-D-glucose (DG) differs from glucose only at the second carbon atom where the hydroxyl group is replaced with a hydrogen, and it is this difference which makes it so useful. DG is able to shuttle in and out of cells via the GLUT1, 2, and 4 transporters just like glucose. It is also able to compete with glucose for the activity of hexokinase; however, this is where the similarities end. While glucose-6-phosphate is able to be converted to fructose-6-phosphate via the activity of phosphohexoseisomerase, and is eventually released as CO_2 through the Krebs Cycle, the absence of the second hydroxyl group prevents DG from being further metabolized. Similarly, it was shown that DG cannot be oxidized by glucose-6-phosphate dehydrogenase and enter the pentose phosphate pathway. Furthermore, as the deoxyglucose-6-phosphatase activity is minimal, the DG will remain in the cell until it is degraded, effectively recording the activity of hexokinase after the DG was injected. As convention does not consider the blood to be a compartment, the DG model represents a two-tissue compartment model.

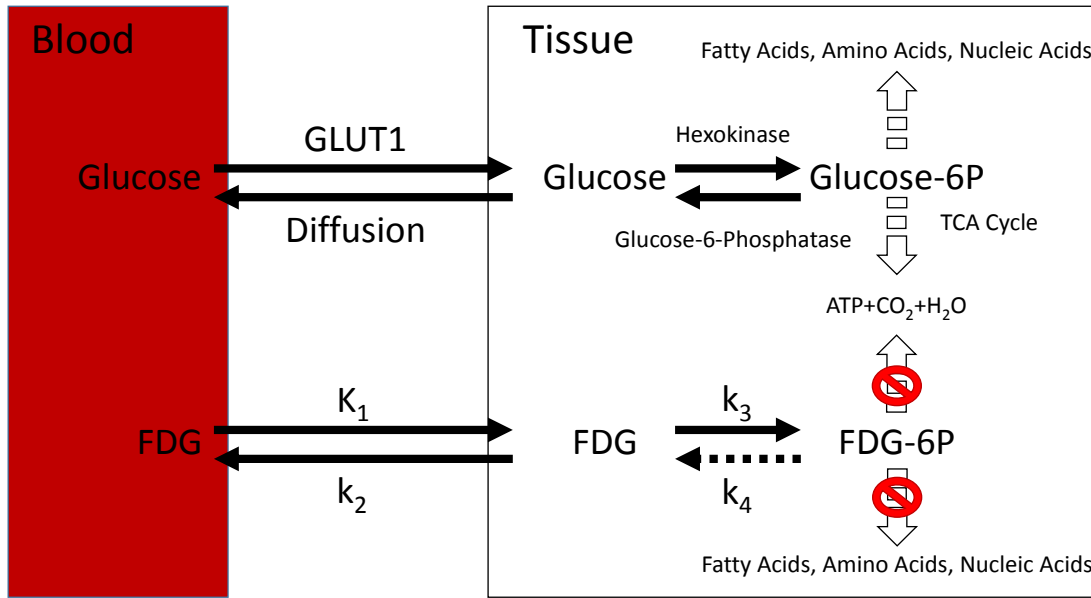


Figure 5 The Two-Tissue Compartment Model. The compartment model allows the derivation of the kinetic parameters that describe a set of sequential biochemical processes. Depicted is the metabolism of glucose and an analogue FDG which allows for the measurement of parameters related to metabolism. The two tissue compartments are free, unmetabolized Glucose/FDG and metabolized Glucose-6P/FDG-6P.

For the models that allow quantitative PET to remain valid, a radiotracer must satisfy

three major assumptions:

1. The amount of injected tracer is small enough that it does not perturb the physiology of the system.
2. The tracer and nonradioactive analogue must continuously be in a steady state, e.g. it has access to all the same compartments as the analogue.
3. There are no isotope effects; labeling the molecule does not change any of its properties.

The PET signal in a voxel in the body is actually the composite signal from the different compartments that exist there (e.g. for FDG a single location may have signal contributions from blood, unphosphorylated FDG and FDG-6P) at a particular time. Application of the model requires one of the compartments be known. The simplest compartment to estimate is the radiotracer input function. In the FDG example, as with most tracers, the input function is that of the radioactivity of the blood, which is typically measured directly via serial blood draws. For tracers that are not trapped in the tissue, and are peripherally metabolized, the signal from the

parent radiotracer must be isolated from the signal of the metabolites (discussed in Section 2.6.6). For some receptor binding tracers it is possible to get a measurement of the available radiotracer from regions where the receptor is not expressed. To simplify the equations, and as a result of the steady-state of the system, the rate constants are considered constants with respect to time.

In practice these rate constants are derived from fitting the model to a dynamic series of PET images that span the imaging time for a region-of-interest or a single voxel. If the variance is high or the SNR is low, the inaccuracy can lead to improper estimates of the kinetic parameters. To address this, a number of simplified graphical methods have been introduced.

1.1.3.3 Dynamic PET Imaging

Following the assumption of a kinetic model with static rate constants, it can be shown that the radioactive distribution at a point in the object can be represented by a function that is continuous in time. The radiotracer distribution as a function of time is referred to as a time-activity curve (TAC). As a result of the stochastic nature of the radioactive decay, Equation (1.7), PET cannot instantly estimate the radiotracer distribution. Thus sampling of the TAC is achieved through serial imaging of a region of distribution over a specified time where the duration of each PET image is dependent on the rate of change of the tracer and the signal-to-noise characteristics in any regions of interest. For neuroimaging, where the region of interest is always in the field of view, dynamic PET acquisition can most easily be performed through continuous acquisition of the PET data where the LOR crystal pairs and corresponding temporal information is retained, a method referred to as list-mode acquisition. List-mode acquisition

differs from LOR acquisition as the latter does not retain any intraframe temporal information. Thus, while list-mode data can be sorted into many different LOR time-series, an LOR time-series cannot be converted back to list-mode or reframed. Once acquired and the duration of each frame is determined, the data can be reconstructed and the activity of a voxel or collection of voxels can be plotted over time yielding a TAC.

1.1.3.4 Radiotracer Input Function Estimation Methods

As previously mentioned, in order to perform quantitative PET imaging the radiotracer input function typically must be estimated. The most straightforward method to acquire measurements of the input function is through direct arterial sampling in which an indwelling catheter is placed in a blood vessel and the catheter line is sampled throughout the duration of the study. While the most accurate site of sampling would be the afferent arterial vessel of the organ of interest, the assumption is made that the radiotracer concentration at all points in the arterial vasculature is the same and the catheter is placed in the radial artery for ease of access.

While direct arterial sampling is the easiest, it is an invasive procedure and requires additional research staff which limit its universal adoption [49]. Furthermore, arterial catheterization presents a disincentive to subjects and complications include temporary occlusion (19.7%), hematoma formation (14.4%), local infection (0.72%), sepsis (0.13%), permanent occlusion (0.09%), and pseudoaneurysm (0.09%) [50]. Alternatively, numerous less-invasive methods have been proposed to measure the arterial input function from the PET data directly [51]. These methods can solely rely on the PET data or can utilize structural imaging modalities to derive the vasculature. PET-only methods can be subject to bias and noise which

can limit their accuracy, especially as it is difficult to differentiate between arteries and veins at the early time points. Additionally, coregistration of sequentially acquired structural and PET images for this purpose has been shown to have limited accuracy due to both the lack of mutual information in the vascular regions and deformations of the vasculature due to the different head positions of each scan [52].

1.1.3.5 Graphical and Simplified Kinetic Model Methods

While full kinetic analysis is ideal, it may not necessarily be feasible. This can be due to obstacles arising from the practicality of arterial lines and serial arterial blood sampling, the imaging duration necessary for kinetic modeling, or the statistics of the data being insufficient for accurate model fitting. Between full kinetic modeling and semi-quantitative imaging there exists a group of methods that rely on simplified models. Typically these simplified models provide less information than full kinetic modeling, often by reducing the number of parameters or by estimating combinations of parameters. These simplified techniques can be broadly broken down into simplified model-based techniques, where compartment models are used however the aspects of the protocol are selected to simplify the modeling, and graphical methods where the data are mapped onto a linear variable and the parameters are estimated. The different methods of analysis and their respective order to regards to some practical requirements can be found in Figure 6.

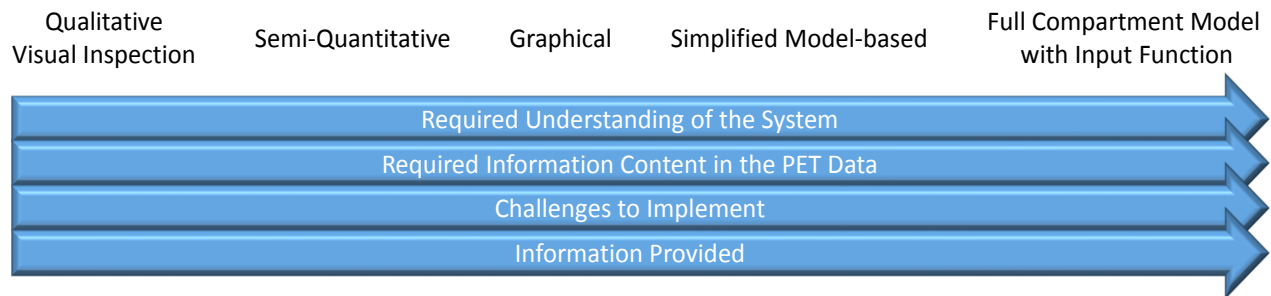


Figure 6 Methods of PET Analysis. Arrows point in the direction of increasing attribute requirement, e.g. full compartment model with input function requires more information be provided than qualitative visual inspection.

Simplified model-based approaches include single-scan and reference region techniques. Single-scan techniques seek to derive kinetic information from a single time-point, typically with the addition of the radiotracer input function. While this may have the most impact in clinical imaging, where time is limited, research scans do not typically face the same time constraints. Reference region methods seek to obviate the need for the radiotracer input function by using the TAC of a predefined region as the reference that other regions of the brain are compared to. Reference region methods have had the most success with ligand-receptor binding studies where a region of completely nonspecific binding can be determined. Additional simplified model-based approaches are those models that seek to induce a true equilibrium in the radiotracer concentration through the administration of a bolus injection followed by a constant infusion of the radiotracer. While the paradigm allows for the simplification of the model, its goal is to allow for stimulus and displacement studies immediately following baseline scans rather than performing the two separate scans.

Graphical methods seek to transform the data to a form that can be represented by a linear function where the slope and the intercept have a physiological interpretation. The benefit of this mapping is three-fold: fitting can be done simply with a linear least-square approach, the quality of the linearity of the data and fit can be determined visually, and as it is

typically applied to later time-points it is less affected by noise that occurs in the initial rapid exchange of the tracer with the tissue spaces. Graphical techniques include Patlak [53-55] and Logan plots [56], which are used for irreversible and reversibly bound tracers, respectively. While an underlying model may be necessary for interpreting the meaning behind the slope and intercept, neither method is model-specific.

A major factor in implementing more sophisticated models is the information content in the data. While there are limitations that cannot easily be improved, such as those set by the inherent Poisson statistics, the formulation and kinetics of the radiotracer, and protocol, the information content of the data can be improved by maximizing the spatial resolution of the data through addressing factors that may degrade the spatial resolution.

1.1.4 Factors That Affect Spatial Resolution

The spatial resolution of a PET system can be determined by combining the individual resolution components,

$$R_{sys} = \sqrt{\sum_i R_i^2} \quad (1.12)$$

where i represents the individual resolution components. Spatial resolution is important as it defines a minimum distance between objects to properly resolve them as well as how large objects must be to get an accurate approximation of their activity concentration. The minimum spatial resolution of a system is defined as its intrinsic spatial resolution.

PET has a rather limited spatial resolution compared to MRI and CT. There are a number of factors that contribute to this low resolution, but they can be separated into intrinsic and extrinsic factors. Intrinsic effects, those which are inherent to the system include:

photodetector resolution, detector-receiver distance sensitivity, and scintillator crystal size including parallax error, with crystal size being the dominant factor in PET imaging. Also included in intrinsic factors are positron range and noncolinearity effects which arise from the physics of emitted positrons and their annihilation with electrons. Extrinsic factors include partial volume effects and stochastic motion effects. Information related to the extrinsic factors is system specific, and will be addressed for the BrainPET in Section 2.1. Particular attention is given to intrinsic factors in the following subsections.

1.1.4.1 Positron Physics-Based Effects

There are two intrinsic effects that limit spatial resolution in PET imaging that arise from the physics of the positron decay and annihilation: positron range and non-colinearity effects. They lead to ambiguity in positron emission positioning, placing a fundamental limit on the spatial resolution of PET systems. The site of the positron emission is different than the electron-positron pair annihilation site due to the kinetic energy the positron acquires during the decay event. The maximum initial kinetic energy is specific for the radionuclide, however the positrons are emitted with a range of energies and only a small fraction are released with this maximum energy. These maximum energies can range from 0.5 to 5 MeV. As the positron passes through matter it interacts with nearby atoms, losing its energy either through radiative or collisional mechanisms along a path until it has lost all of its kinetic energy, interacts with an electron and undergoes an annihilation event, as depicted in Figure 1. This total path is referred to as the positron range; however, of interest to the effect on spatial resolution is the effective positron range which is the perpendicular distance between the decay site and the LOR in

which the annihilation event was recorded. The effective positron range full width at half maximum (FWHM) in water for ^{18}F , ^{11}C and ^{15}O are 0.10, 0.19, and 0.50 mm, respectively [57].

One method to minimize the effect of the positron range is to incorporate it into the reconstruction algorithm at the cost of increasing noise [58]. However, a potentially beneficial result of combining MR and PET is the effect of strong magnetic field on the positron range. Simulations and experimental work has shown that the positron range of high energy emitting radioisotopes like ^{120}I and ^{15}O can be constrained by high magnetic fields, which in turn can provide higher spatial resolution [59-63].

Unlike positron range effects that can conceivably be reduced, spatial resolution losses due to noncollinearity can only be modeled in the reconstruction of the image. For the interaction between the positron and the electron to occur, the pair must have a sufficiently low momentum, however that momentum must not necessarily be zero. As a result, to conserve the momentum of the system, the annihilation photons will not be emitted at exactly 180° relative to one another, but rather in an angular distribution. The angular distribution of the annihilation photons relative to 180° is approximately 0.5 degrees (FWHM) and the blurring effect on the spatial resolution can be approximated by

$$R_{180} = 0.0022 \times D \quad (1.13)$$

where D is the distance between the detectors. For a PET scanner with a D of 40 cm, this results in a resolution effect on the order of 1 mm.

1.1.4.2 Partial Volume Effect

The ultimate goal of PET imaging is to resolve the voxel-wise radiotracer distribution with a uniform accuracy over the entire field of view. This can be challenging for two specific reasons. The first is the voxelization of the image. If a region of high activity is not large compared to the voxel size, or for voxels at the edge of a region, there exists voxels which only partially contain the region of interest. The activity in these voxels is then the combination of the region and of the surrounding region(s). The losses in activity due to edge voxels or objects that are small relative to the spatial resolution are referred to as spill out. If the surrounding regions also contains some radioactivity then the activity of the surrounding regions will contribute to edge voxels of the region of interest which is referred to as spill in. The second reason can be attributed to the inherent spatial resolution of the PET camera and reconstruction algorithm. As the spatial resolution is not uniform over the field of view, PET images are typically reconstructed with a voxel size that is one-third of the sharpest spatial resolution (FWHM) of the scanner. As a result, the image of a uniform spherical source with a diameter the size of a single voxel would be reconstructed as a larger sphere where the activity was highest near the center and decreased along the radius. This effect also contributes to the spill out of the source and to the spill in from surrounding regions. The results of partial volume effects on a non-homogeneous tracer distribution can be observed in Figure 7. As a consequence of these two effects, even if the reconstruction parameters were perfectly modeled and the object of interest was held perfectly still, the PET can only produce a blurred approximation of the true radiotracer distribution. For an accurate estimation of the activity

inside a uniform object (e.g. a ROI), that object must be at least three times the spatial resolution of the system to have voxels that are free from partial volume effects.

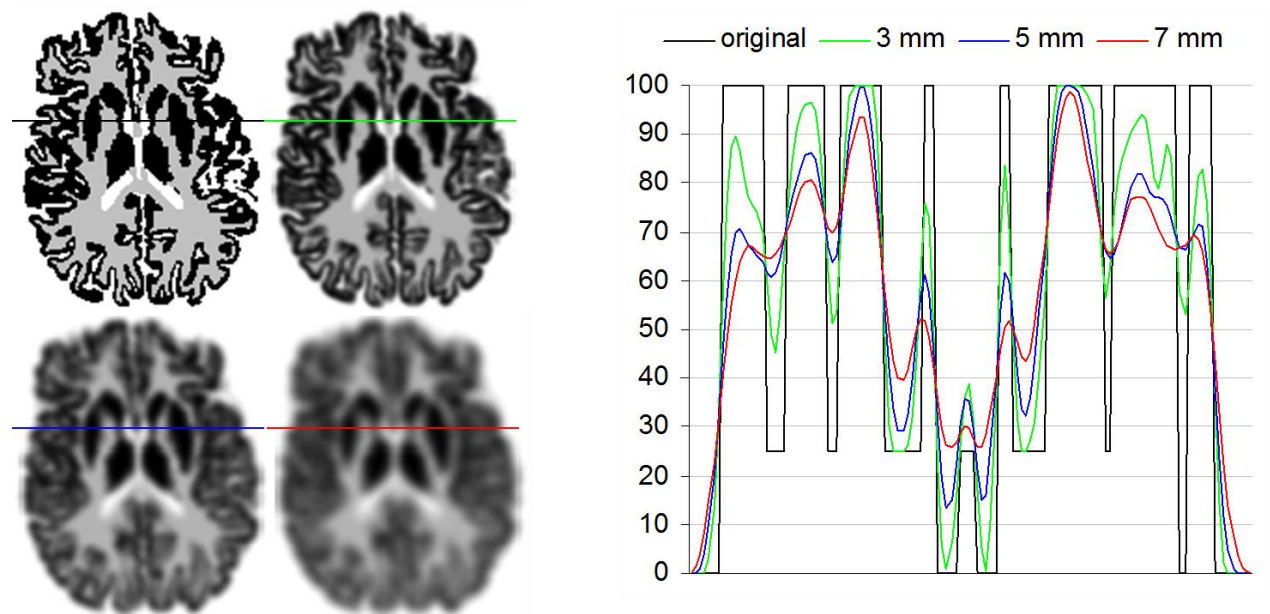


Figure 7 Partial Volume Effect. A simulated brain is shown with regions of low, medium, and high signal with a Gaussian filter to simulate different detector spatial resolutions. The limited spatial resolution of PET cameras introduces a blurring of the data which artificially lower the values in regions of high intensity and raise the intensity of regions of low intensity and is highly dependent on the surrounding intensity values.

A number of methods have been proposed to recover the loss in spatial resolution for neurological imaging. These methods can be separated into those that provide voxel-wise partial volume corrected values and those that provide partial volume corrected ROI measurements. Additionally, methods can be separated into those that are embedded in the reconstruction and those that are performed post-processing. Voxel-wise methods include MAP reconstruction where the anatomical priors provide the boundaries between structures during the iterative reconstruction and are discussed elsewhere [64]. Two very useful voxel-wise post-processing partial volume correction methods include the region-based voxel-wise (RBV) [65] and the modified Müller-Gärtner (mMG) [66] approaches which are voxel-wise extensions of

the ROI-based geometric-transfer matrix (GTM) method and Müller-Gärtner approaches, respectively.

1.1.4.3 Motion Corruption

Motion may be classified as either physiologic, which arises from breathing and the pulsatile flow of blood through the body, or stochastic, which results from the voluntary or involuntary motion of the subject. The latter is more relevant in the case of neurological applications. Most PET based stochastic motion correction methods involve dividing the PET data into smaller frames that can then be registered to a reference frame through a mathematical co-registration algorithm. While these methods can significantly improve the contrast, and by extension the sensitivity, over a “single frame” uncorrected image, the timing resolution is limited to the time necessary to achieve a SNR and contrast that will allow the algorithms to succeed, preventing correction for motion that may have occurred inside an individual frame [67]. These software based methods suffer from two major problems: they rely on noisy data and assume that the distribution of activity does not change with time. More accurate methods, which allow for intrascan motion correction, utilize concurrent hardware systems to track patient motion (e.g. video camera) [68-71]. These methods involve collecting the PET data in list-mode where each detected coincident event is stored and time stamped (as opposed to LOR mode where an entry in a matrix of LORs is incremented each time an event is registered to said LOR). Nevertheless, these external systems require additional hardware and support staff. More importantly, they track the motion of external markers (e.g. attached to the

scalp) that might not always reflect the motion of the brain (or even less so that of internal organs in the case of whole-body imaging).

1.1.5 Application to Neurological Imaging

Since its adoption into clinical imaging, PET has become an important tool in oncology and neurology and is becoming increasingly used in psychiatry. Historically, PET imaging of the brain was most common in oncology where it is still used for the identification, diagnosis of recurrence [72], grading, biopsy/surgical planning, radiation therapy planning [73], response assessment, and prognosis of brain tumors [74]. With regards to neurology, PET is useful in the early and differential diagnosis of dementia, epileptogenic focus localization and subsequent surgical planning in epilepsy, and for the differentiation of Parkinsonian movement disorders. The psychiatric applications for PET are rapidly increasing, partially driven by the growing interest in the development of new tracers, where it is used to better elucidate the mechanism of disease and treatment of psychiatric disorders like depression, schizophrenia, and dementia [75].

While FDG is the workhorse of PET imaging with uses ranging from imaging atherosclerotic inflammation to lymphoma, there are a variety of radiotracers with a diverse sensitivity to different aspect of tissue function. A brief list of such tracers that are used in neuroimaging can be found in Table 1.

Table 1 PET Radiotracers for Neuroimaging

Process	Tracer
Perfusion/Blood Flow	
Diffusible (not trapped)	$H_2^{15}O$, ^{113}Xe
Diffusible (trapped)	$^{13}NH_3$, $^{82}RbCl$, ^{62}Cu -PTSM
Blood volume	^{11}CO
Metabolism	
Oxygen	$^{15}O_2$
Oxidative	^{11}C -acetate
Glucose	^{18}F FDG
Free Fatty Acids	^{11}C -palmitic acid
Osteoblast Activity	^{18}F -
Hypoxia	^{18}F -fluoromisonidazole, ^{62}Cu -ATSM
Cellular Proliferation	^{18}F -fluorothymidine
Protein synthesis	^{11}C -leucine, ^{11}C -methionine
Receptor Systems	
Dopaminergic	^{18}F -fluoro-L-dopa, ^{11}C -raclopride, ^{18}F -Fallypride, ^{18}F -fluoroethylspiperone, ^{11}C -CFT, ^{11}C -NNC112, ^{11}C -Methylphenidate, ^{11}C -RTI, ^{11}C -SCH23390
Benzodiazepine	^{18}F -flumazenil
Opiate	^{11}C -carfentanil, ^{11}C -diprenorphine
Serotonergic	^{11}C -altanserine
Somatostatin	^{68}Ga -octreotide
Estrogen	^{18}F -fluoroestradiol
Amyloid	
Beta	^{11}C -Pittsburgh compound B, ^{18}F -Florbetapir
Tau	^{18}F -T807
Neurofibrillary Tangles	^{18}F FDNP

In addition to the tracers listed and the numerous others that either seek to measure the action, response, or state of some aspect of the body, radioanalogues of drugs can be formulated to study their pharmacokinetics, therapeutic action, and accumulation *in vivo* [76]. Similarly, FDG-PET has gained a roll in understanding the therapeutic mechanism of action of devices like TMS and its effect in schizophrenia [77] and has been combined with other tracers to improve their diagnostic capability [78].

1.2 Magnetic Resonance Imaging

This section provides a basic overview of the principles that govern MRI. A more comprehensive treatment on the physical principles of MRI, its reconstruction, and the development of pulse sequences can be found elsewhere [79-81].

1.2.1 Physical Principles

MRI, or more appropriately nuclear magnetic resonance imaging, is a non-invasive modality that images characteristics associated with the spin angular momentum of nuclei in strong static and varying magnetic fields. While the basic resonance phenomena of MR are exhibited by all atoms with an odd number of protons and/or neutrons, most imaging is optimized for hydrogen as it is the most abundant suitable species in the body.

The angular momentum, \mathbf{L} , of a particle with momentum \mathbf{p} at a radius \mathbf{r} from a reference point is given by

$$\mathbf{L} = \mathbf{r} \times \mathbf{p} \quad (1.14)$$

whose components can be evaluated as

$$\begin{aligned} L_x &= yp_z - zp_y \\ L_y &= zp_x - xp_z \\ L_z &= xp_y - yp_x \end{aligned} \quad (1.15)$$

It can be shown using the Robertson–Schrödinger uncertainty principle that

$$[L_x, L_y] = i\hbar L_z \quad (1.16)$$

And thus

$$\sigma_{L_x} \sigma_{L_y} \geq \frac{\hbar}{2} L_z \quad (1.17)$$

where [...] represents the commutator between the two operators and σ is the uncertainty in the respective components of the angular momentum. Hence for a given z angular momentum there exists a minimum absolute uncertainty between the x and y components of the angular momentum. It can however be shown that the square of the total angular momentum, L^2 , does commute with the components of \mathbf{L}

$$[L^2, \mathbf{L}] = 0 \quad (1.18)$$

allowing for discrete states of the total angular momentum and the z component of the angular momentum to be known. Essentially, if the total angular momentum of the particle and its z component are known, then there always exists some small component in the x-y plane. If the point of reference is taken as the center of mass of the object, then the angular momentum described is the spin angular momentum, \mathbf{S} , and for a system with a total spin magnitude of S , there is always a component that lies in the x-y plane. Protons, neutrons, and electrons—the particles that make up ordinary matter—all have a S_z of $\pm\hbar/2$, meaning the spin angular momentum uncertainty has a magnitude of $\hbar/\sqrt{2}$. As protons and electrons have charge, this spin angular momentum acts like a tiny current loop and induces a magnetic dipole moment μ ,

$$\boldsymbol{\mu} = \gamma \mathbf{S} \quad (1.19)$$

where γ is the gyromagnetic ratio, a constant dependent on the nuclear species. In a magnetic field of strength B_0 , oriented along the z axis, the net effect of Equations (1.16) and (1.18) is that the total spin angular momentum can never lie along the z-direction and the dipole will precess around the z axis with an angular frequency

$$\omega = 2\pi f = \gamma B_0 \quad (1.20)$$

where f is referred to the Larmor frequency. The Larmor frequency for ^1H is $42.58/2\pi$ MHz/T.

In the absence of a magnetic field, the spin +1/2 and -1/2 states are degenerate and there is no preference between states. In a magnetic field oriented along the z axis, however, there is a potential energy difference between the +1/2 and -1/2 states (due to the Zeeman effect, explained elsewhere [82]). This difference leads to a small preference for the +1/2 spin state (about 7 in 10⁶ hydrogen nuclei) leading to a net magnetization of an object, M_0 , given by

$$M_0 = \frac{N\gamma^2\hbar^2 S_z(S_z + 1)B_0}{3kT} \quad (1.21)$$

where N is the number of nuclear spins, and T is the absolute temperature, and k is Boltzmann's constant. At thermal equilibrium the individual x-y components of each dipole cancel each other out and the total magnetization of the object, \mathbf{M} , aligns with the magnetic field. It is this magnetization that provides the signal for MRI. If an additional magnetic field, oriented perpendicular to the B_0 and precessing with an angular frequency given by ω (typically in the RF portion of the spectrum), is turned on for a period of time, it is possible to rotate the magnetization so it lies in the x-y plane. If the additional magnetic field is stopped at this point then the magnetization will precess about the B_0 field leading to a time-varying magnetic field. In the presence of a coil, this time-varying magnetic field will induce a time varying electric field, or free-induction decay, the basic element of the MR signal.

The intensity of the time-varying magnetic field will decay away as the individual spins interact with one another falling out of phase in the x-y plane (T_2 -relaxation), and with the system as a whole (T_1 -relaxation) returning them to either the +1/2 or -1/2 state. The time scales of the two interactions are different by orders of magnitude with the T_2 -relaxation on the order of milliseconds and T_1 -relaxation on the order of seconds and can vary from tissue to tissue. Spatial information can be encoded in the free-induction decay through the use of

additional time-varying magnetic fields known as gradients. The repeated application of RF pulses to excite the spins, gradient fields to encode spatial information, and read-out through a detection coil make up the pulse sequence and the resulting free induction decay data is used to reconstruct the MR image. The nuances of MR imaging, including the use of time-varying linear gradients, pulse sequences organization, relaxation induced contrasts, coil design, and reconstruction are well described elsewhere [80].

MRI is frequently compared to CT in terms of structural imaging modalities. There are a number of benefits to MRI over CT. First, MRI has improved soft tissue contrast. Second, MRI does not involve the use of ionizing radiation, allowing for it to be used for longitudinal monitoring. Similarly, MRI has its shortcomings compared to CT. These include the longer scan time required for MRI, the lower spatial resolution, and the numerous contraindications for MRI. However, in addition to providing structural information, some MRI sequences make use of endogenous contrasts allowing the visualization of a number of functional aspects and highlighting different structural information without the need of externally applied contrasts like CT. The wealth of information available from MR comes from the diversity in application of the time-varying magnetic fields.

1.2.2 Application to Neurological Imaging

MRI has had a profound impact on clinical and research-related neurological imaging. The American College of Radiology suggest that MRI is indicated whenever there is an acute change in mental status, sudden onset of pain, or head trauma [83]. Contrast-enhanced magnetization-prepared gradient echo imaging (MPRAGE) has shown utility in lesion detection

[84]. Additionally, structural imaging has become integral in the assessment and follow up of patients with neurodegenerative diseases like Alzheimer's [85]. In addition to imaging the tissues of the brain, the diverse endogenous contrast provided by MR also allow for the structural imaging of the vessels. Time-of-flight MR Angiography (MRA) has been used for detection of vascular dissection [86], steno-occlusive disease [87, 88], and intracranial aneurysms [89-91].

Aside from structural imaging techniques, functional MR imaging is becoming more routine in clinical management. Blood-Oxygen Level-Dependent (BOLD) MR, which is sensitive to differences in the magnetic properties of arterial and venous blood, has become the tool of choice for imaging neuronal activity and is used in presurgical and language mapping [92]. Magnetic Resonance Spectroscopy (MRS) is capable of measuring the relative abundance of different molecules providing a neurochemical profile of the brain [93]. The cellular and regional tissue architecture of the brain and by extension the health of the white matter can be determined through the diffusion characteristics of water molecules using Diffusion Tensor Imaging (DTI). As many functional MR techniques rely on endogenous contrast, MR has proved itself to be an indispensable tool for characterizing the brain.

1.3 Integrated PET/MR Scanners

Non-invasive imaging has become an indispensable tool in the study, diagnosis, and treatment of human disease. While non-invasive imaging was initially used solely to provide structural information (e.g. x-ray radiography), the quick adoption and popularity of functional imaging modalities like PET and SPECT—which could provide *in vivo* information on such

phenomena as glucose metabolism, gene expression, and receptor-ligand binding—has made non-invasive imaging vital in understanding disease and personalized disease management. While PET possesses an impressive sensitivity, able to detect picomolar concentrations of radiotracers [94, 95], it suffers from poor anatomical detail, making it difficult to correlate areas of activity with their anatomic location. The combination of functional and structural imaging revolutionized PET imaging; only 5 years after the introduction of commercial multimodal PET/CT scanners, they accounted for more than 90% of all PET scanners in use. Similarly, today most manufacturers do not offer standalone PET scanners [96]. Combined PET/CT offered an easy way to correlate functional information with high resolution anatomical information and, due to its side-by-side design, could do so with relatively little interference between modalities. In addition to localizing the activity to a specific anatomical region, CT offered a simple way to determine the necessary PET attenuation correction map; furthermore, CT could do this much faster than the traditional PET transmission scan. However, while CT has demonstrated its efficacy as a powerful tool for providing complementary anatomical information to PET, it suffers from low soft tissue contrast and subjects patients to additional ionizing radiation [97].

1.3.1 History of PET/MR

The idea of combining MR with PET has been around since the late 1980s, about the same time as the investigation of combining PET and CT began. Unlike PET/CT, the integration of PET with MR proved to be much more difficult at the time, and thus received less attention. The major challenges that faced integration of MR and PET could be divided into two categories. The first concern was the technical hurdles necessary to limit interference between

the two modalities. This concern revolved primarily around the need for the PET components to be insensitive to the magnetic fields produced by the MR equipment. In order to have any type of practical multi-modal system, the corresponding technologies would need to be relatively insensitive to one another. In this case, as the MRI produces both strong static and time-varying fields, the PET scanner elements must be primarily built out of non-magnetic components or be properly shielded from the magnetic fields. Furthermore, the PET's operation must be relatively insensitive to the RF fields produced by the MR. With regards to the MR, the PET components must not generate any significant RF noise as it could be picked up by the RF receive coil resulting in MR image degradation. Aside from RF noise, eddy currents are of concern for any device operating inside the magnet. The time varying fields of the gradients can lead to small circular currents in conductive materials. These currents generate magnetic fields that can severely affect the homogeneity of the static magnetic field, leading to poor image quality and artifacts. These eddy currents are also of concern to the PET electronics, as they can lead to resistive heating which can affect the detectors' performance as well as introduce additional noise.

The second challenge to combining MR and PET was designing a system that allowed for simultaneous collection of PET and MR data. Early observations noted the compelling benefits of simultaneous spatiotemporal PET/MR imaging. Unlike PET and CT, which both detect photons through their absorption and conversion to electrical signals, and as a result are likely limited to tandem measurements, MR detects a fundamentally different phenomenon, characteristics of proton magnetic moments. The desire for simultaneous data acquisition in both time and space was noted early in the initial investigation of combining MR and PET, as it

was believed it could potentially benefit both modalities. For example, the functional specificity of PET could be used to strengthen the diagnostic power of MR. On a more basic level, studies of cardiac coregistration methods involving CT and PET or SPECT have demonstrated that coregistration of non-simultaneously acquired data can be problematic and occasionally introduce significant artifacts [98, 99]. Furthermore, the MR data could be used to correct motion artifacts in the PET data. However, in order to simultaneously acquire MR and PET information, the scanners must be concentric, requiring careful consideration in the design of each scanner in terms of space and potential interference. Even though it proved to be too unwieldy a task in the early 1990s, advances in the technology of both modalities have recently led to the development of prototype combined systems.

Creating a scanner with both MR and PET capabilities is a nontrivial engineering and scientific challenge. The limited space inside the MR scanner for PET components, the requirement that all materials be non-ferromagnetic, properly compensating for susceptibility artifacts and eddy currents related to PET components being placed in the magnet, consideration of possible radiofrequency interference between the MR coils and the PET electronics, as well as effects from MR induced heating and vibrations of the PET components must all be considered in constructing an integrated scanner. Most of the material-related artifacts and interference can be minimized by placing the PET components outside of the RF probe [100]. However, this peripheral placement does not resolve all MR artifacts. The PET electronics and related components (e.g. mounting structures, cables, etc.) must remain transparent to the MRI, i.e. extended metal surfaces must be avoided to minimize eddy currents and electrical components must be non-magnetic, or sufficiently shielded, to minimize

their impact on the static field's inhomogeneity. From an MR stand point there is a need to create sufficient space for the PET components and to minimize extraneous RF leakage in their vicinity as it can lead to noise in the PET image.

There have been two avenues investigated to construct combined PET/MR scanners, side-by-side systems and integrated systems. The side-by-side system is simpler from an engineering standpoint; a single table shuttles the patient axially through a MRI followed by a PET scanner (or vice versa) with each device recording information as the patient is passed through the respective scanner. Since these scanners operate independently, their spatial separation, and thus any coupled interference, is only limited by the length of the table track. An obvious drawback to the serial method of collecting PET and MR information is that it lacks temporal simultaneity, and the spatial simultaneity (i.e. the ability to coregister the images) heavily relies on the motion, or lack thereof, of the subject. Additionally, this method requires a longer scan time than a PET/CT system as MR data collection is considerably slower than CT.

Integrated PET/MR systems, while providing the desired spatial and temporal correlation between images, are much more difficult to engineer due to hardware complications. There are a number of approaches that have been investigated for integrated PET/MR systems. Prototype and early scanners relied on photomultiplier tubes (PMTs) that were fiber coupled to crystal arrays that could be placed in the scanner [101-105]. Keeping the PMTs far from the bore was necessary as PMTs operate based on moving electrons through strong electrical fields over a nonnegligible distance which made them particularly sensitive to magnetic fields. Fiber coupling the scintillator with the detector was non-ideal as the lossiness of the fiber—approximately 50 to 75% of the scintillation light was lost—degraded the timing,

spatial, energy, and image resolutions [101, 102, 106, 107]. Moreover, to achieve a spatial resolution and sensitivity that are comparable to clinical scanners requires many scintillator blocks, which in turn would require many fibers whose paths from the crystal to the detector may not be trivial. If these fibers travel along the axis of the magnet they can limit the axial dimension, and thus the axial field of view, of the scanner. An alternative PMT-based design that addressed some of these issues utilized a low-field (1T) split magnet, allowing the fibers to travel radially outward [102]. Eliminating the bends in the fibers improved their light output, and its “double-donut” design allowed for an extended axial field of view. However, to provide a homogeneous region around the MR isocenter this system required special gradients to compensate for its discontinuous nature. In order to extend the axial field of view, the split MR components would exist farther and farther from the isocenter, requiring higher performance gradients to compensate for inhomogeneities. Yet another PMT-based design uses a field-cycled MRI (FCMRI) whereby the static field is replaced by two resistive magnets, a polarizing magnet (on the order of 1T) and a homogeneous readout magnet (on the order of 0.1T) [108]. In this configuration the PMTs can be operated when the field is nearly 0 T. While PET images have been collected using a 0.3 T polarizing field and 94 mT readout field [109], the resistive nature of FCMRI limits its strengths to lower fields and require very stable currents for the readout resistive magnet. Additionally, restricting the operation of the PET camera to when the field is near 0 T restricts the time it can collect data and therefore its sensitivity.

While these first generation systems serve as proofs-of-principle for the combination of PET and MR, and are useful for developing methods for PET/MR, their clinical potential remains limited due to their simple design. Multimodal PET/MR imaging has the potential to be a very

useful tool in both clinical and research applications. While MR has shown promise in the better quantification of PET imaging, PET has the potential to add additional functional information to MR imaging. As the technology is relatively new and still being validated, it is important to design a system that not only integrates both systems, but also utilizes the highest performance capabilities of both technologies, and does so to their fullest, to properly explore the potential impact and ability of multimodal imaging. Not only would such a system provide a new standard in designing PET imagers and allow medicine to answer questions that have remained a mystery, but it would also help facilitate the adoption of multimodal PET/MR imaging.

1.3.2 Applications to Neurological Imaging

PET/MR imaging has the potential to become an indispensable tool for the understanding, diagnosis, and post-treatment monitoring of neurological disease. PET and MR provide complementary anatomical and physiologic information allowing for the improvement of both datasets. PET has been combined with MR for nearly as long as it has been combined with CT. However, while combined PET/CT systems that use a single bed became available in 1998, it was not until 2007 that PET and MR could be performed using a single system. Prior to that studies that sought to combine the high structural resolution and soft tissue contrast of MR with the molecular specificity of PET were performed sequentially and the analysis separately. Furthermore, as the data had to be subsequently coregistered challenges arose for highly specific tracers and short or early frames [110]. Combined systems that allow for the simultaneous measurement of both PET and MR do not require software coregistration of

subject data as any offsets between the scanners isocenters can be measured and calibrated. Without the need for additional coregistration, smaller structures that have previously been challenging to analyze using PET like the inferior colliculus [111], thalamic nuclei [112], and nucleus basalis of Meynert [113] can be studied. Similarly, more conservative margins can be defined in presurgical planning as there is less uncertainty in the relative localization of the data.

PET/MR also has the potential for multiparametric imaging, the collection of quantitative data unique to each modality in an attempt to better characterize a system. An example of multiparametric imaging in the case of a brain tumor can be found in Figure 8. Simultaneous multiparametric imaging improves the efficacy of imaging as a tool for mechanistic imaging, imaging focused on understanding the mechanisms and pathways associated with the development and progression of disease, as the data is collected at the same time. The simultaneous nature of the data collection could be particularly useful for acute conditions like migraine, acute ischemic stroke, and psychiatric disease where physiological, anatomical, and psychological changes can occur on the order of minutes. Alternatively, the scanning duration of the imaging can be extended allowing for more PET and MR data to be collected and yet have the total scan duration still remain shorter than if the imaging had been done sequentially. Additionally, once PET and MR are acquired simultaneously it is possible to validate, calibrate, and quantify MR signals and techniques that have been questioned like the source of the BOLD signal.

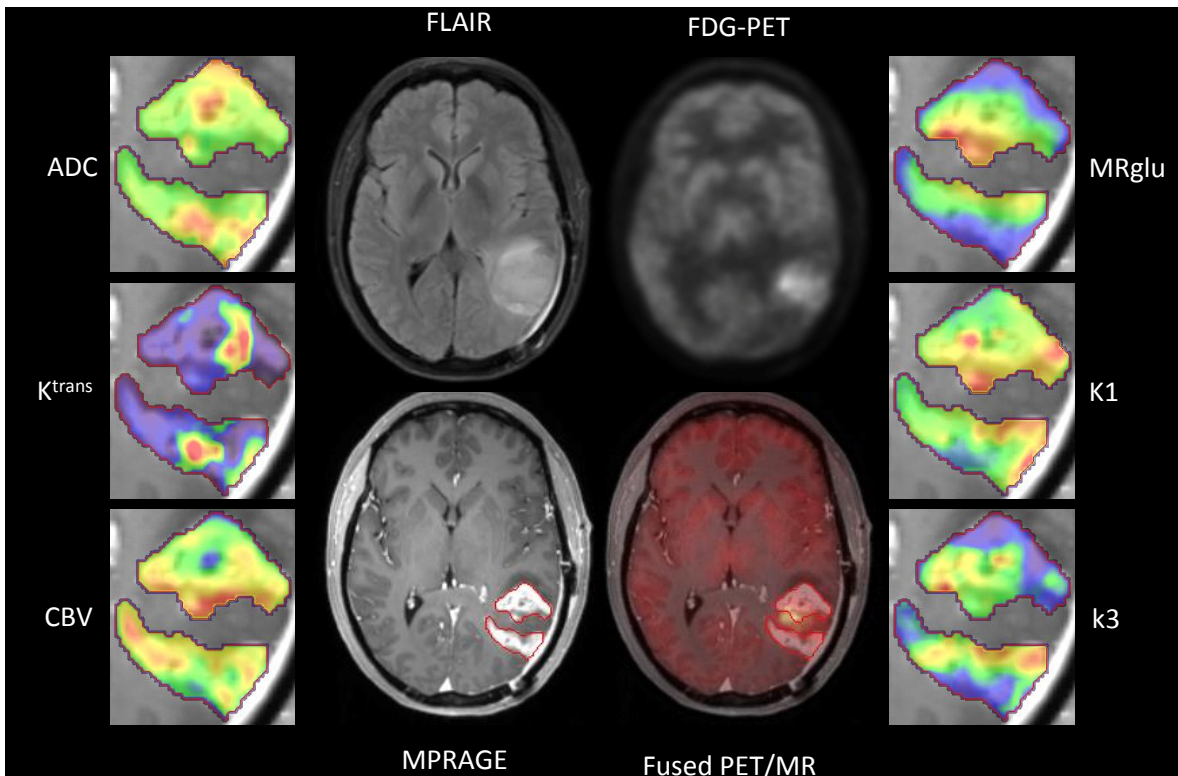


Figure 8 Multiparametric PET/MR Imaging of Glioblastoma. MR and PET are both capable of providing unique information concerning the structure and function of a tissue. In addition to the structural information provided by the MR (FLAIR and MPRAGE), functional information includes the apparent diffusion coefficient (ADC), capillary permeability (K_{trans}), and cerebral blood volume (CBV). Quantitative PET imaging with FDG can provide unique information concerning the metabolism of glucose (MRglu) and kinetic rate parameters (K_1 - k_4).

One potential benefit of simultaneous PET/MR is that it could remove some of the challenges associated with implementing quantitative PET analysis using complete kinetic models. Aside from the simultaneously acquired MR providing information related to the attenuation and anatomical localization, MR can help improve the spatial resolution of PET; assist in region-of-interest delineation and partial volume effect correction; and through the use of MR angiography techniques (Figure 9), provide a simple and robust way to segment the arterial vasculature to estimate the radiotracer input function directly from the data, obviating the need for continuous serial arterial sampling.




Figure 9 Time-of-Flight MR Angiography

Additionally PET/MR provides an opportunity for PET to be used as a tool for longitudinal studies as the additional ionizing radiation from CT can be replaced by MR which does not involve ionizing radiation. As the radiation dose is reduced there is also the potential for application to radiosensitive populations like pediatric patients [114]. In this case, the addition of MR can be useful as it can be used to correct for spurious subject motion over the scan (Section 3.3) obviating the need for sedation and improving the efficacy of PET in other challenging populations like dementia and neurodegenerative diseases.

While combined PET/MR systems present the opportunity to further extend our understanding of disease and improve treatments, there are also some practical benefits for patients. Patients will receive shorter scan times as two imaging studies can be done simultaneously. Acquiring PET and MR data is common for some diseases (e.g. brain tumors) and is becoming increasingly more routine in others. Shorter scan times and a single system capable of two modalities will require less staff and allow for a higher throughput, improving cost efficiency in hospitals. It is also conceivable that the improved diagnostic capability could lead to a reduction in hospitalization and unnecessary tests which will further decrease cost.

Chapter 2 ASSESSMENT AND OPTIMIZATION OF THE PET/MR SCANNER




Figure 10 The Siemens BrainPET. The brain-dedicated PET camera, the BrainPET, operates inside a Siemens TIM Trio 3T MRI allow for simultaneous measurement of MRI and PET signals (left). A rendering of the BrainPET with its external plastic housing turned translucent allowing for visualization of the orientation of the LSO scintillation crystals (right). Image courtesy of Siemens Medical Solutions.

2.1 Siemens BrainPET Prototype

The BrainPET is an MR-compatible brain-dedicated PET scanner prototype, designed to operate inside of a 3T TIM MAGNETOM Trio MR scanner [115]. The BrainPET, which has a

detector design similar to a previously proposed small-animal PET scanner [116], uses avalanche photodiodes (APDs) as the photon detectors in combination with Lutetium Oxyorthosilicate (LSO) crystals. The BrainPET has a transaxial/axial field of view of 32/19.125 cm allowing for full brain coverage in a single bed position. The technical specifications of the BrainPET can be found in Table 2 and an image of the insert can be seen in Figure 10. 3D coincidence event data are collected with a maximum ring difference of 67. APDs, in addition to being insensitive to magnetic fields when compared to the traditional photomultiplier tubes, are small enough that they can be directly coupled to the scintillators inside the magnet, eliminating the need for optical fibers [116]. The preamplifier electronics are also housed in the copper shielding and the electrical output signal is transferred to the filter panel via long cables. The signals are then sent to 48 electronic processing modules (EPM) where the events are passed between EPMs to check for delayed or prompt coincidence between events. The prompt coincidence window is 12 ns while the delayed coincidence window used for randoms estimation is set to 120 ns. When an energy qualified coincidence event is found, the data is written to the hard disk.

Table 2 BrainPET System Specifications

	Characteristic	Value
Scintillator	Crystal Material	LSO
	Crystal size (mm ³)	2.5x2.5x20
	Crystal per detector block	12x12
Photon Detector	Hamamatsu APD	8664-55
	Size (mm ³)	5x5
	APDs per detectors block	3x3
System	No. of detector blocks	192
	No. of crystals	27,648
	Axial field of view (cm)	19.25
	Transaxial field of view (cm)	32
	Gantry inner diameter (cm)	36
Dataset	Gantry outer diameter (cm)	60
	No. of LORs	226,934,784
	No. of sinograms (3D)	1,399
	Sinogram size	256x192
	Image size	256x256x153
Performance	Recon voxel size (mm ³)	1.25x1x25x1.25
	Spatial resolution (FWHM mm)	2.5-4.5
	Energy resolution (%)	15
	Photon Sensitivity (%)	5.61
	Scatter Fraction (%)	43.6

As the insert is placed between the internal body coil and the subject, the body coil cannot be used to transmit the RF pulse for MR imaging. Instead a dual-coil system is used where an outer 1-channel-transmit/1-channel-receive circularly polarized (CP) bird-cage coil is used to excite the image and an inner receive only coil is used to record the raw MR data. At this time, there are two coil designs used for human brain imaging, an 8- and a 31-channel coil [117].

2.1.1 PET Data Format

The BrainPET can store data in two formats: *list-mode* and *line-of-response*. The BrainPET signal processing electronics are built on the Siemens QuickSilver format [118-120]. The tomograph sends each event as a 48-bit data packet that contains an 8-bit header and a 40-bit payload. The 8-bit header encodes a counter that ensures no data were lost (referred to as the Gray Code) and also describes the type of event. Count related events include either

coincident events, which encode the two crystal identifiers where the events were detected, or single counts, events that were detected, but without another event in the coincident time window (encoded as two singles event per payload to conserve space and improve efficiency). In addition to count related events, the system sends time marks at 200 μ s intervals to provide time information to the data file and an assortment of user defined gates which can encode information like when MR pulses are initiated, when a button is pressed, respiratory gates, etc. Storing the data in list-mode format maintains the entire 48-bit packet. If the data are stored in LOR format the coincident and singles events are stored separately as prompt and randoms LOR files, respectively, and all timing and gating information is discarded. When the data are transferred from the tomograph and written to the console each data format is accompanied with a header file that includes when the data collection was initiated. This information can be used to stitch multiple list-mode files together with the proper Gray Code headers and supplying the correct number of time marks between frames. This provides a robust way to reconstruct scans where the subject must leave and return to the scanner without the need of additional time recording; rather, the only requirement is that the acquisition is stopped before the subject is removed from the scanner. The reconstruction of dynamic data is further explained in Section 2.2.3.

While the list-mode format contains all the necessary information for the reconstruction of a PET image, the LOR format cannot differentiate between prompt and random events. As a result the prompts and randoms must be stored separately. LOR files can either be stored in unsigned 16-bit integer format (uint16) or IEEE double floating-point number (64 total bits with 11 exponential bits). The uint16 format is used for prompt LORs as fractional counts are not

possible. The floating-point format is used for storing LOR data related to the system normalization and for data that are unrelated to prompt data (e.g. LOR format of the MR hardware for attenuation correction with motion correction). While the randoms could be stored in LOR format, it is much more efficient to store the data in another format, the delayed coincidence maps (DCmap), as it is directly calculated from the singles rate in each crystal. The DCmap stores the singles rate in each crystal as an unsigned 32-bit integer. Given two crystals with singles rates S_i and S_j , the rate of random coincidences R_{ij} is given by

$$R_{ij} = 2\tau S_i S_j \quad (2.1)$$

where τ is the coincidence window time. The BrainPET is not capable of measuring both single and coincident events at the same time. Instead, the singles rates are derived from the delayed coincidence rates [121].

For reconstruction, the list-mode files are first sorted into LOR space and subsequently further compressed into sinogram space (span=9). The PET data are reconstructed with a standard 3D ordinary Poisson ordered-subset expectation maximization (OP-OSEM) algorithm using both prompt and variance-reduced random coincidence events [121] as well as normalization, scatter [34], and attenuation sinograms. As the BrainPET lacks a transmission source, attenuation sinograms must be estimated, either through the additional collection of CT data or directly from the MR data [23, 122, 123], PET data, or both.

In order for the PET electronics to operate inside the bore of the MRI they must be shielded and isolated from the RF waves produced by the MRI components. This isolation is achieved by surrounding the PET cassettes with 10- μ m copper shielded boards [124]. To accommodate the shielding there is a gap of 6.74 mm between blocks. As there are no crystals

in these gaps, photons that pass through these areas cannot be detected. The net result on the continuous sinograms is an overlying structure of zero counts, which is illustrated in Figure 11. On average 38% [range 37%-49%] of the sinogram bins are empty due to the gap structure. As a result of the missing data, Filtered Back Projection cannot produce accurate images and iterative reconstruction methods are required for image generation [125]. While methods have been presented for interpolating sinograms, the gaps are too large to allow for such interpolation methods to succeed in the case of the BrainPET.




Figure 11 BrainPET Crystal Orientation and Sinograms. A single detector ring of the BrainPET shows the gaps between adjacent detectors while the axial view illustrates the gaps between adjacent rings (left). As a result of the gaps, a simulated sinogram (middle) would have a structure of zero counts bins in the sinogram representing virtual LORs.

While compressing the sinograms to remove the gaps could conserve file space just as it is done for the LOR files, the gaps are maintained because event-by-event motion correction schemes can place events into the virtual LORs (further explained in Section 3.3.2) [123]. An example of a sinogram with events placed inside the gaps due to motion correction can be found in Figure 12. The degree to which the gaps are filled is due to the activity, the extent of the motion, and the length of time spent in each position.




Figure 12 Sinogram Gap Filling Due to Motion Correction

The sinogram data are reconstructed with an isotropic voxel size of 1.25 mm and a volume consisting of 153 transverse slices of 256×256 pixels. The data are stored in a format similar to an interfile. The file format consists of a volume file (*.i) where the value of each voxel is stored as an IEEE double floating-point number. For a standard 256×256×153 image, the total file size is 40,108,032 bytes. Accompanying the volume file is a header file (*.i.hdr) that stores the dimensions of the file and the voxel size in millimeters. As the file format does not contain any orientation information or reference to the coordinate convention, care must be taken in preparing and interpreting the files. The BrainPET reconstruction uses a Radiological convention (LAS) requiring that all data used in the reconstruction conform to this coordinate system. Typically this requires a left-right flip in most images files and metadata as they are typically stored in the Neurological convention (RAS). This left-right flip convention is preserved in all of the output files so third-party software that have a preference for RAS may be used such they produce results that can be directly applied, e.g. the transformation parameter output of coregistration algorithms can be fed into the reconstruction.

As the reconstruction of a sinogram multiplied by a constant is equivalent to the multiplication of the reconstructed image by a constant, all count rate corrections are performed in image space. The exception to this rule is when motion correction is performed

(Section 3.3.2.1) as a single frame is broken down into subframes which are then corrected for motion and recombined.

2.2 PET Reconstruction Optimization

As the goal of this work is to improve the quantification of the PET data in a simultaneous PET/MR system, it is vital that either distortion free PET images can be generated, or that the distortion and effects of the reconstruction are properly characterized. Distortions (e.g. stretching/compressing) in the PET reconstruction may reduce the reliability of ROI measurements due to mismatch between MR-generated ROIs (Section 2.6.2) and the reconstructed PET data. While larger ROIs, like cerebral white matter, may not be significantly affected by distortions, smaller regions, such as arterial ROIs that may be used to generate an image-derived input function (Section 4.3), will be significantly affected. Similarly, and perhaps more importantly, the spatial offset between the MR and PET isocenters and the transformation required to align objects in MR and PET space is determined using an image based alignment method (Section 2.4.1). This section sought to optimize the reconstruction of the PET data by minimizing spatial distortions.

2.2.1 Regional Spread Function Model

To account for the non-uniform distortions caused by the PET photodetectors, the OP-OSEM algorithm [121] was modified to allow for the addition of resolution modeling parameters via image space techniques [126]. The modeling parameters are applied to the iterative image estimates before they are compared with the observed sinogram data, allowing

the nonuniformities in the field of view to be minimized when the scanner is accurately modeled. The model and the parameters are equivalent to the regional point spread function (rPSF) of the scanner, and are important for partial volume effect correction (Section 3.2) as well as for accurate estimation of the radiotracer distribution. As the rPSF is applied to the iterated PET image after its estimation, the resulting PET image is not smoothed by the resolution of the camera and smoothing of the image may be necessary for subsequent image analysis. The rPSF of a scanner is heavily dependent on the aspects of the PET system (crystal face size and length, scanner radius, and signal and timing electronics) in addition to the positron physics-based factors.

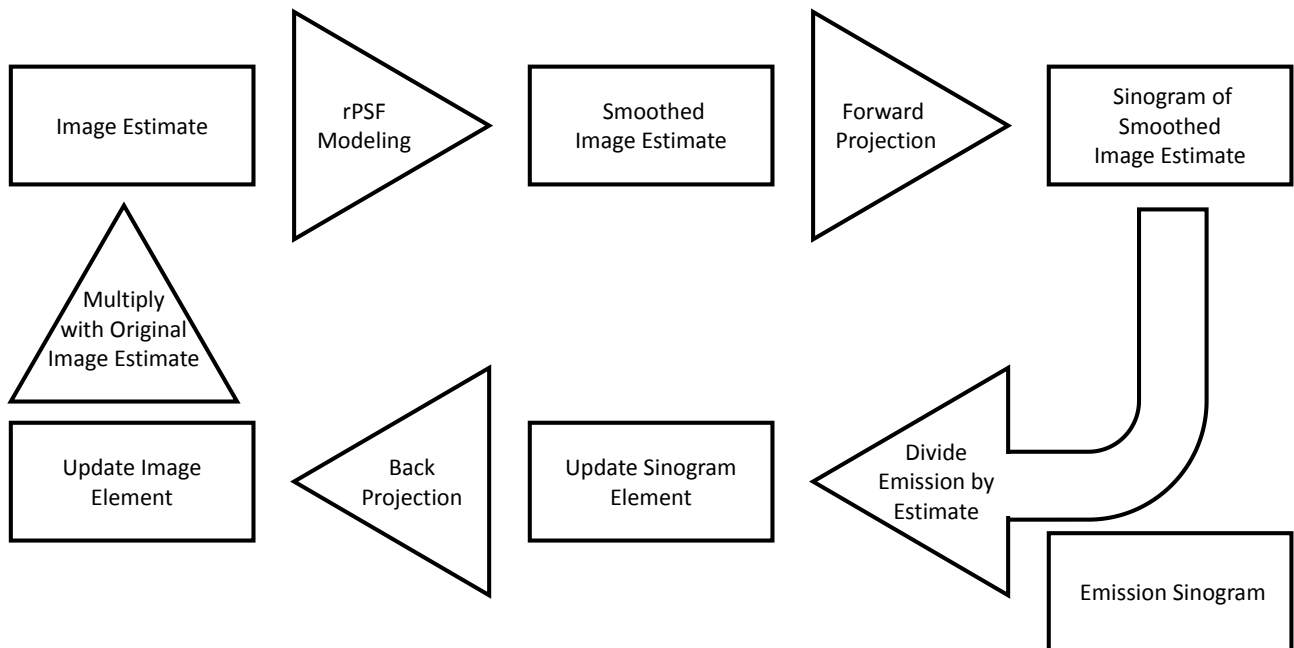


Figure 13 Reconstruction with Resolution Modeling

To estimate the rPSF model for the BrainPET scanner, a separate device that does not have an accompanying MR was used. A Ge-68 point source (as to minimize the effect of attenuation and scatter) was affixed to a robot capable of accurately positioning it in the PET field of view. The point source was moved along a 21x21 grid at one cm intervals with a two

minute PET acquisition at each position. The LORs for the individual acquisitions were then combined and reconstructed using OP-OSEM. The expected and observed images of the 21x21 grid can be found in Figure 14. Based on the acquired images, the rPSF was modeled with two independent functions. The first is a 3D Gaussian kernel which accounts for the symmetric spatial resolution loss due to crystal size, photon noncolinearity, and positron range. The second is a radial motion kernel, used to account for parallax effects [127] and implemented using the formulation

$$I_s(\mathbf{r}) = \frac{1}{l(\mathbf{r})} \int_0^{l(\mathbf{r})} I \left(\mathbf{r} + \frac{l' \times (\mathbf{r}_e - \mathbf{r})}{l(\mathbf{r})} \right) dl' \quad (2.2)$$

$$\mathbf{r}_e = \mathbf{r}_c + (\mathbf{r} - \mathbf{r}_c)b \quad (2.3)$$

$$l(\mathbf{r}) = \|(\mathbf{r} - \mathbf{r}_c)b\|_2 \quad (2.4)$$

where $I_s(\mathbf{r})$ is the smoothed image, $I(\mathbf{r})$ is the original image, \mathbf{r} is the coordinate, $l(\mathbf{r})$ is the integration length, \mathbf{r}_e is the end coordinate, \mathbf{r}_c is the center of the image, b is the radial blurring factor, and $\|\cdot\|_2$ is the L²-norm. With regards to the motion blur, those events at a larger radius are more affected. For the algorithm used, b is parameterized by the variable m where

$$b = \frac{1}{1 - \frac{m}{256}} \quad (2.5)$$

And the symmetric 3D Gaussian was parameterized by its standard deviation.




Figure 14 BrainPET Point Spread Function. (From left to right) The expected image from a 21x21 grid of point sources with 1cm spacing; observed image with spatial distortions; the observed image with a radial blur modeled in the reconstruction; optimized image with both radial blur and 3D Gaussian modeled in the reconstruction.

The values for the rPSF were determined by reconstructing the image and measuring the spatial resolution (FWHM) at each point. Initially, a rotation of 0.6° was observed in the BrainPET data used to model the PSF aspects of the camera. The rotational offset was included in the OSEM to remove the effect of this rotation. An accurate estimation of the rPSF cannot be made without optimization of the radius of the true imaging volume, which as the volume is cylindrical can be parameterized by a transaxial radius. For a two dimensional transaxial sinogram like that shown in Figure 2, the radius represents the maximal spatial displacement of the sinograms; hence for a fixed voxel size and volume, a larger displacement will have the net effect of stretching the image. As the exact spatial position and spacing between the points was known due to the use of the robot, the radius of the scanner could be estimated. The effective scanner radius was estimated to be 18.4 cm. Once the proper radius was determined, the data were then reconstructed with a series of Gaussian and motion blur values. The best spatial resolution was found with a Gaussian blur of 2.5 mm and a radial blur of 16. These results confirm those of Kolb *et al.* who measured the spatial resolution of the scanner at the center of the field of view to be on the order of 3mm FWHM along the three principle dimensions using a line source in air [124].

2.2.2 OP-OSEM Optimization

As the Massachusetts General Hospital's BrainPET currently operates inside of a 3T MRI, the strong magnetic fields prevent the use of the robotic apparatus for measuring certain aspects of this particular PET system including the radius of the scanner. As such, the reconstruction parameters and system model were optimized using a series of phantom measurements.

2.2.2.1 Blur Optimization

The subsequent optimization of the reconstruction parameters relies on the repeated registration of PET and MR data and fitting equations to profiles of the data to determine the parameters in a systematic fashion. In order to accurately fit the data, it must be well approximated by the model function, which requires spatial distortions in the data to be minimized. Section 2.2.1 demonstrated how spatial distortions in the PET data could be removed through modeling the rPSF in the reconstruction, however the exact parameters may vary from system to system.

To determine the optimal Gaussian and radial motion blur parameters, a series of point-source phantoms were used. The point source array was created using 28 glass capillary tubes attached to the end of a wooden ruler with their bases fixed at 1 cm intervals. Prior to affixing the tubes, 5 ml of water containing 10 mCi of F18 was drawn into a syringe. Using a 2" 21G needle, each capillary tube was filled by forming a droplet at the needle then drawing it into the capillary tube. Both ends of the capillary tube were sealed by pressing them into polymer clay. The ruler was placed in the scanner, proximal to the bed and just outside of the field of view to

minimize attenuation effects from the wood while supporting the capillary tubes with the activity near the center of the field of view. PET data were collected in list-mode for 1800 seconds. No MR imaging was performed.

To determine the optimal system model, the data were reconstructed using a range of Gaussian and radial motion blur parameters, 0 to 6 mm and -18 to 18 mm, respectively. As the effective scanner radius mainly “moves” points radially and has little effect on the shape of the reconstructed point, a radius of 18.53 cm was used. Similarly, as attenuation from the glass tube and clay cap and scatter were believed to be small, they were ignored and only images corrected for the normalization of the scanner were generated. The final PET volumes were all rotated about the z axis using a constant such that the points laid along the x-axis. After rotation, events outside the profile were set to zero, activity near the edge of the field of view was set to zero, and the images were smoothed with a 3D Gaussian filter with an isotropic standard deviation of 3mm. To enhance the contrast, the slices were summed along the z-axis, creating a 2D image. The approximate position of the points was determined by projecting the data along the x-axis and finding the peaks. For each of the peaks, the 2D image was then windowed to a 8.75x13.75 mm field of view with the center of the field of view coinciding with the peak along the x-axis, and an asymmetric 2D Gaussian was fit to the region using a least-squares algorithm. From the best-fit 2D Gaussian, the standard deviation of the Gaussian along the x and y-axes (σ_x and σ_y) were determined for each point. To determine the optimal blurring parameters, the mean of σ_x and σ_y as well as their ratio were determined for each set of parameters.

Example point-source images for different sets of Gaussian and motion blur parameters can be found in Figure 15. The number of estimated peaks ranged from 28 to 30, due to the appearance of additional local maxima in the reconstructed point-source images corresponding to the points to the far right of Figure 15 near the edge of the field of view. To eliminate any bias that might be introduced due to these artifactual points, only the central 24 points were considered.




Figure 15 Gaussian and Radial Blur Images. A phantom of point sources was used to determine the optimal Gaussian and motion blur parameters for the MGH BrainPET scanner as the MR prohibited the use of an external robot. The phantom reconstructed without resolution modeling (top) shows a blurring that increases with distance from the center.

The quantitative results from fitting the 2D Gaussians to the images for 3 mm can be found in Figure 16. The motion blur which gave the most symmetric Gaussian, as determined by the average of the ratios of the Gaussian standard deviations, was $m=13$ which had a mean ratio of 1.00 ± 0.05 . However, the standard deviation of the Gaussian along both axes improved with more extreme motion blur values so 16 was chosen since the ratio was large but the FWHM of the 2D Gaussian was smaller.

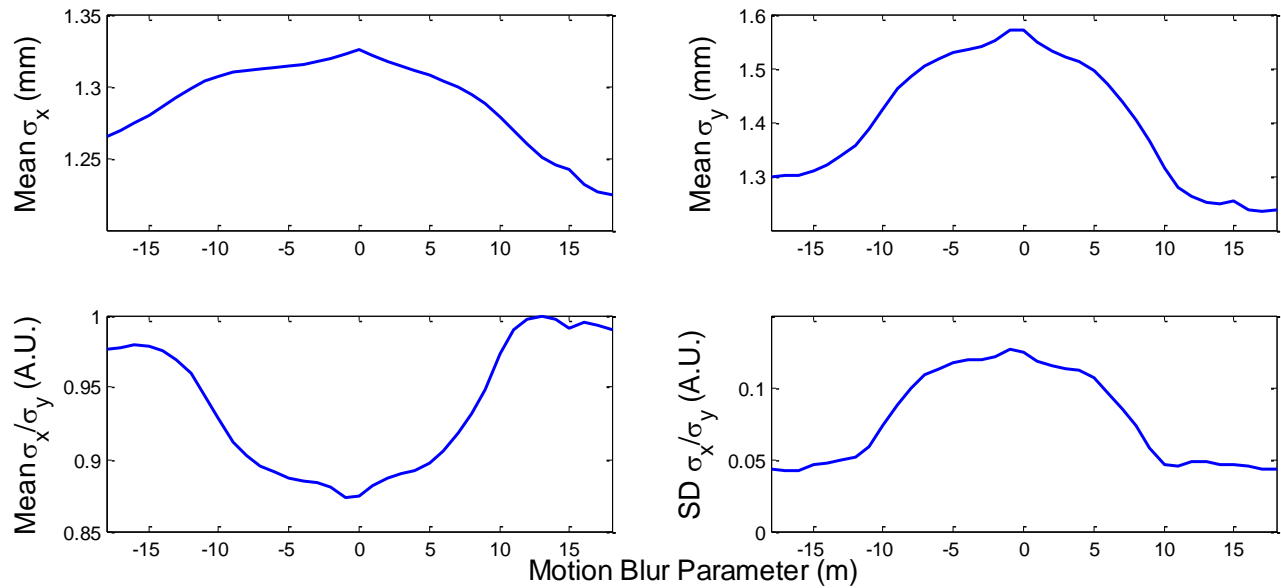


Figure 16 Determination of Optimal Radial Motion Blur. The ideal reconstruction of the point source would be a symmetric Gaussian in the x-y plane. To determine the motion blur parameter which gave the most uniform fit, the data were reconstructed with a series of motion blurs and a 2D Gaussian was fit to each point along the phantom, allowing the average FWHM along each direction and the ratio to be determined where a ratio of 1 represents an isotropic Gaussian.

2.2.2.2 Reconstruction Optimization Methods

To optimize the remaining reconstruction parameters for general imaging, a fillable Dorenzo phantom with holes of diameter 6, 5, 4, 3.5, 3, and 2.5 mm was filled with a mixture of water and ethanol to minimize the formation of bubbles. 2 mCi of F-18 was added to the Dorenzo to provide a PET signal. The phantom was placed in the CP+8ch coil for simultaneous PET/MR imaging. MR images were acquired with a Fast Low Angle Shot **[FLASH]** (0.5 mm isotropic, TR/TE =24/14 ms, TA=24.34 min) sequence to provide the spatial location of the rods. One hour of list-mode emission data were collected simultaneously with the FLASH acquisition. PET images were then reconstructed with different effective PET scanner radii, iterations, and blurring parameters and the optimization of these parameters was determined through image based methods.

As the PET image is influenced by a number of factors including partial volume effects, effective scanner radius, and blurring parameters, determining the centers of all the rods can be challenging. Similarly, fitting a two-dimensional image with a model would require a large number of degrees of freedom and thus would be heavily dependent on the starting conditions. As an alternative, a central slice of the MR image was selected and a 1D profile was chosen that passed through the center of two groups of larger rods (6 and 4 mm). A one dimensional function was then fit to the profile allowing for estimation of the center of each circle. To minimize bias arriving from the fitting procedure, the centers of the rods in both the MR and PET profiles were determined using the same procedure. Assuming the MR image is free of distortions, the optimal effective PET parameters are the parameters where the distance between the centers of the rods in the MR and PET images is minimized. A slice of the FLASH image along with the MR and PET profiles, for the minimum and maximum PET scanner radii that were considered, can be found in Figure 17.

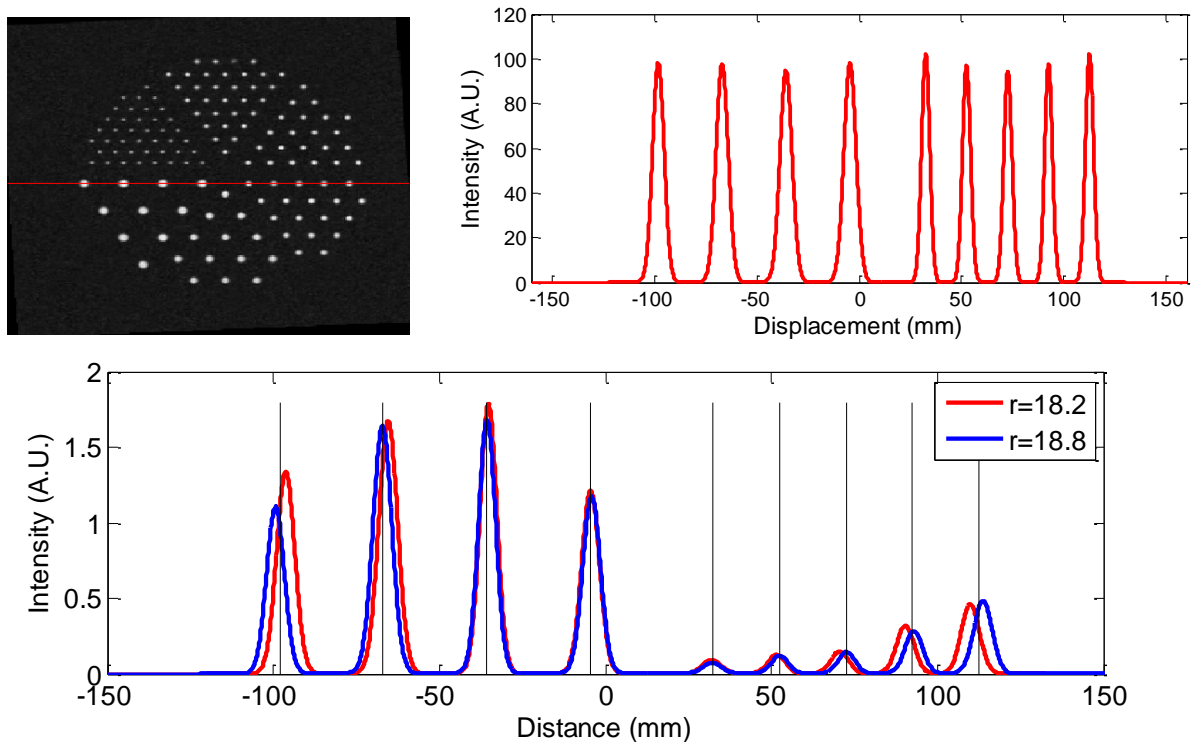


Figure 17 PET Reconstruction Effective Radius Estimation. To determine the proper radius of the MGH's BrainPET system a fillable Derenzo phantom was used. A FLASH image of the Derenzo phantom was used as the true image (top-left). A profile through the centers of a set of holes was measured and fit with 9 Gaussian functions where the peak positions were used to represent the centers (top-right). The PET data were reconstruct (not shown) with varying scanner radii [range 18.2, 18.8] and fit with similar Gaussians (bottom). The vertical lines represent the peak positions as determined by the MR.

In the following subsections the difference between the centers of the circles in the MR and PET images (referred to as the Error) will be used to optimize the PET reconstruction parameters. As an imaging parameter is modulated through a range of values, the individual error for each circle will be considered independently. Each circle is considered independently to minimize the contribution of a single inaccurate fit due to noise and background activity. Graphs will be presented in the following subsections where the error is the ordinate and the parameter of interest is the abscissa. The ordering of the plots and their relation to the Derenzo image can be found in Figure 18.

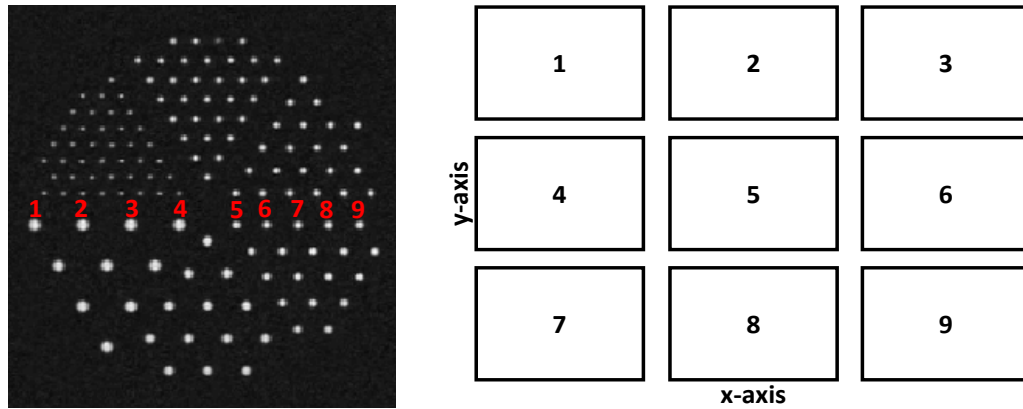


Figure 18 Reconstruction Optimization Graphical Layout. An axial slice of the acquired Derenzo image with the rods along the profile enumerated (left). The layout of the graphs for changes in error as reconstruction parameters are altered where the error is plotted as the dependent variable and the parameter as the independent (right).

2.2.2.3 PET Reconstruction Optimization Data Acquisition

3600 seconds of list-mode data were collected and sorted into a single LOR dataset. The LOR data were then sorted into a set of prompts and randoms sinograms using different effective scanner radii. Normalization sinograms were generated using the same radii from floating-point LOR files. Normalized images, i.e. images that were only corrected for the scanner sensitivity (including the attenuation of the coil), were then generated with specified radial and Gaussian blurs, iterations, and subsets. To correct the PET data for attenuation effects, the normalized PET image was first registered with a FLAIR image which had been previously registered to an attenuation map. The attenuation map was then forward projected into sinogram space and combined with the coil. The resulting attenuation corrected image was used as an input for estimation of the scatter. To account for the effects of missing blocks on the scatter profile a block-wise correction factor must be converted to sinogram space from LOR space using the effective radius of the scanner. The fully corrected image was then

registered to a FLAIR image that had been oriented axially and resliced to an isotropic resolution of 1 mm.

2.2.2.4 PET/MR Profile Modeling

To estimate the centers of the circles in the MR and the PET images, the profiles were fit with a series of nine Gaussians using a least-squares algorithm with the functional form

$$f(x) = a_1 e^{-\frac{(x-b_1)^2}{c_1^2}} + a_2 e^{-\frac{(x-b_2)^2}{c_2^2}} + a_3 e^{-\frac{(x-b_3)^2}{c_3^2}} + \dots + a_9 e^{-\frac{(x-b_9)^2}{c_9^2}} \quad (2.6)$$

where a_n is the amplitude, b_n is the center, and c_n is the variance of the respective Gaussians.

Three separate functions were fit to the PET profiles to estimate the peak position: a set of nine Gaussians, representing the nine hot rods along the profile; a set on nine Gaussians with the addition of a constant term between the peaks of the first and last peak to account for any additional activity of low spatial frequency which may arise from partial volume effects or scattered photons; and the nine Gaussians with the constant term and an additional Gaussian to account for additional observed activity in the profile near the 4 mm hole closest to the edge of the field of view.

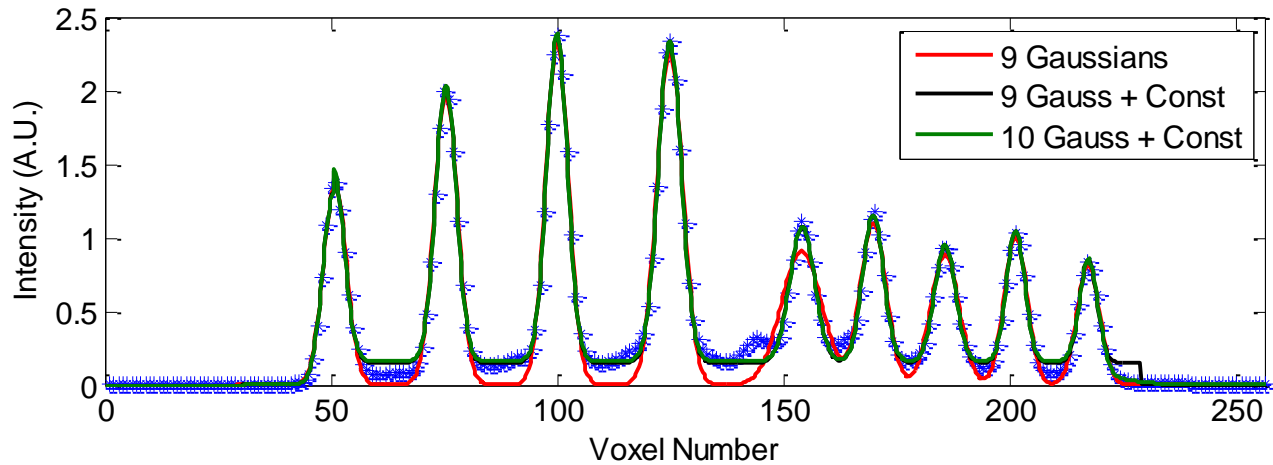


Figure 19 Derenzo Profile Models. Three different models were applied to the PET profile data (blue): nine Gaussians each with a degree of freedom corresponding to peak height, standard deviation, and peak position; the nine Gaussians with an additional top-hat function between the peak position of the first and last Gaussian and a degree of freedom representing its height; and a model of 10 Gaussians with a top-hat function where the additional Gaussian was used to better model the background activity between the 4 mm rods.

2.2.2.5 Radius Optimization

The 9-Gaussian model was fit to the FLASH image of the Derenzo phantom. The parameters of the fit can be found in

Table 3. The signal-to-noise ratio of the MR image was significantly high such that the addition of a background component did not significantly affect the peak position.

Table 3 Parameters of Derenzo Profile

Parameter	Gaussian no.								
	1	2	3	4	5	6	7	8	9
a^\ddagger	99.55	97.47	99.78	97	94.41	99.69	105.2	102.2	93.43
b^\ddagger	-97.9	-66.9	-35.7	-4.4	32.5	52.4	72.6	92.4	112.5
c^\ddagger	4.1	4.3	4.2	4.3	3.0	2.9	2.8	2.8	3.1

For evaluation of the PET reconstruction, a FLASH MR image was taken of the Derenzo phantom and fit with 9 Gaussians. For the MR and the PET data, the position of the peak, b , was used to determine the position for the center of the circle while the variance, c , was used to determine the optimal subset and blur values for the PET. ‡ measured in arbitrary units. ‡ measured in mm.

The PET data were reconstructed with the rPSF modeled with a Gaussian blur of 3 mm and a radial motion blur of $m=16$. The scanner radius was varied from 18.2 to 18.8 cm and the normalized (N), normalized + attenuation corrected (NA), and normalized + attenuation

corrected + scattered corrected images were reconstructed (NAS). The three functional models described in Section 2.2.2.4 were then fit to the profile and the error was calculated. The resulting error in the centers for the rods as calculated using the 10 Gaussian + Constant model are shown in Figure 20.

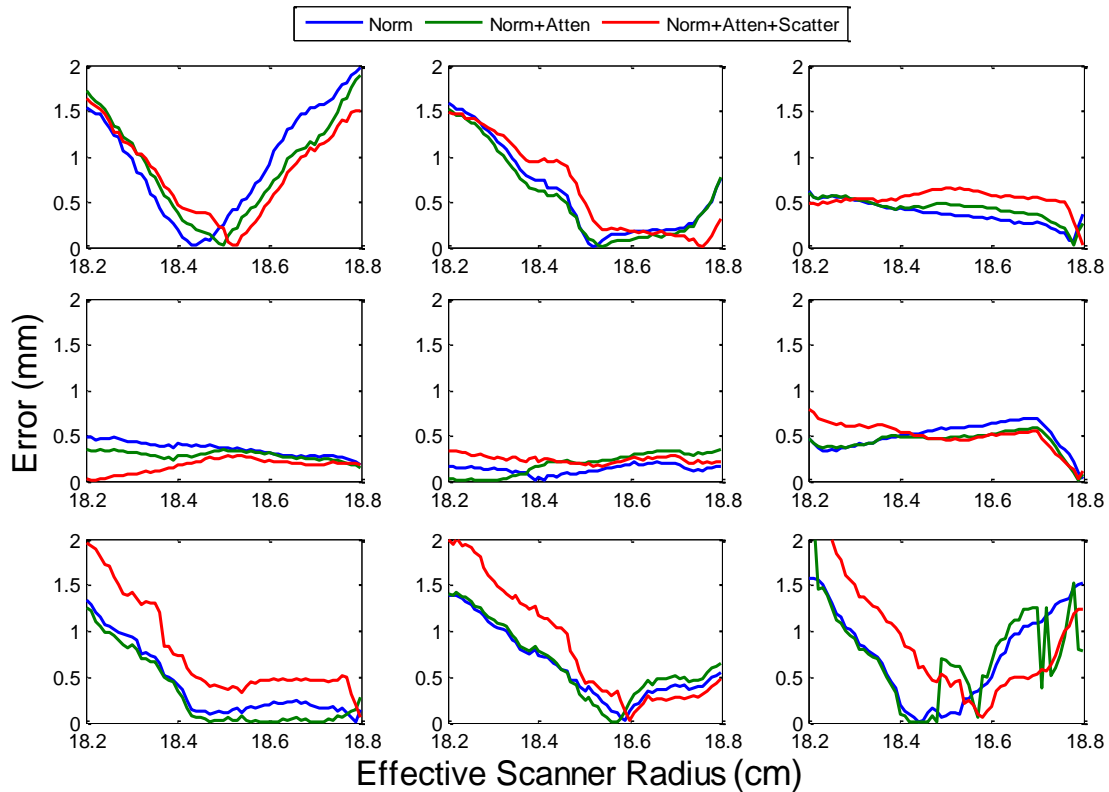


Figure 20 Effective Scanner Radius Estimation. The effective scanner radius was estimated using a Derenzo phantom. Three models of increasing complexity were applied to profiles of the Derenzo phantom to estimate the radius. Depicted is the results for the most complex model where the error is the difference between the peak position as estimated from the FLAIR MR image and the position as estimated from the PET image reconstructed with the given effective scanner radius. A perfect match between the datasets would have an error of 0 mm.

A majority of rods showed an approximate minimum error for a scanner radius of 18.53 cm for the N, NA, and NAS images. A small trend of increasing radius with additional correction was observed between the N, NA, and NAS images. The largest effects are seen in the outer rods, suggesting that the trend may be due to errors in the model for the asymmetric

background in the N and NA images or a reduction in the background and improved contrast due to attenuation and scatter enhancement.

2.2.2.6 Iteration Optimization

As previously mentioned, due to the gap structure in the sinograms, the BrainPET is not capable of accurately reconstructing images using Filtered Backprojection. Rather, the distribution must be estimated using an iterative reconstruction algorithm, OP-OSEM. As the number of iterations increase, the estimate sinogram better reflects the measured sinogram and eventually converges. This is reflected in Figure 21 - after one iteration the estimate poorly reflects the true image, however as the number of iterations increase the structures are better defined.




Figure 21 Effect of Iteration Number of Image Quality. Increasing the number of iterations better resolves structures as is visible in the sum of all axial planes for multiple iterations (top row). However, with increasing iterations the noise also increases, as is apparent from a single plane of phantom passing through a uniform region of the Derenzo phantom (bottom row).

Using the effective scanner radius of 18.53 cm, as determined in Section 2.2.2.5, the appropriate number of iterations was determined. The Derenzo phantom was reconstructed with a Gaussian blur of 3 mm and a radial motion blur of 16. The estimated PET images for the first 128 iterations of the OSEM code were then generated and processed using the steps outlined in Section 2.2.2.3. The peak position of each rod was then determined and the error along the profile was calculated. The optimal number of iterations should resolve all the peaks and minimize the error in the centers positioning as estimated from the MR and PET images.

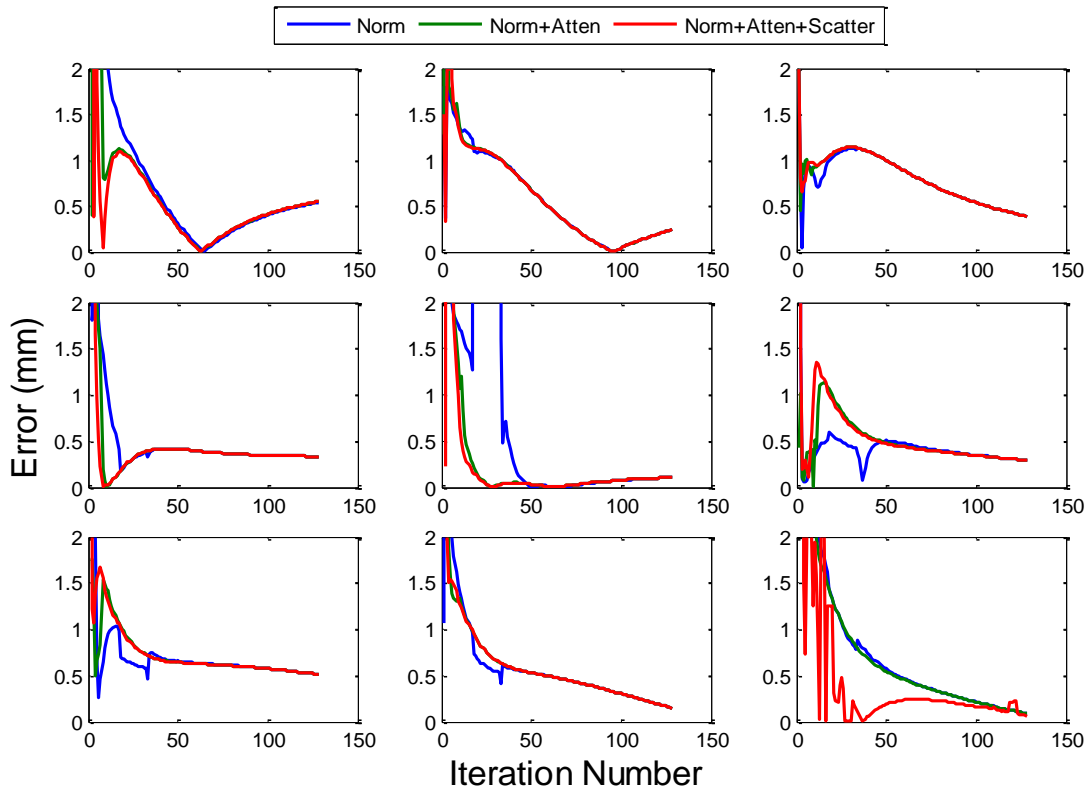


Figure 22 Iteration Number Optimization. Images of the Derenzo phantom were generated using between 1 and 128 iterations of the OSEM code. Data were reconstructed with a Gaussian blur of 3 mm and a radial motion blur of 16.

The results of the reconstruction and subsequent error calculation in the peak position for the nine rods in the profile can be found in Figure 22. After approximately 50 iterations the N, NA, and the NAS images all estimate the same value for the peak position of each rod, except

for the 4 mm rod furthest from the center of the phantom. The furthest and second furthest 6 mm rods showed the smallest error at 64 and 96 mm, respectively. After the number of iterations passed the value that minimized the error for each rod the error began to increase. The errors in the other rods were either asymptotically approaching a constant or decreasing very slowly as compared to the two 6 mm rods. Based on these results, the number of iterations should be between 64 and 96.

OSEM has been shown to have a tradeoff between noise and bias with increasing iterations. This tradeoff is visually apparent in Figure 21, where the small rods are better resolved with increasing number of iterations, but the noise in the uniform region also increases with the number of iterations. The optimal number of iterations is a balance between bias and noise as both voxel-wise and ROI-based measurements will be performed. To determine the optimal number of iterations, a 33.75 mm square ROI was placed in the uniform region of the Derenzo phantom depicted in the bottom row of Figure 21. The bias and the noise were then calculated within that region over the range of iterations. As only the relative bias was relevant, the 128 iteration ROI measurement was taken as the reference or true value. The bias and the noise of the i^{th} iteration were determined as

$$Bias_i = \left(1 - \frac{ROI_i}{ROI_{true}} \right) \times 100 \quad (2.7)$$

$$Noise_i = \left(\frac{SD_i}{ROI_i} \right) \times 100 \quad (2.8)$$

where ROI and SD are the mean and standard deviation within the region of interest.




Figure 23 Percent Bias and Noise Versus Iteration Number

A plot of the bias and noise versus iteration number can be found in Figure 23. As expected the bias asymptotically approached zero with increasing iteration number while the noise increased. Thirty-two iterations gave a bias±noise of $-9.72\% \pm 3.997\%$, 64 iterations gave $3.03\% \pm 4.8\%$, and 96 iterations gave $-1.063\% \pm 5.33\%$.

In addition to the noise increasing with number of iterations, computation time increases. While the reconstruction algorithm has been optimized to run on multiple nodes, it still takes approximately 2-4 minute to generate an image. While this additional time may not be significant in the generation of motion-corrected dynamic PET images (Section 3.3.2.1), where the majority of time is spent combining sinograms in specific ways to generate the motion-corrected prompts, randoms, normalizations, and attenuation sinograms, it is the rate-limiting step in PET-based motion estimation and complementary frame reconstruction (Section 3.3.1.3 and Section 4.3.1.2.1.1, respectively) where many short PET images must be generated. A major step in shortening the reconstruction time was the modification of the EM algorithm to only use portions of the sinogram data for each update. In the ordered-subset expectation

maximization algorithm (OSEM) the backprojection to determine the image update is performed on subsets of projections rather than the entire sinograms, significantly improving speed [128]. As only portions of the data are considered at a time, the backprojection is more efficient and one iteration of OSEM is roughly equivalent to one iteration of EM. Furthermore, for a single iteration of the algorithm, the image has been updated a number of times equal to the number of subsets. To achieve the same point along the convergence as one iteration of OSEM with n subsets, an EM algorithm would have to undergo n iterations.

2.2.3 Dynamic Imaging Accuracy

As the overall goal of this work is to derive kinetic and parametric information from the PET data, it is important that the reconstruction of a time series of PET images is accurate. To this end, it is important that the resulting images reflect the biology without significant artifacts from the reconstruction.

The data from a healthy volunteer who underwent simultaneous PET/MR imaging for another project (Image-derived Input Function; Section 4.3) was used to test the reconstruction of dynamic data. The subject was placed in the head-first supine position and simultaneous PET/MR imaging was initiated. The subject was subsequently injected with approximately 5 mCi of FDG via an indwelling venous catheter using a hand-bolus injection of FDG and imaged for 90 minutes. The data was then reconstructed using constant framing schemes ranging from 600 seconds to 60 seconds. The data were corrected for attenuation and scatter effects, decay, and branching fraction. FreeSurfer was used to generate anatomical ROIs (the specifics of

which are explained in Section 2.6.2.1) and the time activity curves of the regions were determined. A plot of two regions for a series of reconstruction times can be found in Figure 24.

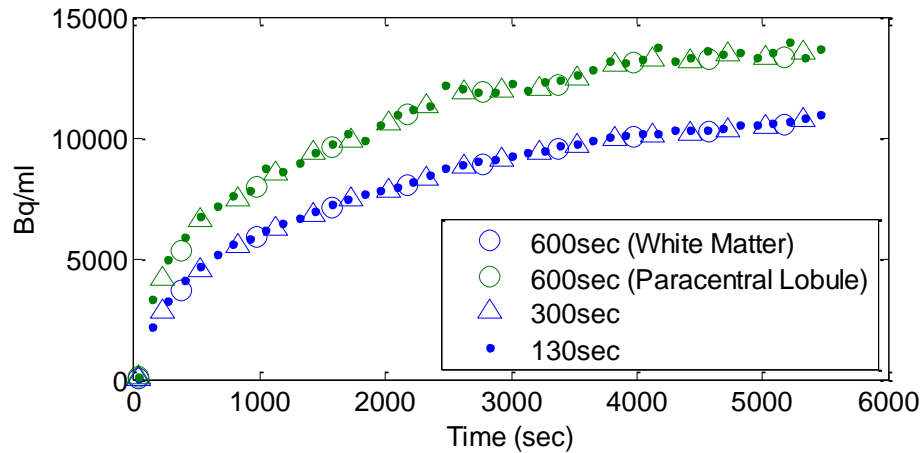


Figure 24 Dynamic Reconstruction Effects. Time activity curves generated from a White Matter mask and a mask of the Paracentral lobule for multiple PET reconstruction framing schemes following a bolus hand-injection of FDG. No significant artifacts or deviations are noted over the range of framing schemes.

No statistically significant deviations were found over frame lengths ranging from 600 sec to 30 sec and all points were within the error bars of nearby points, suggesting that the data can be properly read, reconstructed, and the decay correction is accurately applied.

Data can also be reconstructed using time-varying frame lengths, e.g. 8×60 sec, 2×150 sec, 2×180 sec, and 8×300 sec. To reconstruct the data with time-varying frame lengths the desired frame lengths must be defined in the form of a text file. This information can be saved and reused on other datasets and is part of the automated reconstruction pipeline described in Section 2.5.

2.3 MR Gradient and RF Effects on BrainPET Performance

A number of technical challenges have to be considered in the construction of a simultaneous PET/MR system. While several exotic MR designs have been suggested for

PET/MR systems, the best results, and current industry standard, is to place the PET electronics inside the MRI making it vital for the PET electronics to be insensitive to magnetic fields and the RF fields produced by the MR. As the RF components are placed inside the bore of the PET scanner, the major potential source of distortion and noise are the gradients [100]. Weirich *et al.* previously reported on MR-based interference on the BrainPET [129] requiring a global look-up table based correction factor be applied to the data [130]. The interference manifested as a drop in count rate. Their analysis suggested that the count rate reduction was due in part to eddy currents resulting from compromised shielding, however additional MR-based interference was believed to be a contributing factor. To determine the integrity of the shielding, as well as characterize the effect that the simultaneously acquired MR data would have on the PET electronics and signal detection, interference studies were performed at a detector- and system-wide level.

2.3.1 Assessment of Detector-level MR Interference

2.3.1.1 Materials and Methods

To evaluate potential MR-based interference on the aspects of the individual detectors, including energy discrimination and event localization, the energy spectrum of each crystal in the system was measured without simultaneous MR imaging and while a number of different MR sequences were running. A 0.42 mCi $^{68}\text{Ge}/\text{Ga}$ rod source was suspended in the center of the CP + 8-ch MR coil using surgical tape to minimize the effects of attenuation on the coincident photons. The BrainPET was set to operate in XYE mode (Wide Open Mode) where the energy of every single photon is recorded along with their location on the photodetector

face. Since the singles rate is considerably higher than the coincident event rate, sufficient counts were recorded with 30 seconds frames, allowing for multiple samples to be acquired during a single MR sequence. The MR sequences which were performed included diffusion tensor imaging (DTI), turbo spin-echo (TSE), and BOLD. As the APDs' performance is known to have strong dependence on temperature, the effects of temperature drift were also examined. To reproducibly evaluate the crystal level performance the following imaging XYE-PET/MR paradigm was performed:

1. **No MR imaging:** Baseline PET spectra for subsequent scans
2. **DTI:** Vibration effects
3. **TSE:** Specific Absorption Rate effects
4. **BOLD (A>P acquisition):** Temperature effects
5. **BOLD (R>L acquisition):** Phase-readout effects
6. **No MR imaging:** Temperature effects

The paradigm was repeated on five separate days.

2.3.1.2 Results

Once the XYE data were acquired, the events had to be associated with their respective crystals. Each recorded single event carries three numbers, its position along the x axis, the y axis, and its energy in the form of the associated multichannel analyzer bin. The 9 APDs that make up part of the photodetector provide a position based on Anger Logic with a spatial resolution considerably finer than the crystal size. Additionally, the positioning of the events may be distorted from a Cartesian grid. To properly assign events to their associated crystals, a look-up table is used that is specific to each block and each position in space to a crystal. These look-up tables were determined by collecting a sufficient number of events, creating a 2D

histogram of the data (referred to as a flood histogram), and then performing a watershed segmentation of the histogram.

After the events were organized by crystal, a spectrum was generated for each crystal from all the event energies assigned to it. This spectrum should have two major contributions, one from coincident photons and another from photons that have been scattered. A third component, the background due to the LSO, has a similar range of energies as the scattered photons. While in theory all of the unscattered photons related to coincident events should have an energy of 511 keV and thus appear as a delta function, statistical fluctuations in the amount of light produced by the LSO per keV of incident photon energy leads to a Gaussian-like distribution. The peak (and mean value) of this distribution is referred to as the photopeak. To determine the photopeak of each crystal, two Gaussians were fit to each spectrum, where one Gaussian represented the unscattered photons and one represented the scattered and background photons, with the unscattered photopeak being at a higher energy level than the other. A spectrum of the photopeak positions for all 27,648 crystals that make up the BrainPET was then created.

The resulting system-wide spectrum of crystal photopeaks in the absence of simultaneous MR data acquisition and while numerous MR sequences were acquired can be found in Figure 25. Two pronounced Gaussians were present, the lower energy distribution associated with edge crystals and the higher energy distribution associated with the inner crystals of each block. Fitting the system-wide spectrum with a pair of Gaussians showed no significant deviation from a peak value of 300 when any of the tested MR sequences were run.

Although temperature effects on the order of 1°C were observed over the entire scan, they did not have an influence on the system-wide photopeak spectrum.

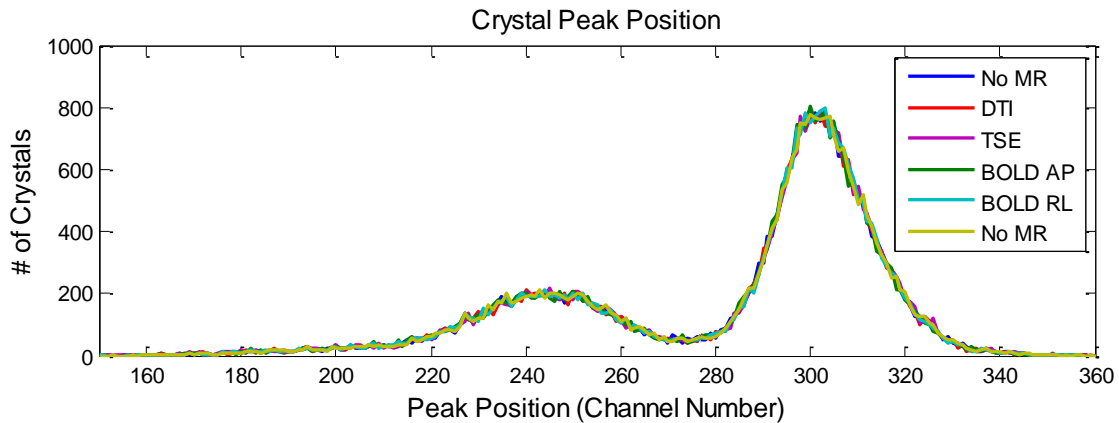


Figure 25 BrainPET Photopeak Spectra During MR Imaging . The PET camera was run in XYE mode and the energy profile and flood histograms for each crystal was measured in the absence of simultaneous MR imaging and during a number of MR sequences. No deviations in the peak position were noted during MR imaging suggesting minimal interference at the detector level.

2.3.2 Assessment of System-level MR Interference

2.3.2.1 Materials and Methods

Phantom experiments were first performed to evaluate the effect of running MR sequences on the PET data acquisition. Sequences included MP-RAGE and dual-echo UTE sequences, two of the sequences that were reported by Weirich *et al.* to have observable effects on the count-rate performance of their BrainPET. The complete scanning paradigm and sequence specific information can be found in Section 4.1.1.1.2. Particular attention was given to temperature effects. The BrainPET manages temperature changes using two mechanisms. The first method is a hardware based temperature controller which activates a closed-loop chiller. When sensors in the cassettes measure temperatures outside the optimal operating range the chiller is activated, providing a heat sink for the forced air that directly cools the

cassettes. The second method is a software-based photopeak tracking algorithm which adjusts the crystals blocks 511 keV photopeak positions based on the cassette's temperature. For these measurements, a uniform cylindrical Ge-68 phantom (~5 MBq, 25 cm length, 20 cm inner diameter) was placed inside the head coil and simultaneous PET/MR imaging was performed using the MR protocol described above. PET data were acquired in list-mode format for one hour and 30-second frames were subsequently generated. From the prompt and delayed events, total true coincident events were determined for each frame. The measurements were repeated five times. To determine the reproducibility, each frame was normalized to the average number of counts per 30 seconds for that trial. Relative change from the count averaged baseline was determined as well as the difference between the minimum and maximum values. To evaluate changes in the reconstructed images, 300 second frames were reconstructed and the maximum relative change from the average was determined over the imaging session. For comparison, the same experiment was also performed with the software photopeak tracking feature disabled.

2.3.2.2 Results

The normalized count rate versus time and the corresponding temperature time-courses for five representative trials with and without photopeak tracking can be found in Figure 26 demonstrating very high reproducibility when photopeak tracking was either enabled or disabled. With photopeak tracking enabled, the maximum change in count rate was less than 1.5%. Reconstructing this data into 300 second frames yielded a maximum drop of $1.08 \pm 0.39\%$ in a large ellipsoidal ROI. This drop in counts occurred during simultaneous DCE

imaging, and shortly after DCE imaging was concluded the count rate returned to baseline. When photopeak tracking was disabled the normalized count rate decreased linearly with time with a maximum change from baseline the order of 25%. Comparatively, DCE imaging didn't seem to have as proportionally strong of an influence on the count rate when photopeak tracking was disabled.




Figure 26 PET Count Rate Fluctuation . PET count rate fluctuation over the course of the protocol (top) with corresponding temperature fluctuation (bottom) with and without temperature correction (left and right, respectively). The highly reproducible drop occurred during simultaneous MR-DCE imaging. The maximum drop in count rate was less than 1.5% over the duration of the protocol. (right) The PET count rate fluctuation without photopeak tracking shows a steady decline in counts over the duration of the protocol. The maximum difference from the average count rate is found to be on the order of 25%, while a change on the order of 45% is observed over the duration of the protocol.

These phantom studies were performed to characterize the hardware-related effects of the MR on the PET data quantification for a scanning paradigm that would be used to examine the effects of MR on cerebral FDG utilization (Section 4.1). While count rate changes were found to correlate with temperature changes, total count rate changes were found to be relatively insignificant. The robust operation of the BrainPET is likely due to its ability to address temperature changes induced in the PET cassettes as a result of eddy currents. With both the photopeak tracking software- and the hardware-based temperature controllers activated, a drop in counts on the order of 1.5% was observed which was determined to have little impact on the PET data quantification; however, when photopeak tracking was disabled

there was a drop in true counts on the order of 45%. This severe change suggests that, in addition to hardware-based temperature control, software-based photopeak tracking is essential for simultaneous imaging using the BrainPET and even for sequential imaging as heat is not immediately dissipated.

2.4 Multi-modal Alignment

While the BrainPET system allows for simultaneous acquisition of MR and PET signals, the data are not correlated by default. This discrepancy takes two forms: a spatial discrepancy due to the fact that both data sets are not being collected by the same detector and thus the center, field of view, and principle axes may be slightly off from one another; and a temporal discrepancy as the data are acquired using two separate computers which do not directly communicate with one another and thus do not share the same clock. The following sections will illustrate how the data are synced, both spatially and temporally.

2.4.1 Spatial Coregistration

As previously stated, the MR and PET data are not aligned by default as shown in Figure 27. This is due to differences in the centers of the fields of view of the respective systems. As the BrainPET is removable, the offset between the MR and PET images can change significantly. Accurate measurements of the offset is necessary as the MR data, and results from it's processing, are used for attenuation correction (Section 3.1), motion correction (Section 3.3.1), ROI-based corrections and analyses (Section 2.6.2).




Figure 27 Spatial Offset of MR and PET. Simultaneously acquired MR (gray-scale) and PET (green) images of a Derenzo phantom show spatial mismatch between the modalities. Coregistration of the data can yield the spatial offset between the systems.

The BrainPET's spatial alignment can be measured using either a fillable, cylindrical, hot-rod lucite Derenzo phantom or a cylindrical, oil phantom with removable ^{68}Ge rods. Regardless of the phantom, the method for measuring the offset remains the same. In the case of the hot-rod phantom, the phantom is filled with between 0.5 and 1 mCi of activity with a half-life greater than the scanning duration (e.g. ^{18}F). The phantom is then placed in the head coil such that the rods run parallel to the bore of the MRI. The setup is then placed inside the BrainPET. FLASH MRI is performed while five minutes of PET data are collected. Next, the phantom is reoriented and MR and PET data are collected again. This process is repeated several times. The registration of the MR and PET volumes is performed in two passes. First, the PET images are reconstructed without accounting for the attenuation or scatter of the object, only the attenuation of the coil is considered as its position in space is considered fixed relative to the center of the PET scanner. The FLASH images are then registered to the PET images using SPM8's rigid normalized mutual information registration algorithm. As the attenuation map of the Derenzo phantom is symmetric, it cannot be directly registered to the individual frames. Rather, using morphologic filtering an attenuation map was derived from a high resolution

FLASH. After the individual MR scans are registered to their respective PET images, the high-resolution FLASH is registered to the individual FLASH images and the transformations are applied to the corresponding attenuation map to generation a position specific map. This attenuation image is used to reconstruct the PET images and correct for attenuation and scatter. The FLASH images are then registered to this corrected PET image. As these are independent measurements of the offset, the components may be considered independent and averaged independently. An example of the values derived from five measurements can be found in Table 4.

Table 4 Mean Multimodal Offsets

	Rotation (degrees)			Translation (mm)		
	X	Y	Z	X	Y	z
1	-0.149	-0.028	0.61	-1.843	0.664	6.655
2	-0.092	0.038	0.562	-1.824	0.642	6.582
3	-0.021	-0.121	0.55	-1.866	0.585	6.574
4	-0.017	-0.016	0.556	-1.835	0.67	6.479
5	-0.067	-0.014	0.581	-1.867	0.64	6.544
	-0.0692	-0.0282	0.5718	-1.847	0.6402	6.5668

As the offset between the MR and PET scanners is derived in image space, any factors that affect the MR or PET images, e.g. effective scanner radius and blurring parameters used in the PET reconstruction, require the scanner offset to be recalculated. Care should also be made that there are no uncorrected distortions or shifts in the MR image.

2.4.2 Temporal Correlation

A beneficial aspect of the independent operation of the MR and PET consoles is that there are no requirements on when or how the data are to be collected, which could for

instance allow for MR flow measurements to be performed during the injection of the tracer followed by the MR sequences necessary for deriving the PET attenuation map.

As the MR and PET systems are independent, they do not share a clock, and thus timing ambiguities are introduced into the data. In order to align the data temporally, a sample-and-hold circuit was designed and constructed using a monostable multivibrator (Texas Instruments SN7412N). Some MR sequences output a TTL pulse every time they repeat an RF pulse sequence (TR). This 5-volt pulse is narrower than the PET electronics can handle (10 μ sec), so the circuit extends the signal. The TTL is then recorded in the PET list-mode data as a trigger.

The system was tested using a Ge-68 line source with an MR-visible phantom and a series of MR sequences (e.g. EPI, CLN) with varying TRs (e.g. 20 ms to 5 s). The number of expected triggers was consistently found to match the expected value given by the acquisition time divided by the TR.

To determine the actual time shift, first the acquisition time of each MR EPI sequence is decoded from the DICOM metadata along with the TR. The sequences are then organized into blocks representing the duration of the EPI sequence. The unique frequencies of the TRs are then stored and used to reject triggers in the PET data that do not correspond to EPI sequences. The frequency of a PET trigger is determined by the time of the successive trigger. If that period corresponds to one of the TR frequencies, that portion of PET data is tagged with that period and added to that block. Successive portions of PET data that share the same frequency are then grouped together to represent a sequence where the end of the sequence is when the frequency changes to a value outside a tolerance range (i.e. a trigger chain of TRs of 3 s with tolerance of 0.5 s would consider a successive trigger that occurred 3.5 s after the last to still be

part of the chain but would consider 3.6 s to be a new chain). If the TR frequency is different than the event before, a new block is started. All blocks are weighted equally. The temporal offset between the datasets is the value that maximizes the cross correlation between the signal chains. An example is shown in Figure 28.

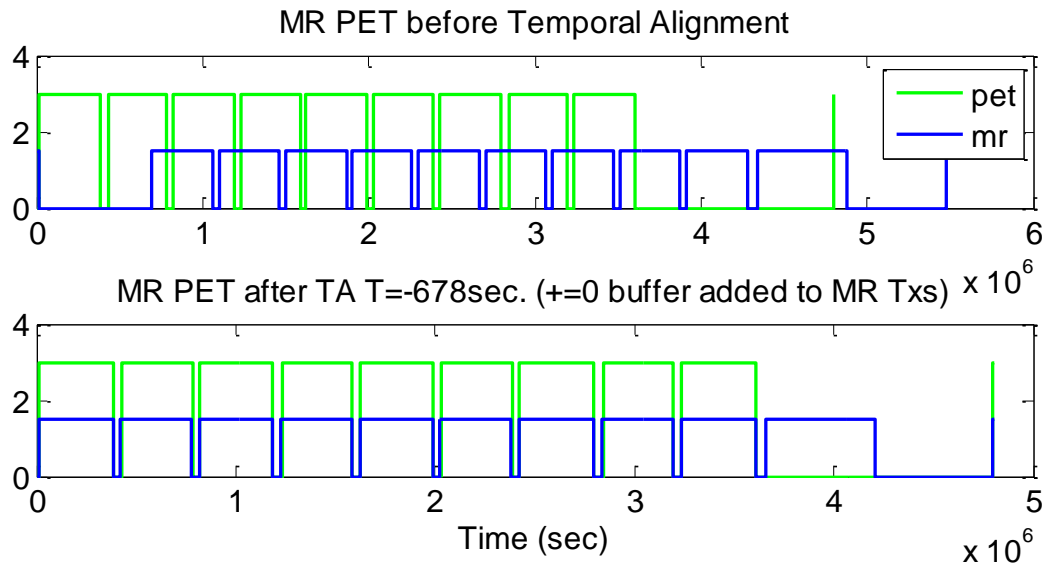


Figure 28 Multi-modal Temporal Alignment. The MR and PET computers do not share a clock which can introduce offsets when trying to correlate the data in post processing. By having the MR produce trigger pulses that can be written into the PET data this offset can be visualized and then corrected for by maximizing the cross-correlation of the “times” when the MR sends the pulses and when the PET camera receives the pulses. Cross-correlation is robust enough that features appearing exclusively in either the MR or the PET data are largely ignored in the estimation.

In the absence of any MR triggers in the list-mode data it is possible to approximate the temporal alignment using any information that is present in both the MR and PET data. If the spatial offset between the MR and PET cameras is known, and the data are collected simultaneously, it is possible to use the position information to temporally align the data. Finding these spatially aligned images can be challenging, especially as the MR and the PET imaging frames are very different. However, extending this idea, it is possible to simplify this search by looking at the change in position between successive frames, i.e. the motion of the subject. How the motion profiles of the subject are estimated from the MR and PET data is

discussed later (Section 3.3.1), however by maximizing the information overlap between the datasets, the time offset can be approximated.

2.5 BrainPET Automated Data Processing and Image Reconstruction

Accurate reconstruction of PET images requires considerable knowledge of the associated physics and signal processing electronics. Prior to reconstruction, the state of the system must be known and deviations from normal system behavior must be detected and compensated for in the image reconstruction process. Several steps are involved in generating a PET volume. For instance, correction factors that account for variations in detector efficiencies and photon attenuation must be applied to obtain a first estimate of the radiotracer concentration, which is then used to estimate the scatter correction sinogram and, finally, generate the quantitatively accurate PET volume. Similarly, if data are stored in list-mode format, the temporal location of the data relative to the total time span of the file must be determined; if multiple list mode files are present where reconstructed images will span across multiple files, their relative times must be derived. To properly handle image generation for PET studies, where 1-150 individual images may need to be generated for a single subject, a substantial amount of automation is necessary.

In addition, the reconstructed PET and MR are not aligned by default already discussed. Similarly, different MR sequences may have different user specified imaging orientations resulting in MR images which are also not registered in space. While software co-registration algorithms (e.g. based on mutual-information) may be used to align different images, they may

lead to inaccurate registration and in some cases have very limited efficacy, especially when the two images have few similar features. A more reliable method is to apply the predetermined transformations to the volumes. Similarly, interpreting and applying the information stored in the image header associated with the metadata can be used to move an MR image back to its orientation in real space. These standard transformations must be applied correctly for both visualization and subsequent analysis and may be subject to change as the BrainPET is able to be removed and repositioned in the MRI. Having a unified tool for reconstruction, derivation of the offsets and associated operation state, and analysis can significantly simplify the generation of high-quality PET images.

2.5.1 Detection of Malfunctioning Detector Blocks

As previously mentioned, the BrainPET consists of 32 cassettes, each containing 6 blocks of 12×12 arrays of LSO crystals, which are each read out using 3×3 arrays of APDs. The relative orientation of the crystals and APDs, along with the estimated singles events in a representative a block can be found in Figure 29. Hardware problems with the APDs can introduce artifacts in the images and can affect the quantitation if they are not accounted for. The BrainPET does not have a hardware-based method of detecting APD failure and the individual APD levels are not written to the list-mode data so APD fidelity must be determined using software. The implemented method for checking the state of the detectors is performed by examining the delayed coincidence maps. The delayed coincidence maps (DCmaps) are an estimation of the singles events detected in each crystal. Using Anger logic, the crystal corresponding to where the photon was stopped is determined by comparing the electrical signal detected in each APD

and a lookup table. The lookup table is acquired using a uniform source and is part of regular maintenance, however hardware issues can affect APD gain between maintenance, which can in turn affect localization.




Figure 29 PET Block Design and SD Regions. From left to right: The APD orientation for a single detector block; the corresponding count rate measured in a block when it is irradiated with a source; the two oblique orientations use to determine that the block is functioning properly, additionally the symmetry along the line that bisects the block horizontally and vertically are used (no depicted).

Symmetry is important for Anger logic, and in an ideal system the behavior of each APD would be similar. If the APDs of the block have similar performance and the block is irradiated uniformly, there should be four lines of symmetry in the detector, a horizontal line that bisects the array, a vertical line, and two diagonals (shown in Figure 29 on the right). An example of a uniform DCmap where all of the detectors are behaving properly can be found in Figure 30. The image is scaled relative to the crystal with the maximum count. The figure is organized such that the row refers to the axial position, with 0 being closest to the torso of the subject when placed in the head-first position, and the column represents the each axially oriented cassette with cassette 0 as the one furthest from the ground. The block is referred to using a column:row notation so 16:0 is the block in the cassette closes to the floor and is located closes to the patient torso.




Figure 30 Normal DCmap

The most common issue with the APDs arises from changes in their gain. If the gain is set too high, it will amplify the signal from low energy photons and mistake them for potential coincident event photons leading to an increased countrate in that block. Alternatively, if the gain is too low, it will distort the flood histogram by biasing all the events in that block away from that crystal, rendering the LORs passing through the block unusable. Similarly, if there are any significant changes in the gain of a photodetector's APD, then it will no longer be accurately characterized by the normalization sinogram. An example where the gain on detector 23:0 is inappropriately high can be found in Figure 31. If a block fails prior to a scan then it can be deactivated and imaging and reconstruction are not significantly influenced. However, if the block malfunctions while a scan is being performed then it must be addressed in the reconstruction or artifacts will appear in the resulting images. If a block is marked as malfunctioning, then it is ignored during the reconstruction, wherein all events that would be assigned to LORs that intersect the crystals of the block are discarded. Similarly the normalization sinograms used for the reconstruction are also modified to remove the LORs associated with the block.




Figure 31 Faulty APD's Effect on the DCmap

Whether or not a block is marked as malfunctioning is determined from the cumulative DCmap of the scan. As middle-edge APDs defects do not have significant effects on the block count rate (shown in Figure 32), they are ignored, reducing the system to four sensors. To determine the fidelity of the corner detectors, the block is separated into two sets of two triangles by dividing the detector along the diagonal and two sets of rectangles (size regions in total). As a first pass, the mean of each block is detected, and any mean that is found to be lower than a percentage of the global mean (set to 15% for the BrainPET) is killed. Then, the relative deviation in the mean and standard deviation of each triangle or rectangle from the entire block's mean and standard deviation are determined, and if one triangle's relative change is found to be larger than a threshold (50%), the entire block is marked. This is performed by recording the block id (written in cassette:detector format to a file named dead_blocks.txt), and zeroing all sinogram bins containing those crystals by supplying it to the sinogram sorting function. Similarly, a new normalization sinogram is generated from the float-format LOR data corresponding to the normalization scan. The number of "live" blocks is used in the determination of dead-time correction.




Figure 32 DCmap of Defect in Middle-Edge of Block

2.5.2 Automated Bolus Arrival Time Determination

It is important in dynamic PET imaging to sample the TAC so that its shape can be accurately determined, i.e. where the change in the tracer distribution and activity is low, longer frames can be employed, and when the change is high, shorter frames are necessary. Typically short frames (as low as 1 second) are employed immediately after the tracer injection to accurately measure the dynamics of the system. As the PET camera is not synchronized with the injection, ambiguity can be introduced. Similarly, picking a consistent reference time for when the injection occurred can be challenging and difficult to reproduce, especially with hand-delivered bolus injections and variable catheter lengths.

There are a number of possible approaches to minimize the error introduced in timing. One approach is to include a number of very short frames that span the time when the tracer was believed to be injected. However, it is difficult to predict a time wherein all injections would have occurred. A more elegant approach is to determine the reference point directly from the data. When the PET data are recorded in list-mode format, it is possible to determine the number of true coincidence events for each second that the camera is active, referred to as the head-curve. It is possible for the operator to determine the reference time as the time when true counts begin to rise as the radiotracer enters the camera's field of view. However,

determining that specific time can be difficult and may vary between users. The true coincidence rate is a combination of events arising from random self-activation of the LSO crystals, coincidence events due to scattered photons, and true, unscattered annihilation events. As the radiotracer is injected, most of the recorded events are either due to self-activation or scattered photons, with the proportion due to scattered photons growing as the radiotracer gets closer to the camera producing an increase in activity before the bolus enters the field of view.

2.5.2.1 Three-part Linear Piecewise Bolus Arrival Time Model

To resolve this problem, the radiotracer activity may be modeled using a piece-wise linear function that approximates the physical system, similar to the approach taken by Singh *et al.* [131]. The bolus can be modeled as a piece-wise function of three linear equations,

$$f(x) = \begin{cases} c & t < \tau \\ m_1 \times (t - \tau) + c & \tau \leq t < \tau + \lambda \\ m_2 \times (t - \lambda - \tau) + m_1 \times \lambda + c & t \geq \tau + \lambda \end{cases} \quad (2.9)$$

where c is the background counts, τ is the bolus arrival time, λ is the time-to-peak, and the m s represent the slopes of the lines. To minimize the error that could be introduced to the fit by the exponential decay, only the first 300 seconds are included. Of all the studies evaluated, no study had more than 300 seconds between initialization of the PET camera and the arrival of the radiotracer bolus in the PET field of view.

The model is fit to the data using a least absolute residuals algorithm. The initial estimates of the fit parameters are determined from the data. The initial estimates [ranges] of the parameters are as follows:

- τ (bolus arrival time): Time corresponding to the peak count-rate [0, 300].
- λ (time-to-peak): 10 seconds [0, 300].
- c (background counts): minimum count-rate measured [0, max count rate].
- m_1 (rising slope): the slope corresponding going from the minimum count rate to the maximum in 10 seconds [negative of same but in one second rather than 10, positive of the same but in one second rather than 10].
- m_2 (post-injection slope): slope of the line connecting the peak count rate to the last count rate measured [same as m_1].

An example of the automated bolus arrival time detection model fit to an FDG bolus

hand-injection can be found in Figure 33.

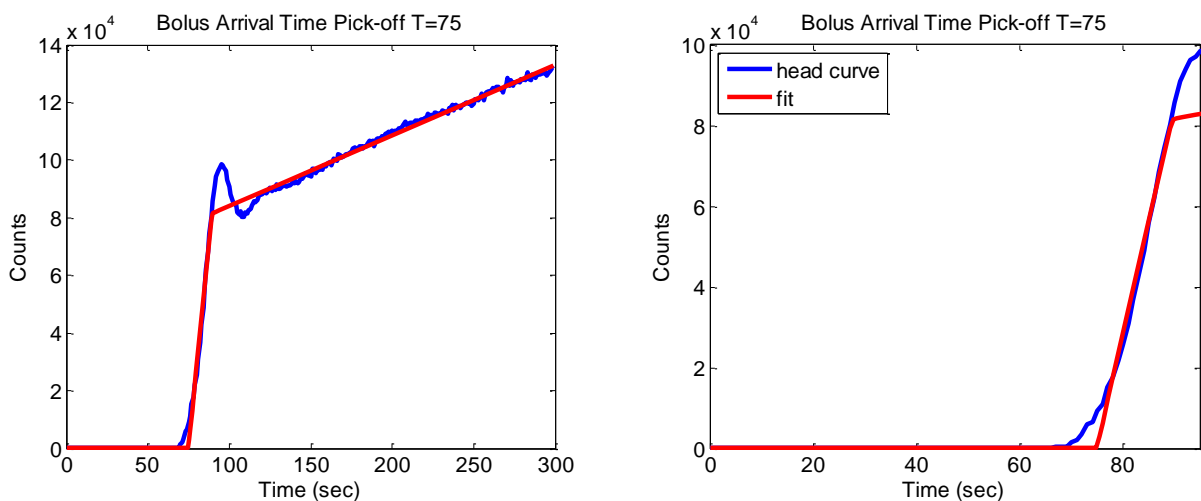


Figure 33 Automated Bolus Arrival Time. Result from a hand-bolus infusion of ~ 5 mCi FDG. Data is shown on a linear scale over the first 300 sec (left) and over the first 100 sec (right). The initial rise in activity is due to scatter events as the bolus moves towards the field of view.

The above model assumes that the rising edge of the bolus can be approximated as a line. While this may be true for a perfect hand injection, if the rate at which the bolus is delivered changes for any reason, this approximation may no longer be valid. An example of a nonuniform rate can be seen in Figure 34. In these cases the BAT is underestimated. This error can propagate into uncertainty in the kinetic parameters, especially those that are heavily influenced by early time points. Additionally, for short half-life tracers like ^{15}O (half life=122.24 seconds), accurate timing is necessary.




Figure 34 Nonuniform injection. Poor bolus delivery or incomplete delivery prior to the saline flush can present as a nonuniform rise in the count rate. This in turn can lead to under estimation of the injection time in the three-part model.

2.5.2.2 Four-part Linear Piecewise Bolus Arrival Time Model

A generalized bolus arrival time model is also considered. The generalized equation includes an additional line segment. This additional segment allows the model to account for changes in the flow rate during the injection. In this case the function is given by

$$f(x) = \begin{cases} c & t < \tau \\ m_1 \times (t - \tau) + c & \tau \leq t < \tau + b \\ m_2 \times (t - \lambda - \tau) + m_1 \times \lambda + c & \tau + \lambda \leq t < \tau + \lambda + \gamma \\ m_3 \times (t - \gamma - \lambda - \tau) + m_2 \times \gamma + m_1 \times \lambda + c & t \geq \tau + \lambda + \gamma \end{cases} \quad (2.10)$$

where the time to peak is now the sum of λ and γ parameters. As the additional degrees of freedom make the fitting more sensitive to the initial estimates, the fitting is done in two steps. First the data is fit with the simpler three-line model. The estimates for the parameters are then used to initialize the fitting of the four-line model. The boundaries of the variables are the same as they are for the 3-part model, while γ is allowed to vary between 0 and 300 and m_3 is allowed to take any value. The 4-part model fit to the data from Figure 34 can be found in Figure 35.




Figure 35 Nonuniform injection fit with 4-part model. This is the same data as presented in Figure 34, however with a 4-part model. Note the improved correspondence with the true data.

Additionally, the 4-part linear piecewise model better approximates the uniform bolus injection. An example of the 4-part model fit to a uniform bolus can be found in Figure 36.




Figure 36 Uniform injection fit with 4-part model.

2.5.2.3 Comparison of Three- and Four-part BAT Models

To compare the effect of using the 4-part BAT model versus the 3-part, a retrospective study was performed. Head-curves from 100 PET studies consisting of ^{11}C -PBR, ^{18}F -FDG, ^{18}F -Fallypride, ^{64}Cu , ^{11}C -NNC112, ^{11}C -Temozolomide, and ^{11}C -Dipernorphrine datasets were processed and the two models were fit to determine the BAT. Both models returned the same

BAT in 62% of cases. The 4-part BAT was found to be later than the 3-point BAT in 34% of cases.




Figure 37 Comparison of 4- and 3-part BAT models. Retrospective study of 100 patient studies where the BAT is compared using the 3-part linear piecewise model and 4-part model. Similar estimates were found for 62 of the studies. The 4-part BAT was found to be greater than the 3-part BAT in 34 cases.

On the reconstruction computer, a Dell Precision T7600 (specifications can be found in 2.6) no significant change in the fit time was measured, with times of 0.94 ± 0.22 s and 1.53 ± 0.23 s for the 3-part and 4-part BAT algorithms, respectively. While the difference in estimation may be small, and will have minimal effects on long-lived, low extraction tracers like FDG, accurate timing has been shown to be important in ^{15}O -water studies [132].

2.5.3 BrainPET Image Reconstruction

The automated reconstruction requires the list-mode data, their corresponding headers, and a study file that contains metadata for the study, referred to as the “Dose info” file. This file contains the activity in the syringe before and after injection, along with the corresponding measurement times so the values can be decay corrected to the same time to determine the

injected dose, as well as the approximate time of injection (Series Time) and the time the PET camera was initiated (Acquisition Time). If a single injection is performed (i.e. bolus, infusion, or bolus + infusion) after the PET scanner has been initiated then the Series and Acquisition times should be the same time, the time the camera was initiated. This time can be found in the corresponding list-mode file header. Also included in this file are subject data necessary for the calculation of SUVs (e.g. weight) and corrected SUVs (e.g. gender, height, blood glucose) as well as additional data (birth year, calibration date/time, blood pressure, pulse). Additionally, the automated reconstruction requires that the radioisotope is noted so the proper half-life and branching coefficient can be used, that the corresponding normalization file be selected, and that a framing scheme is selected (e.g. single 40-90 minute frame, dynamic framing, etc).

The automated reconstruction of the PET data relies on a heavily stereotyped folder organization and naming scheme. The parent folder in which all of the data are maintained and results are output is typically given a study/subject id. For a simple reconstruction two subfolders are necessary, *PET*, which contains the list-mode data and the dose info file, and *MR_PET*, which contains a folder with the attenuation maps (*mu_maps*). Additional folders at the same level of *MR_PET* and *PET* include: *MR*, which contains all the MR dicom data; *CT*, which contains any raw CT data that can be used for attenuation correction; and *Blood*, which contains the blood data.

When the reconstruction is initiated, the list-mode data are first concatenated into a single file with the addition of data corresponding to any missing time marks so a single list-mode file is generated that spans the entire imaging session (data formatting is explained in Section 2.1.1) and is stored in a subfolder of *PET* named *Sum*. The reconstruction is performed

using DOS executables, so batch files that contain the commands are then copied over. The selected framing information is also copied in the form of a text file. After the necessary files have been copied the entire list-mode file is converted into a LOR/DCmap file pair and the count rate information for the duration of the scan, referred to as the head-curve, is determined. The Block Detection algorithm is performed (Section 2.5.1) and the dead block information is stored in a file that is to be used throughout the reconstruction.

The deadtime correction, framing, and bolus-arrival-time are all determined from the head-curve, which has information relating to the prompts, delays, and singles rates. The difference between the prompts and delays yields the true head-curve which is used to determine the BAT following the algorithm outlined in Section 2.5.2. The deadtime and true framing are calculated from the singles rates.

As mentioned in Section 2.2.3, time-varying framing can be used for the reconstruction of PET data allowing prescribed framing files to be saved and used for all subjects in a particular study, eliminating any bias introduced into the kinetic analysis from using different framing schemes and eliminating the potential for user error. With regards to dynamic imaging of data composed of multiple list-mode files with gaps between the files, the correction for missing data is done in the software. Based on the list-mode data headers and from counting the time marks in the list-mode files, the time that the scanner is active can be determined. Based on the list-mode data and the gray code corrected filler files that account for when the scanner was turned off, the elapsed time of the scan can be determined. A vector of the total time, referred to as the series duration, is composed where each position represents one second. If the scanner was active during a given second, the corresponding bin in the series duration is set to

1. If the scanner is off, the bin is set to 0. The framing of the data for a particular study is determined consecutively from the selected prescribed framing. Specifically, first the BAT (Section 2.5.2) is determined and a marker is moved to that position in the series duration. The first desired frame time is then selected. If the scanner was on for the entire length of time then the marker is moved to the new position further down the series time vector and the next desired frame time is selected. If the scanner was not on for the entire desired time, then the marker is moved to the position just prior to the break and a frame is generated from that data up to the point when the system was turned off. A frame is then generated that spans the duration of which the scanner was turned off and the marker is moved to the point before the scanner was turned on. To determine the frame time of the first frame after the break, the amount of time that has elapsed since the BAT is determined. Once that time is determined, the prescribed frames are added one after another until the sum of the prescribed frames is greater than or equal to the elapsed time. The next prescribed frame time in the list is then used and the list is continued from there. This means that if a subject left the scanner and returned at a point in the middle of what would have been a five minute frame had they not left and the next frame would be a 10 minute frame, a 10 minute frame would be started at the moment they returned to the scanner.

2.6 Integrated Tools for Advanced Automated PET Analysis

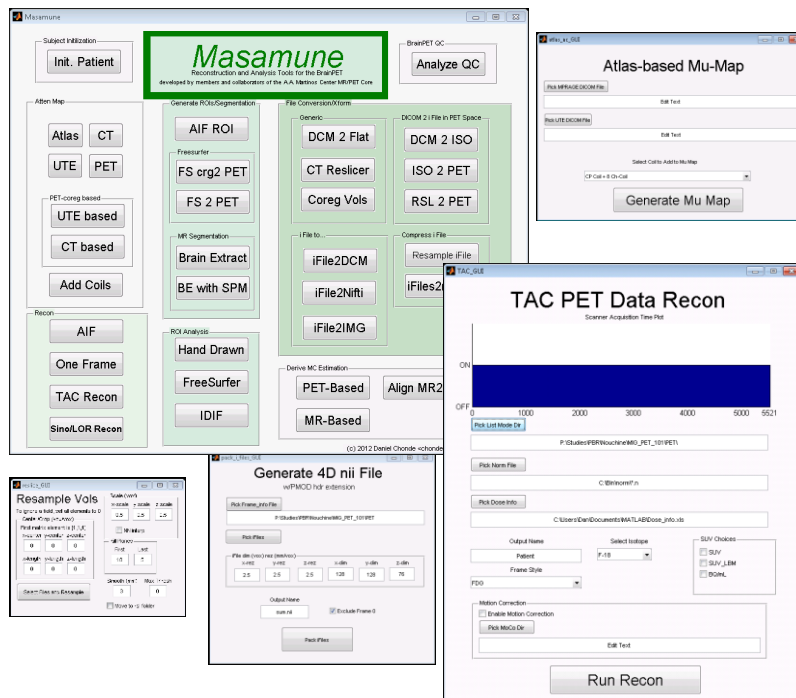


Figure 38 PET/MR Graphical User Interface. The Masamune graphical user interface is built using Matlab and can interface with MR and PET analytic tools in a stereotyped manner for standard users or can run via command line for advanced users.

2.6.1 Graphical User Interface for PET Reconstruction and PET/MR Analysis

The BrainPET GUI, Masamune, was implemented in Matlab along with Windows executables used for the reconstruction steps. The main interface panel, along with some of its associated functional screens can be found in Figure 38. The GUI allows users to: reconstruct both dynamic and static PET images in an automated manner, automatically perform advanced correction to the data (inter- and post-reconstruction motion correction, partial volume effect correction); convert files between different file types (e.g. NifTi, Analyze, DICOM) so that they

can be easily loaded into analysis software (e.g. PMOD, SPM, FreeSurfer); generate properly oriented attenuation maps from MR, PET, and CT images; and segment anatomical MR-images.

The GUI and associated reconstruction tools run on a Dell T7600 with 256 GB of RAM. Memory intensive and complex jobs are automatically sent to a high-performance computing cluster (HPCC) that consists of 127 nodes, each node with 8 CPUs and 56 GB of RAM and then returned to the workstation upon completion. All FreeSurfer based commands are run using the HPCC. Communication between the workstation and the HPCC is performed through Matlab using the ssh2 toolbox [133].

2.6.2 Region of Interest Labeling

A common practice in the analysis of PET data is to group voxels in an image that are believed to behave similarly or belong to the same anatomic structure. By pooling the data into ROIs it is possible to determine the mean activity as well as estimate the standard deviation in that region. In ROIs where the variance is low (e.g. are homogeneous in their uptake), the estimated mean intensity will be a less biased estimation of the true tracer distribution. Similarly, the calculated standard deviation can be used as a set of weights for kinetic analysis and fitting models to the data.

Typically, under the postulate that form follows function, anatomical ROIs are used for PET analysis. Anatomical segmentation of PET data can be difficult as different structures may appear similar; however, segmentation of T1-weighted MRI is common practice. While MR images can be segmented by hand, this task typically requires a trained neuroanatomist, is very time consuming, and may have a high degree of variability. In an effort to address these issues,

a number of automated methods have arisen to automatically label regions of the brain. These automated methods can be divided into three basic algorithms [134]:

1. Expert knowledge-driven methods that utilize a priori information such as shape, location of structures, population-based statistical thresholds, etc. to determine the classification of each point in space.
2. Probabilistic-based atlas segmentation methods that rely on *a posteriori* probability maximization for labeling each point in space.
3. Deformable template based methods which attempt to register a subject to a template image that corresponds to a predetermined atlas.

Masamune uses two different algorithms, FreeSurfer and SPM8, to automatically derive anatomical labels, both of which relying on the T1-weighted MPRAGE MR image.

In addition to using ROIs to define mean activities for regions, they can be used to correct for partial volume effects (Section 3.2) and restrict smoothing to within a tissue class (Section 2.6.4).

2.6.2.1 FreeSurfer-Based




Figure 39 FreeSurfer-derived ROIs. FreeSurfer is capable of automatically segmenting cortical and subcortical structures from an MPRAGE image. Here the labels are shown on a subject-specific 3D representation of the cortex (left) as well as a spherical projection of the brain (middle). Additionally FreeSurfer allows for user-defined ROIs, here shown on an inflated brain (right).

Non-rigid body transformation of the data to an atlas image requires degradation of the spatial resolution of the image which may in turn introduce bias into the analysis. FreeSurfer generates subject specific anatomic segmentation of the brain using high-resolution T1-weighted MPRAGE images [135], which can produce subject specific ROIs for TAC analysis (Figure 39). Similarly FreeSurfer allows for cross subject registration by inflating the subject-specific cortical parcellation and mapping the inflated brain to a normalized space, which provides better inter-subject mapping than Talairach [136]. As the MR and the PET data are collected simultaneously in time and space, the transformations applied to the MPRAGE image can be directly applied to the PET image without the need of external coregistration algorithms, allowing for improved analysis of PET tracers which do not provide significant anatomical detail.

The FreeSurfer segmentation pipeline begins by converting the MPRAGE image into a $256 \times 256 \times 256$ matrix with an isotropic resolution of 1 mm and is realigned to a standard Talairach brain and resliced into a coronal format [137]. All subsequent volumes share the resolution, orientation, and voxel size of this image, which typically differs from the original

image. Subsequent to the registration an intensity normalization and bias-correction is performed through a seed-based approach using the white matter [138]. Subsequently the skull and surrounding soft-tissue structures are removed from the image [139] and the white matter and gray/white matter boundaries are determined [137]. The resulting surfaces are then corrected for topological differences [140], used to determine the pial surface [141], and finally used to segment cortical and subcortical anatomical regions of interest [135, 142-144]. The cortical segmentation is performed using a surface based approach by inflating the brain [145]. The surface-based labels are then propagated along the rays used to convert the cortical ribbon into a surface generating a 3D cortical label that can be combined with the subcortical labels.

For application to the PET data the resulting labels must first be reoriented such that they correspond to the anatomical locations in the original input dataset. The labels are then reoriented such that they correspond to the anatomical locations of the structures in the MPAGE image which is aligned to the PET data. This is performed either through registration of the MPAGE image with the corresponding skull-stripped brain that is aligned to the labels and copying the orientation metadata to the header of the labels file or by applying the transforms used to reorient the original MPAGE image to the PET data if it is available. Once the labels have been moved such that they are spatially aligned to the PET data, they must be resliced and interpolated such that they share the same dimensions. The reslicing and interpolation is performed using SPM and the resulting images are interpolated with a nearest-neighbor interpolation scheme to a $256 \times 256 \times 153$ matrix with an isotropic resolution of 1.25 mm. Once interpolated, voxel indices of the two images correspond to the same locations in

space and the labels file can be used to construct a look-up table relating voxel locations to anatomical regions. The FreeSurfer segmentations are stored in the *MR_PET>ROIs>FreeSurfer* folder.

Care should be taken with the interpretation and application of the results of any automatic segmentation algorithms. This is especially important with FreeSurfer because unlike other software packages (e.g. New Segment tool in SPM) that provide probability of tissue class (bone, air, soft tissue, cerebro-spinal fluid, gray matter, white matter), FreeSurfer segments specific regions. While this can be useful in determining the overall volumetric and morphologic change in a structure over time, if the specific region contains functionally distinct regions then it can introduce bias in ROI measurements if it is assumed to contain only a single functional region [146, 147]. This has been reported for the Thalamus ROI, where the definition extends into the lateral thalamic nuclei and thus may contain a contribution from white matter as well as grey matter [146].

2.6.2.2 SPM-Based

The SPM8-based method relies on an intensity image and a corresponding labeled atlas. The intensity image, or template, must be registered to an intensity image of the subject (typically an MPRAGE). Once the transformation between the template and the subject is known, it can be applied to the accompanying atlas. The Montreal Neurological Institute MNI-ICBM152 template is used in SPM [148-150]. Anatomic atlases corresponding to the Automated Anatomical Labeling [151] and Brodmann areas [152] are also available. Recently, higher resolution MNI templates have been produced with nonlinear registration methods which can

improve registration between the subject and template [153]; however, atlases corresponding to these improved templates have not yet been produced. The linear and nonlinear MNI152 templates can be found in Figure 40.




Figure 40 SPM-derived ROIs. Linear MNI152 template (A) shows less anatomical detail than the “6th generation” nonlinear MNI152 template (B). Subject specific and grey matter limited ROIs derived from the Automated Anatomical Labeling atlas (C) and Brodmann Atlas (D) corresponding to the linear MNI152 template. The subject-specific ROIs are derived through nonrigid registration of the anatomical MRI the stereotactic MNI152 reference.

Generation of the subject-specific ROIs is derived thusly: First the T1-weighted MEMPRAGE MR image is registered to the MNI152 template using SPM8’s New Segment tool which is a unified nonlinear registration, intensity-bias correction, and tissue segmentation algorithm [154]. The transformations corresponding to the nonlinear registration between subject and template is then saved and the inverse transformation is applied to the atlas, mapping it to the subject image. A nearest-neighbor interpolation is used when applying the inverse transformation to preserve the labels. Due to the low-resolution of the template, the labels typically include varying portions of white matter. To remove white matter voxels from

the subject-specific ROIs and to confine the ROIs to grey matter, the subject's segmentation probability maps are used. Any voxel that has a white matter probability of greater than 50% is assigned to white matter and similarly for CSF. This produces a grey-matter-limited set of ROIs that can be combined with the CSF and white matter ROIs. The subject-specific AAL and Brodmann ROIs for a representative subject can be found in Figure 40.

To determine the impact of restricting MNI derived cortical labels to subject-specific grey matter regions for ROI analysis, the AAL and Brodmann ROIs with and without gray-matter restriction were determined for ten healthy controls. The size of each ROI was determined before and after gray matter correction and the relative size was calculated. A bar plot of the mean gray matter contribution for each of the segmented Brodmann Areas can be found in Figure 41. The mean gray matter fraction across all ROIs for the four subjects was calculated to be $53.51\% \pm 10.98\%$ [21.05%-74.49%] for the Brodmann ROIs and $55.70\% \pm 15.03\%$ [4.78%-86.34%] for the AAL ROIs.




Figure 41 Mean Gray Matter Contribution to Brodmann Atlas

The minimal grey matter overlap in the AAL atlas was observed in the posterior cingulum ($32.81\% \pm 4.42\%$), pallidum ($6.77\% \pm 1.57\%$), cerebellum X ($22.17\% \pm 4.94\%$), and vermis X ($11.97\% \pm 4.28\%$). A low gray matter fraction is expected in the cingulum as it is a white matter structure. Similarly, it has been shown that the SPM registration has a low reliability with

regards to nonrigid registration to the cerebellum [155-157]. These low grey matter fractions of the ROIs corresponding to grey matter structures demonstrate that there is significant white-matter contribution (nearly 50%) to the uncorrected Brodmann and AAL ROI sets and care should be taken in interpreting the ROI estimates.

In addition to the Brodmann and AAL ROI sets, other whole brain ROI sets are available, including the Eickhoff-Zilles [158], Talariach Daemon [159], Harvard-Oxford [142], CC200 and CC400 [160], as well as region specific atlases like the probabilistic Choi2012 atlas of the caudate and putamen, a subsegmentation of the Thalamus, and a probabilistic atlas of the Cerebellum. The SPM8-based segmentations are stored in *MR_PET>ROIs>SPM*.

2.6.3 Intensity Normalization

Another application of the anatomical ROIs is masking and global normalization of the data. While semi-quantitative units like SUVs attempt to normalize subjects across body types and injected dose, they do not account for changes in blood flow, baseline metabolism, and/or binding. Intra- and intersubject variability in the case of longitudinal studies can affect the results of group analysis. Intensity normalization is commonly performed to account for these effects, whereby the images are scaled by a constant value specific to each image that give enforce the mean within a ROI (e.g. the whole brain) to be identical across images. An example of two subjects before and after intensity normalization is shown in Figure 42. To perform the normalization the average FDG uptake was calculated over the entire brain as defined by a binary mask and each voxel in the image was divided by the calculated value.




Figure 42 ROI-based Intensity Normalization. The 40 to 60 minute FDG-PET images of two subjects before (left) and after (right) intensity normalization. Subjects were normalized to a whole brain mask. Subjects are normalized to MNI space allowing identical planes to be shown in both subjects.

The choice of region varies by tracer and convention, however typical regions that are used to normalize FDG data include global gray matter, the cerebellum, and whole brain [161].

With normalization to whole brain there is a debate over using a mask that includes the ventricles, or limiting the estimation to solid tissue regions of the brain. Including the CSF regions contained in the brain would lower the average as they are fluid filled and typically cold.

To determine the effect of including CSF-filled structures would have on the normalization scaling factor, a retrospective analysis was performed on 58 subjects who underwent FDG imaging following transcranial magnetic stimulation. The PET images were reconstructed from the data acquired 40-60 minutes post injection and an MPRAGE image obtained simultaneously was aligned to the PET image. The FreeSurfer derived labels were also aligned to the MPRAGE, resliced to the spatial resolution of the PET image and interpolated using nearest-neighbor interpolation. Once the labels have the same dimensions and spatial resolution as the PET image, the mean of a region was calculated as the mean of the PET voxels that corresponded to the label tag in the label image. The mean was determined for three regions, the whole brain including regions of CSF (any label greater than zero), the whole brain

without the CSF structures, and the cerebellum (regions corresponding to cerebellar white and gray matter). The results can be found in Figure 43.




Figure 43 ROI-driven FDG-Intensity Normalization Across Subjects. Whole Brain estimates are plotted as circles to demonstrate the similarity between Whole Brain estimates and the Whole Brain estimates that included internal CSF structures like the ventricles. Cerebellum and injected dose are shown for reference.

There was such significant overlap between the whole brain and whole brain + CSF estimates that it was necessary to represent the whole brain estimates with a marker rather than a line. The whole brain value was consistently greater than the whole brain + CSF value with a relative difference of $-0.55\% \pm 0.41\%$. Based on these results, there is minimal bias in taking the entire FreeSurfer derived brain as a volume to use for intensity normalization of static FDG-PET data. It is important when selecting a region used for intensity normalization that the region be stable across groups, unaffected by the process being studied, and not strongly affected by external stimuli [162].

2.6.4 Surface-based Smoothing

A common application of simultaneous PET/MR imaging of the brain is exploratory analysis over the whole brain [163-166]. Exploratory analysis may either be performed on a voxel- or ROI-basis with the goal of defining a statistical or parametric map. While each method

is sensitive to different effects, they are both improved by a reduction in the local variance within the activation region. Typically this reduction in variance is achieved by smoothing the data with a Gaussian function. Whether performed on the entire volume as a 3D volumetric smoothing or slice by slice as a 2D planar smoothing function, conventional volumetric smoothing mixes tissue classes (e.g. CSF, white matter, grey matter, bone) that may have different radiotracer distributions which can lead to bias in the images, especially at the interface of tissues with different radiotracer distributions (e.g. bone and gray matter). Additionally, due to the folded structure of the brain, cortical regions that have disparate functional and anatomic origins may be located near one another and thus lead to a mixing of signals due to their proximity.

An alternative approach to volumetric smoothing is to separate tissues into their respective classes and smooth across classes. While this removes the effects of mixing between the classes, spatially localized regions on the cortical ribbon may still contribute to one another. To further reduce the error introduced by smoothing along the cortical ribbon, surface based smoothing techniques have been developed where values are mapped onto vertices of the inflated polygon mesh-model of the cortical surface and neighboring vertices are averaged together [167, 168]. A number of studies have performed surface-based smoothing on PET data using FreeSurfer [169-173]. While these methods have demonstrated a reduction in bias and variance in the kinetic parameters of the data, they have not been used for the more routine application to semi-quantitative imaging or applied to subcortical structures.

The combined surface/volumetric-based smoothing begins with the independent reconstruction of the PET data and FreeSurfer driven segmentation of the anatomical MPAGE

MRI data. The previous instances of surface-based smoothing have utilized sequentially collected MR and PET data requiring coregistration. While a registration scheme such as the boundary-based registration (BBR) [174] has shown accuracy with the registration of MR and PET data, it requires the cortical surface be visible in the PET dataset, limiting its application to tracers where the boundary can be delineated and to frames of a sufficiently long duration. In a simultaneous PET/MR system where both data are acquired simultaneously there is no need for software-based coregistration. Furthermore, the effects of mismatch from dynamic imaging where there may be motion either before or after the MPRAGE is acquired can be reduced with the application of motion correction schemes in the PET reconstruction (Section 3.3).

In the case of PET motion correction, the transformation between the original orientation of the MPRAGE data (as read from the DICOM) and the motion corrected position is saved. If the MPRAGE was only corrected for the offset between the MR and PET isocenters, as would be valid for an animal with a headpost or if the subject was fixed in the scanner by some alternative mechanical means, then the transformation would be recalled from the most recent offset measurement. The last possibility would be an unknown orientation in which case the original MPRAGE is registered to the MPRAGE registered with the PET using software coregistration and the transformation is saved. The reconstructed PET data is then converted to a Nifti format and the inverse of the transformation is applied, and the left-right flip is applied, moving the PET into registration with the original MPRAGE.

The combined surface/volume smoothing implementation requires that the FreeSurfer parcellation was performed on MR data that had the same spatial and slice orientation as the original DICOM. This is necessary because, as previously mentioned, an early step of FreeSurfer

is to reorient the data. If the original orientations were used in the parcellation then the transformations, which are saved in FreeSurfer, were used and the PET data can be mapped to the FreeSurfer volumes. Once in alignment with the FreeSurfer volumes the transformations which convert the cortical ribbon volume to the surface are applied to the data with the vertex volume defined as the mean active through the particular ray orthogonal to the surfaces of the ribbon sampled in step sizes of 1% of the total ray length.

In parallel with the surface-based smoothing, a volumetric based smoothing is performed on the on the noncortical ribbon data.

2.6.5 Image Reconstruction and Analysis Workflow

The workflow for the automated reconstruction and analysis is shown in Figure 44. The reconstruction of PET data can be organized thusly:

1. **Organize Data:** As the MR and PET consoles are separate, the data are stored on different machines. The data must be moved to a location where both are accessible to the reconstruction code. The use of a stereotyped folder structure and file format allows for the reconstruction to remain predominantly automated. Additionally, other relevant information, e.g. FreeSurfer derived anatomical labels, can be organized for their subsequent application for PET data correction.
2. **Estimate the Subject Motion:** Patient motion is unavoidable for dynamic imaging and can still be significant for static images depending on their duration. Estimating the motion allows for subsequent correction of the PET data and fills in the sinograms, both of which improve the effective contrast in the image.
3. **Prepare Necessary files for PET Data Processing:** PET data processing requires information regarding the attenuating material in the field of view, which includes the subject and any MR hardware. Similarly the state of the detector and information related to the quantification of the scanner must be determined for accurate PET quantification.
4. **Select Reconstruction Parameters:** Information regarding the framing of the data, tracer type, and study design are necessary for the accurate reconstruction and correction of the data. Additionally the format of the data (e.g. SUVs, Bq/ml, etc.) must be selected in accordance with the analysis goals.

5. **Reconstruct Data.**
6. **Post-process PET Data:** Post processing includes correction for partial volume effects, smoothing the data (surface/volumetric or volumetric), and the additional reconstruction of PET data allowing for the derivation of the PET radiotracer input function.
7. **Export of PET Data:** Many PET analytic software packages have strict requirements on the format of the data (e.g. image format, list delimiters, etc.). The data, and the relevant metadata, must be properly formatted to the specific needs of the software in an automated fashion to the specifics of the software to minimize the chance of user error.
8. **Analyze PET Data.**

Initially the MR data are sorted into folders corresponding to the sequence and individual scan order. Motion estimation algorithms are then applied to the MR data (explained in Section 3.3.1) to estimate the head displacements over the duration of the scan relative to the position at the bolus arrival time and to align the MPRAGE MR image to that reference frame. When MR-based motion estimates are unavailable, the estimates are derived from the PET data. The motion-corrected MPRAGE (and associated MR data) may then be used to generate mu maps for the PET reconstruction. The Time-of-Flight MRA is then aligned to the motion-corrected MPRAGE image and an arterial mask is derived for the derivation of the PET radiotracer input function using an image based technique (IDIF) and the complementary frame-based PET reconstruction method (Section 4.3). Motion-corrected dynamic PET images can then be reconstructed and combined surface and volumetric smoothing can be performed on the data. The smoothed PET images can then be used with the IDIF to generate parametric images.

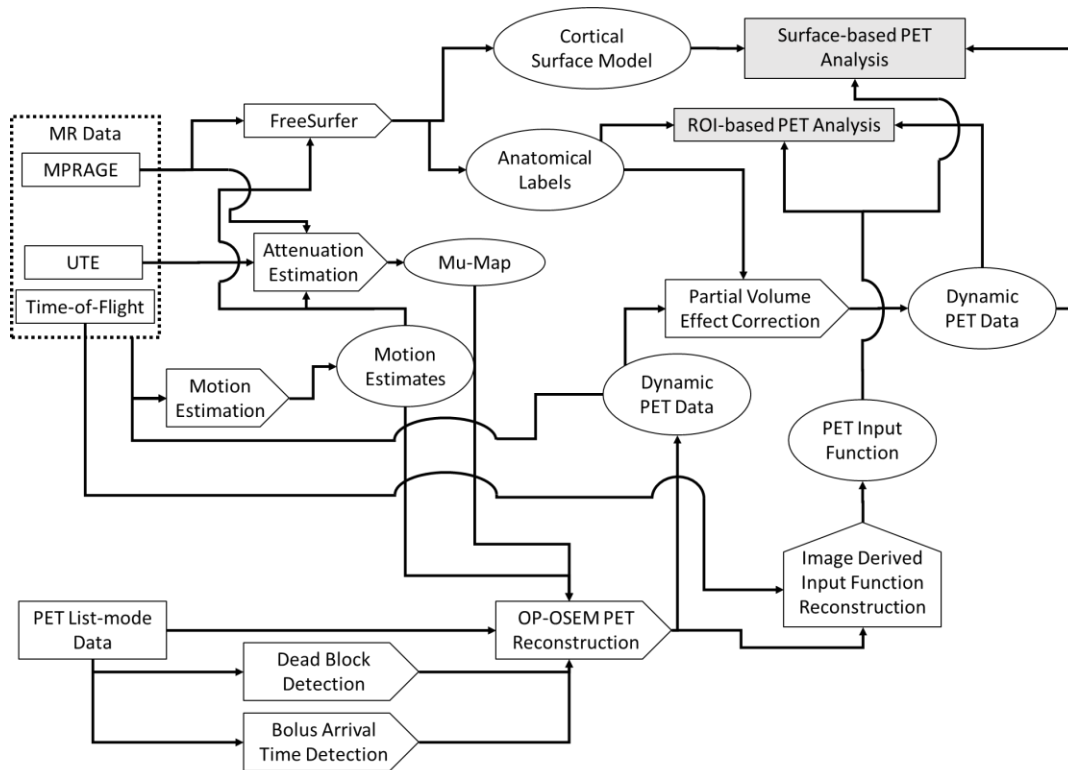


Figure 44 Automated Reconstruction and Analysis Workflow

2.6.6 Blood-based PET Radiotracer Arterial Input Function (AIF) Estimation

Situations arise where the acquisition of the PET radiotracer input function is necessary: studies where kinetic parameters are of interest, requiring full compartmental kinetic modeling; studies where the injection cannot be performed inside the scanner because an additional activity must be performed, e.g. transcranial magnetic stimulation; studies where the blood must be further processed, e.g. metabolite analysis and correction; and studies directed at the validation of a new radiotracer. In these situations the radiotracer input function is typically determined from serial measurements from the blood. Typically the blood is drawn by hand via an indwelling arterial or venous catheter and the activity is determined using an automatic gamma counter (PerkinElmer). The analysis of the gamma counter results to produce a

radiotracer input function requires a number of steps and corrections (e.g. detector calibration factor, decay correction, background subtraction) making it prone to user error. To standardize and optimize the processing of the blood data an automated pipeline was constructed.

The gamma counter uses cassettes that have a capacity of 10 sample tubes and writes the output to a text file. To assist with the analysis, the user must provide necessary information for the analysis including the type (e.g. arterial whole-blood, venous plasma), collection time, and volume of each sample into a template file. The first step of the processing is to convert the measured sample activities to Bq/ml that have been decay corrected to the time-of-injection. To simplify the time keeping and to minimize errors and ambiguity in the temporal offset of the input function, the time of injection is determined automatically by estimating the bolus arrival time (Section 2.5.2) from the list-mode data of the PET data. If metabolite correction is necessary, the metabolite analysis is also performed. The arterial sampling is not continuous and the frequency of sampling may change throughout the duration (e.g. 10 second samples for the first two minutes, followed by 5 samples every one minute, then samples at twenty, forty-five, sixty, and ninety minutes post-injection), requiring interpolation of the data. The most basic form of interpolation used is linear interpolation. However, linear interpolation of the data is sensitive to noise in the measurements. To address this issue, the input function and metabolite curves are fit with functions to provide smooth, continuous estimates. The bolus radiotracer injection is modeled using a piece-wise linear and three-exponential functions

$$f(t) = \begin{cases} 0 & t < \tau_s \\ A_1(t - \tau_s) & \tau_s \leq t < \tau_p \\ (A_1\tau_s - A_2 - A_3)e^{-\lambda_1(t-\tau_p)} + A_2e^{-\lambda_2(t-\tau_p)} + A_3e^{-\lambda_3(t-\tau_p)} & t \geq \tau_p \end{cases} \quad (2.11)$$

where A_s are the amplitudes, λ s are the decay constants, τ_s is the bolus arrival time, and τ_p is the peak time. Additionally, for fast bolus injections the model presented by Feng *et al.* is also implemented [175],

$$f(t) = \begin{cases} 0 & t < \tau \\ (A_1(t - \tau) - A_2 - A_3)e^{-\lambda_1(t-\tau)} + A_2e^{-\lambda_2(t-\tau)} + A_3e^{-\lambda_3(t-\tau)} & t \geq \tau \end{cases} \quad (2.12)$$

where the initial estimates of the parameters are performed using a three-stage procedure [176].

Manual sampling by hand may prove time consuming, prone to errors, and unnecessarily subject study staff to additional radiation. As an alternative an MR-compatible automated blood sampler was built following the work of Breuer *et al.* [177].

Chapter 3 MR-BASED METHODS FOR ADDRESSING SPECIFIC CHALLENGES TO PET QUANTIFICATION IN PET/MR

3.1 Attenuation Correction

As explained in Section 1.1.2.1, estimation and correction of the attenuation is vital for accurate PET quantification as it not only impacts the data by accounting for absorbed or scattered photons, but it is also used in the modeling of the scatter correction and can directly impact the contrasts of an image. As the BrainPET does not have a transmission source, the map of linear attenuation coefficients (LACs), or mu map, must be acquired separately (e.g. CT scan) or derived from the MR and/or PET data, all options supported by Masamune. Examples of mu maps generated using the currently implemented methods described below can be found in Figure 45 for a representative subject.




Figure 45 Implemented Mu Map Options. Continuous mu maps methods (top-from left to right) include CT, Pseudo-CT, and continuous UTE-Atlas-based mu maps. Segmented mu map methods (bottom-from left to right) include UTE-Atlas-based, UTE, and Single Tissue Model.

3.1.1 Head Attenuation Map Estimation

3.1.1.1 Single Tissue Model-based

The simplest method for estimating the head mu map is to fit an ellipse to the brain and/or head and assign a single linear attenuation coefficient (e.g. corresponding to soft tissue) to all the pixels inside the ellipse. The fitting could be performed manually on reconstructed uncorrected images [178, 179] or through the use of automated algorithms that estimate the contour of the head from the sinogram data [180, 181]. While sinogram based methods were less prone to subjective bias and considerably faster than manual slice-by-slice registration of ellipses, they also made the assumption that the emitting object was well approximated by an ellipse.

PET image quality has significantly improved with the advent of higher resolution full 3D PET acquisition without septa and iterative reconstruction methods. Similarly improved CPU

performance allows for real time processing. Rather than assuming an elliptical shape for the head or relying on the sinogram data, the uncorrected PET image can be used to generate the mu map using a series of morphologic functions. The accuracy of the PET image derived mu map may depend on numerous parameters, including the radiotracer, frame duration, and shape of the object being imaged, typically requiring manual intervention and optimization. To simplify this process we have constructed an interactive GUI that relies on morphological functions to generate a single tissue model-based mu map from an image in real time (Figure 46). This method has been used for animal studies performed on the BrainPET [164, 165, 182]. Since the skull of the small animals (e.g. rat, rabbit) is thin, the bias introduced by ignoring the bone is small when using this method. Similarly, this method is favorable in nonhuman primate imaging as no atlas or MR-based methods have been implemented.

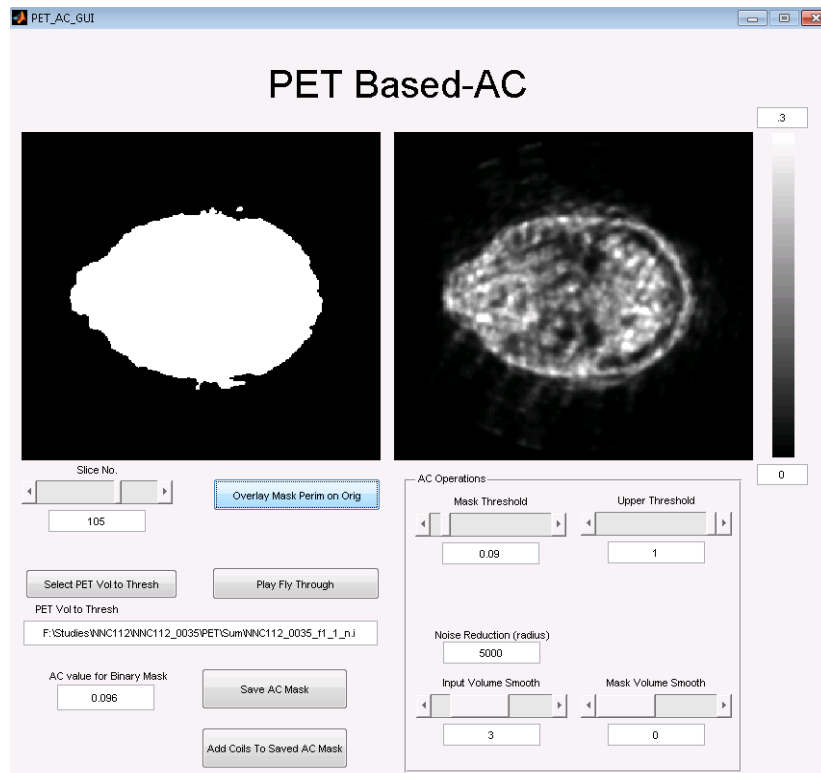


Figure 46 PET-based Attenuation Correction GUI

3.1.1.2 CT-based

Similar to PET, the contrast in CT is influenced by properties of the electron density of the object; however, unlike PET, in CT the penetrating photons do not come from within the body, but rather from an extracorporeal source that is rotated around the object. CT contrast is directly related to the attenuation properties of the tissues in the field of view for photons in the 40-140 keV energy range. Unlike CT, which has a range of photon energies, the photons in PET are nearly monochromatic with an energy of 511 keV. The attenuation properties of materials change at different energies so a transformation must be applied to the CT values to convert them to the linear attenuation coefficients for 511 keV photons. This scaling can be performed using a bilinear approach where the CT values below a certain threshold are scaled with one constant, while another constant is used for the values that exceed the threshold [12, 183].

For the generation of a CT-based mu map for use with the BrainPET, first the CT values are scaled to 511 keV. As the thresholds and linear scaling factors vary with beam energy and manufacturer, a lookup table of equations based on the experimentally determined values for the Siemens [184] and GE scanners [185] are used for the BrainPET. The original CT image is then registered to the subject's MPRAGE using a 12 degree of freedom affine transformation as determined by SPM8's normalized mutual information coregistration algorithm. The affine transformation is then applied to the scaled mu map. Finally, the repositioned scaled mu map is resliced to match the dimensions of a PET image, a 256×256×153 array with an isotropic voxel size of 1.25 mm, and the mu map of the MR hardware is added to the image.

3.1.1.3 Advanced MR-based Methods

Unlike CT, MR measures aspects of the proton density rather than electron density. As a result, MR cannot directly measure the attenuation properties of an object and its attenuation map must be estimated by indirect means. MR-based attenuation correction (MR-AC) techniques can be divided into two broad categories: segmentation-based techniques, which attempt to classify a voxel as a tissue type (e.g. bone, air, fat, soft tissue, etc) based on a single or a set of MR images and then assign it a linear attenuation coefficient based on its tissue class; and atlas-based techniques, which attempt to register a predefined template to morphologic MR images of the subject. As both categories have their drawbacks, hybrid methods that include aspects of both techniques also exist. The approaches developed for the BrainPET include a direct segmentation-based method based on the signal intensities from a pair of MR images (Section 3.1.1.3.1), a hybrid approach that segments the tissue based on an *a priori* anatomical atlas and the MR signal intensities (Section 3.1.1.3.2), and a direct atlas based technique where the subject-specific mu map is created through nonrigid registration of an population-based atlas to the subject (Section 3.1.1.3.3).

3.1.1.3.1 Ultra-Short Echo Time MR-based Methods

One of the most straightforward methods to derive the attenuation map of an object using MR is to simply segment the object into tissue classes and then assign an average linear attenuation coefficient to each tissue class. This presents a challenge in traditional MR imaging as the intensity of the signal is based on the hydrogen nucleus and cortical bone is mostly composed of atoms with higher atomic numbers like calcium leading to lower signal intensities

and shorter transverse relaxation times compared to soft tissues. This low contrast makes distinguishing air from bone, respectively the lowest and highest attenuating tissues in the head, quite challenging.

One method to address this challenge is with the use of ultra-short echo time (UTE) sequences which are capable of imaging tissues with short transverse relaxation times. If the MR signal intensity at a point in space is given by I_0 at time 0, the intensity at the echo time (TE) is given by

$$I = I_0 e^{-TE/T_2} \quad (3.1)$$

where T_2 is the transverse relaxation time. Thus given a pair or more images with different TEs, bone and air can be segmented. Keereman *et al.* performed this segmentation using empirically derived thresholds on the relaxivity maps (R2), given by

$$R_2 = \frac{\ln I_1 - \ln I_2}{TE_1 - TE_2} \quad (3.2)$$

where the subscripts one and two represent a short and longer echo time, respectively [21]. An alternative approach is to generate two sets of images with two empirically defined thresholds, one where bone is enhanced and one where air can be extracted, leaving all remaining voxels as soft tissue. To accomplish this with two UTE images, the bone enhancing (B) image can be generated as

$$B = \frac{I_1 - I_2}{I_2^2} \quad (3.3)$$

And the air image (A) as

$$A = \frac{I_1 + I_2}{I_1^2} \quad (3.4)$$

For the BrainPET a dual echo UTE sequence was implemented with echo times of 0.07 and 2.46 ms. Linear attenuation coefficients of bone and tissue were set to 0.151 cm^{-1} and 0.096 cm^{-1} , respectively [186]. While PET images generated with a UTE-derived mu map were shown to agree with those generated using a scaled CT, the misclassification of voxels within the brain was apparent and the thickness of the skull was grossly underestimated. These errors were in part due to spatial distortions between the two UTE echoes and were most prominent at the boundaries between tissues, e.g. tissue/cerebrospinal fluid. Similarly, coil sensitivity was not the same across the images leading to significant hotspots which occasionally introduced bone on the outside of the head. The use of non-rigid registration and intensity bias-field removal could potential mitigate these problems.

3.1.1.3.2 Probabilistic and Atlas-based Methods

A common artifact in UTE-based mu maps is voxels in the brain being misclassified as bone or air. In fact one major shortcoming of the purely data driven segmentation methods is that they do not include *a priori* knowledge of the shape or structure of the head. To address this issue and to improve the segmented mu maps, a probabilistic atlas-based segmentation approach was implemented [122].

The probabilistic-atlas classifies a voxel as air (0 cm^{-1}), tissue (0.096 cm^{-1}), or bone (0.151 cm^{-1}) based on image intensity and position information. The position and intensity information is derived from a set of MR images, an MPRAGE and two UTE images (with echo times equivalent to those in Section 3.1.1.3.1). The spatial-priors atlas portion, which provides the *a priori* position information was generated from a set of CT and MPRAGE images. The CT images were first registered to the MR images using a 6 DOF mutual information coregistration

algorithm. Next the corresponding MR images were registered to the MNI152 T1 template using a 12 DOF coregistration and the resulting deformation field was applied to the CT data. The transformed CTs were then segmented and the probabilistic-atlas was determined by binarizing the labels and normalizing the occurrences of a tissue class at a given voxel location by the total number of datasets used to generate the atlas. The intensity portion of the atlas was determined by training a classifier on the 3D histogram of the MR images. The UTE images were moved to the MNI152 space similar to the CT. The A and B images were derived from the two UTE images, providing each location in space with three intensity values (I_{MPRAGE}, I_A, I_B) which determined a histogram bin. Using the MNI registered segmented CT, the tissue class label, L_m could be determined and the intensity-based probability space mapped.

The mu map using the probabilistic atlas requires the acquisition of two UTE echos and an MPRAGE for the subject. The data are registered to the atlas and the probability of each tissue is given by

$$P(L_m | I_{MPRAGE}, I_A, I_B) = \frac{P(I_{MPRAGE}, I_A, I_B | L_m)P(L_m | X_n)}{\sum_m P(I_{MPRAGE}, I_A, I_B | L_m)P(L_m | X_n)} \quad (3.5)$$

Where X_n is the voxel location, $P(L_m | X_n)$ is the tissue class probability calculated from the spatial-priors atlas for the voxel, and $P(L_m | I_{MPRAGE}, I_A, I_B)$ is the probability calculated from the intensities.

While the probabilistic-atlas showed better results compared to the UTE-based method with regard to skull thickness and the misclassification of voxels inside the brain, it is still limited to segmented tissue classes.

3.1.1.3.3 *Pseudo-CT-based Methods*

The template-based approach was first suggested by Kops *et al.* [24]. In their method, a nonlinear spatial normalization algorithm was used to register MR and PET transmission images to their respective templates independently to generate a T1-atlas and the corresponding transmission atlas. The T1-atlas was then registered to the specific subject's T1 images and the transformation was applied to the transmission atlas to generate the subject specific mu map. The correlation between the atlas-based subject specific transmission data and the measured transmission data was improved upon by Malone *et al.* by first coregistering the transmission images to the T1 images, determining the transformations for only the T1 images, and then applying the same transformation to the transmission images to generate the atlas [187].

We developed an improved method for generating the mu map from an MPAGE volume using SPM and FreeSurfer. Briefly, a unified segmentation/spatial normalization approach is performed on an MPAGE image to create probability maps of tissue classes (bone, air, noncerebral soft matter, gray matter, and white matter) from the MR image. To improve the segmentation and registration, the intensity bias field was removed from the MPAGE image using FreeSurfer. The segmented probability maps can then be used for a more accurate registration between a subject and an atlas.

An atlas was generated from the images of 15 glioblastoma patients by registering the tissue probability maps to a common space using DARTEL, a fast diffeomorphic image registration algorithm [188]. The transformations mapping the individual atlas subject was then applied to the scaled CT image to create an average mu map. To generate a subject specific mu map, first the individual subject tissue probability maps are registered to the atlas of tissue

probability maps using DARTEL to derive the nonrigid transformation. The inverse transformation is then applied to the atlas mu map to generate a subject specific mu map.

This method generates continuous valued mu maps using only the MPRAGE image which is acquired for anatomic localization as well as for use with FreeSurfer. The major drawback to this method is that it presupposes the subject has an anatomy that is well represented by the atlas, as elements (e.g. missing anatomy) cannot be completely removed from the atlas by nonlinear registration but only distorted. Also it may require the generation of a separate atlas for pediatric patients.

3.1.2 MR-Hardware Attenuation Map Estimation

In the case of the MR-BrainPET scanner, the mu map of the MR hardware (e.g. RF coils) has to also be considered. As the hardware does not change and its position relative to the PET field of view remains constant due to locks that engage within the MR table, the hardware mu map only needs to be derived once per piece of hardware, from a CT of the MR hardware. An explanation of how the CT of the coil is converted to a mu map is provided in Section 3.1.1.2. As the MR hardware does not emit coincident 511 keV photons nor does it show up in the MR images, the position of the coil in the PET field of view cannot directly be determined. Instead, small spherical fiducial markers were fixed to the object and filled with water or F-18 solution for CT and PET imaging, respectively. After the resulting volumes were coregistered, the CT-derived mu map was positioned in the PET space. The mu maps of several coils (e.g. CP coil, an 8 channel receive coil, a 31 channel receive coil [117], a large and a small primate brain coil) have been generated and can now be selected from the Masamune GUI.

The inner diameter of the BrainPET gantry is 35 cm while the imaging field of view is 32 cm, which leaves 3 cm of space between where the imaging field of view ends and the real boundary of the scanner. This is typically occupied by the CP coil and its elements, which have a nonnegligible contribution to the coincident photon attenuation. To account for the material that lies outside of the 32cm imaging field of view, the coils are stored in $320 \times 320 \times 153$ matrices with an isotropic voxel size of 1.25mm. To generate the complete mu map for data reconstruction, the hardware mu map is cropped to the standard matrix size and added to the head mu map. To derive the attenuation sinogram, the head and the extended hardware mu maps are forward projected into sinogram space and then multiplied.

3.2 Partial Volume Effects Correction

Even with an accurate estimation of the attenuation, the reconstructed distribution of the PET radionuclide may be inaccurate and introduce bias into ROI and voxel-wise analyses, even in phantom scans where there is minimal motion. These inaccuracies come from partial volume effects that arise from the low spatial-resolution of the PET detectors, the voxelization of the object, the regional point-spread function (rPSF) of the PET camera, and the relative contrast in the radionuclide concentration across the object of interest. Typically the PSF of a PET camera can be measured and is typically approximated as a Gaussian, which can then be used to correct for these partial volume effects post-reconstruction. While post-reconstruction partial volume correction can recover much of the information content that is lost, modeling the system in the reconstruction can significantly improve the initial image, however in

modeling the rPSF in the reconstruction, the residual rPSF of the resulting image, RrPSF, may not be accurately modeled by a Gaussian.

3.2.1 Applying a PSF to an ROI

In Section 2.2 the optimization of the OSEM algorithm for the BrainPET was discussed. The reconstruction was optimized by including a PSF model in the algorithm immediately prior to the forward projection of the iteration estimate. By modeling the rPSF of the camera, it is possible to minimize its contribution to the reconstructed image. If the model perfectly described the scanner then the reconstructed image would reflect the true activity once it was smoothed to account for the noise introduced by the reconstruction. Fully modeling the system is challenging and would require complete control over the scanning environment and that the performance of the system be fully characterized, which is unreasonable. Thus, even with the application of the rPSF in the reconstruction there is some residual blurring of the data. This residual ambiguity in the distribution can, however, be removed with the application of a partial-volume effects correction algorithm.

Given a rPSF, the measured activity in an isolated region of uniform activity will be corrupted by spill-out effects. To account for spill-out, ROI-based partial volume effects correction techniques require the ROIs be blurred by the scanner rPSF. To determine the spill-out effects for a composite set of ROIs, they are first separated into individual files, each containing a single ROI of the set. The images are then blurred using the rPSF of the scanner (e.g. 3 mm Gaussian and b=16 motion blur) to generate the expected image from the scanner if the resolution was not modeled. The blurred ROI is then forward projected into sinogram

space and reconstructed using OSEM with resolution modeling and reconstruction parameters that match the data to which the correction will be applied to. An example of the original binary ROI for white matter, as well as the ROI blurred by the rPSF, and finally the image blurred by the RrPSF resulting from forward projecting and reconstructing the blurred ROI can be found in Figure 47.

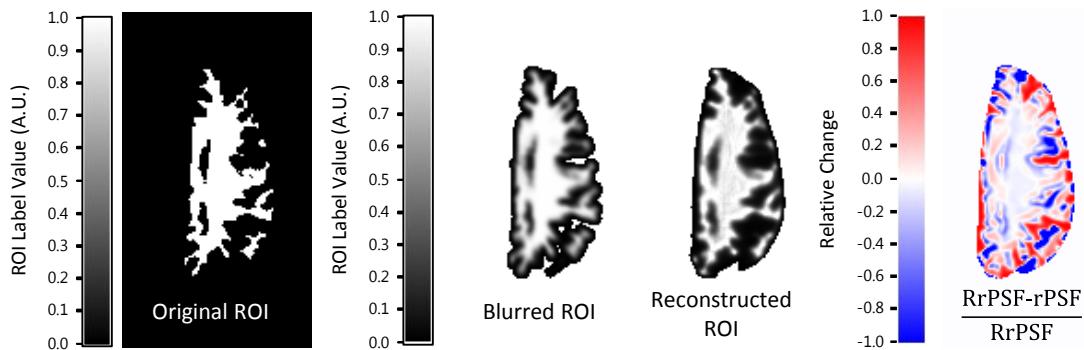


Figure 47 Partial Volume Effect Correction Processing for Region of Interest. Processing stream (from left to right): Original ROI must be blurred by the point spread function of the system for accurate PVEC; if resolution modeling is implemented in the reconstruction, then the blurred ROI must be forward projected and reconstructed with resolution modeling to yield the necessary blurred image. The relative change between the rPSF blurred image and the residual rPSF (RrPSF) image shows considerable differences at the edge of the ROI.

While partial volume effects have an influence on any ROI measurement, regions that are larger than three times the spatial resolution of the camera contain voxels near their centers that are not significantly influenced by partial volume effects. However, in brain imaging most of the regions of interest resemble sheets rather than spheres so partial volume effects may be significant. A number of ROI sets have been defined on the MNI brain allowing the distribution to be estimated in a number of structural and functional regions. To accurately estimate the activity inside of a ROI, partial volume effect correction must be performed which requires knowledge of the PSF of the system. The PSF has been modeled for the scanner space, but is not currently sophisticated enough to allow it to be spatially warped to MNI space, requiring the ROI estimation be performed in subject space if partial volume effects are to be

considered. An algorithm was discussed in Section 2.6.2.2 which converted the MNI labels into subject specific labels.

3.2.2 Geometric Transfer Matrix

A geometric transfer matrix (GTM)-based partial volume effect correction algorithm scheme has been implemented on the BrainPET. The GTM PVEC was implemented as it was more general than the Muller-Gartner approach, where the white matter ROI is assumed to provide an unbiased estimate. The GTM method presented by Rousset *et al.* relates the measured mean activity in a set of ROIs, $ROI_{N,mea}$, to the true mean activity in those ROIs, $ROI_{N,true}$, through a system of linear equations,

$$\begin{bmatrix} ROI_{1,mea} \\ ROI_{2,mea} \\ \vdots \\ ROI_{N,mea} \end{bmatrix} = \mathbf{W} \times \begin{bmatrix} ROI_{1,true} \\ ROI_{2,true} \\ \vdots \\ ROI_{N,true} \end{bmatrix} \quad (3.6)$$

where W is the GTM matrix [189]. If each ROI is considered a functionally distinct region of uniform activity, then the measured mean activity inside the region is given by

$$ROI_{j,mea} = \frac{1}{n_j} \sum_{i=1}^N ROI_{i,true} \int_{V_j} RSF_i(\mathbf{r}) dr \quad (3.7)$$

where n_j is the number of voxels in ROI_j , V_j is the volume bound by ROI_j and $RSF_i(r)$ is the region spread function, the result of convolving a unit mask of the i^{th} ROI by the rPSF over the field of view. The RSF for a given region was shown in Section 3.2.1 and is the spill-in of ROI_i into ROI_j . The measured activity of ROI_j is the combination of the spill-in factors of the ROIs multiplied by their mean values,

$$ROI_{j,mea} = \sum_{i=1}^N \omega_{ij} ROI_{i,true} \quad (3.8)$$

where ω_{ij} is

$$\omega_{ij} = \frac{1}{n_j} \int_{V_j} RSF_i(\mathbf{r}) dr \quad (3.9)$$

The i th row and j th column of the \mathbf{W} matrix is then given by ω_{ij} . The diagonal terms in the \mathbf{W} matrix represent the contribution of the ROI to itself after the spill-out is accounted for while off-diagonal terms account for spill-in to other regions. As the number of voxels in a given region may vary, and due to the isotropic blur introduced with the motion blur, \mathbf{W} may not be symmetric.

3.2.3 Region-based Voxel-Wise Correction

While the GTM approach can be used for correcting ROI-based measurements, it cannot directly provide voxel-wise measurements like the Muller-Gartner method, which in turn limits analysis of the data from the application of group-wise statistic parametric maps. An extension of the GTM-based partial volume effect correction capable of providing voxel-wise estimates was proposed by Thomas *et al* [65] - the region-based voxel (RBV)-wise correction.

The RBV method is fundamentally similar to how an isolated ROI would be corrected for spill out effects. First, GTM-based partial volume correction must be performed on a PET image using a set of ROIs, providing a set of mean values for the ROIs. A composite image of the GTM corrected activity is then created where all the voxels of a given ROI are assigned the mean value of that ROI. The composite image is then blurred with the PSF following Section 3.2.1. The voxel-wise corrected image is then given by

$$PET_{RBV}(\mathbf{r}) = PET_{mea}(\mathbf{r}) \times \frac{PET_{GTM}(\mathbf{r})}{PET_{GTM}(\mathbf{r}) \otimes PSF(\mathbf{r})} \quad (3.10)$$

where the mea, RBV, and GTM subscripts represents the type of PET images.

As the RBV correction method is general for any PSF, it is possible to correct raw PET images, images where the rPSF for the scanner is not modeled in the reconstruction, for both partial volume effects and spatial distortions introduced by the rPSF of the scanner.




Figure 48 Rim Artifacts in RBV Images Without rPSF Modeling. Data were reconstructed with rPSF and without rPSF in the reconstruction and then partial volume effect correction was performed with the corresponding rPSF, or lack thereof, to account for additional error introduced in the reconstruction. A significant edge artifact (black arrow) is present around the rim of the RBV brain that was reconstructed without rPSF modeling in the reconstruction.

While RBV can correct images for both partial volume effects and spatial distortions, it may introduce a hyper-intensity artifact at the rim of the volume with significant bias on the order of 50% or greater when compared to images where the rPSF is modeled in the reconstruction (Figure 48). As such it is ideal to include the rPSF in the reconstruction.

3.3 Head Motion Compensation

The diagnostic power of PET (and MRI) can be significantly reduced as a result of patient motion [67], either from physiologic or from stochastic motion, which can both lead to image degradation and artifacts. Furthermore, some patient populations (e.g. pediatric patients, Alzheimer's disease patients, non-compliant individuals) are particularly prone to intra-scan

stochastic motion making it difficult to accurately derive quantitative or, in some cases, even qualitative information. To minimize the effects of stochastic motion in neurological studies involving MRI and PET, a number of different head restraints have been proposed (e.g. molded plastic masks, orthopedic collars and straps, bite-bars, and vacuum-lock bags). The application of these methods in neurological studies has been shown to be largely unreliable and not always well tolerated by patients; in some cases inducing patient motion to ameliorate discomfort [69, 190, 191]. An alternative method is to allow for the free motion of the patient, but to track and correct for it using a software approach. A number of methods to correct for stochastic motion have been proposed for both MR and PET. It is important to note that even with the implementation of a software motion correction approach, clever design of the transmit and receive coils can be used to significantly reduce the possible range of motion while remaining comfortable to the patients [192].

3.3.1 Rigid-body Motion Estimation

3.3.1.1 MR-based Motion Estimation Methods

Even though MR possesses the ability to collect high contrast-to-noise volumes faster than PET, motion can still be present and a number of MR based methods have been developed to address this issue. The simplest of methods involves post-processing the data. If anatomical images are collected throughout the session, they can be coregistered to a reference image using numerous published and well characterized algorithms to provide estimates of the patient motion [193-199]. These image-based methods however suffer from a poor time resolution equal to the time between the start of subsequent anatomic scans and thus has no

way to detect motion that may have occurred in any interwoven non-anatomic scan. Similarly, motion during the collection of the anatomic scan can lead to image degradation that may in turn adversely affect the software's ability to successfully coregister the anatomic volumes. With the advent of fast imaging sequences (e.g. EPI) it was possible to collect whole volumes with high temporal-resolution, which could be compared to a reference for estimating the motion. The EPI images could be collected between subsequent scans or interwoven into scans. An even more robust method of motion tracking is to use navigator-based "real-time" tracking methods. The earliest techniques involved acquiring single lines of k-space, performing a fast-Fourier transform, and comparing it to some reference projection [200]. While these 1D navigator techniques were sensitive to motion, they could only decode it along a single axis and could not easily decode translations and rotations. Subsequent techniques sought to minimize navigator time by choosing optimized k-space trajectories which could provide intrascan motion estimates and to provide information in all three spatial dimensions [201].

3.3.1.1.1 *Echo-planar Imaging and Prospective Acquisition Correction*

EPI sequences are capable of acquiring high signal-to-noise images of the whole brain within seconds making it appealing for functional MRI where aspects of the brain could be rapidly imaged during the presentation of a stimulus. Misregistration in the functional data due to spurious subject motion has been shown to lead to variations in the time course as well as to introduce artifacts in the data [202-204]. A number of techniques have been implemented for the rigid-body post-processing realignment of EPI data and while realignment has been shown to minimize motion related artifacts, they may still be present and affect subsequent analysis [204].

Many registration techniques do not account for the effects of through-plane motion on the images or distortions due to susceptibility effects resulting from the repositioning of the head in the magnetic field. Image-based post-acquisition correction for these effects would require substantial information about the system and modeling. However, due to the high temporal-frequency at which the images are collected and fast reconstruction techniques, it is possible to estimate the rigid-body motion during acquisition. This is the basis of the prospective acquisition correction (PACE) implemented on the Siemens scanners, whereby every MR repetition period, an EPI series is acquired and registered to the first volume of the series. From these estimates the gradient fields can be adjusted for the next EPI acquisition [205]. The use of PACE has been shown to reduce the presence of motion artifacts [206, 207].

In addition to updating the gradients, PACE records the motion estimates and stores a motion corrected EPI dataset for subsequent analysis. Due to time constraints required to optimize the registration algorithm for speed, the estimates are coarse compared to the resolution that is possible when performing the registration after acquisition, when time is less of a factor. Additionally, the motion estimates recorded by PACE are with respect to the MR reference position, which is slightly different than the PET reference position. To derive the MR-based motion estimates, the PACE corrected EPI volumes that have not been reoriented are used to estimate the motion of the subject.

SPM8 was used for the registration of EPI volumes as it allows automatic processing through Matlab-based scripts, can be extended to include the effects of image distortions due to field inhomogeneities resulting from motion, and was shown to have a good overall performance in the registration of fMRI data [208]. The SPM8 algorithm uses a least-squares

approach and performs the registration in two steps. First the images are aligned to the initial image in the time series. A mean image is then created from the aligned images and all the images are aligned to the mean image. Registering the first image to the mean image may lead to a slight offset in the first image which must be accounted for as the goal is to estimate the subject motion over time, not to simply align the MR data. The specifics of how to correct the data for this effect are outlined in Section 3.3.1.2. The least-squares registration implementation allows selectable parameters associated with quality (corresponding to the number of voxels used), the sampling separation, interpolation schemes, and the use of a 3D Gaussian blur prior to registration. The registration was optimized for accuracy using quality=1, fwhm=1, sep=1, and interp=5.

In sequences where the contrast changes greatly between frames, and thus the first image may have substantially different contrast than the rest, as is the case in DTI, a normalized mutual information based-approach was found to provide less noisy estimates. A similar two-pass procedure is followed where the images are registered to the first image in the series followed by the generation of a mean image and the registration of all the images to the registered image.

3.3.1.1.1.1 Effect of Distortion Correction on the Motion Estimates

Derived from EPI Images

A majority of the MR-based motion estimates derived from MR data on the BrainPET come from EPI sequences. Geometric distortions arising from magnetic field inhomogeneities are a well known limitation of EPI. These nonuniform distortions lead to the compression and expansion of parts of the brain, influencing the signal intensity and shape of the brain, which

has been shown to make accurate registration to anatomical MR images difficult [209, 210]. An example of an EPI with these nonuniform geometric distortions can be found in Figure 49.




Figure 49 Effect of B_0 inhomogeneity on EPI registration. Coregistration of an uncorrected fMRI EPI (left) with a MPAGE (right) with contours of EPI image in red. Note the pronounced dropout of EPI signal in the axial view of the center of the volume when compared to the anatomical image.

The susceptibility of EPI to geometric distortions due to field inhomogeneities arises from the low bandwidth of the MR data acquisition in the phase-encode direction. It has been shown that the distortion is a result of local warps in the phase encode direction and the distortions can be corrected for with a local mapping of the field [210, 211]. While field maps have been shown to improve the geometric distortions and alter the resulting statistics, they are not always collected before every EPI due to protocol constraints like total scan time, the region of the brain that is being analyzed, and investigators preference. From a simultaneous PET/MR point-of-view, the addition of field maps before every EPI scan would add considerable time to MR imaging during which the PET signal would continue to decay (e.g. for C-11 the 63 second acquisition time leads to a reduction in count rate by 3.5%). As the addition of field maps would reduce the PET SNR and in turn make correlation with simultaneous MR-based

functional measurements more challenging, it is important to determine the added value of correcting the geometric distortion. One alternative option is to collect only one field map for each EPI type and use it for all subsequent EPIs of that type (as the field map must be collected using parameters that match the individual EPI type). While this option may not be appropriate for correcting the fMRI data for statistical analysis it may be suitable for registration.

To determine the utility of field maps for registration, a volunteer was placed inside the BrainPET in the head first supine position and MR imaging was performed using the CP + 8-channel MR Tx/Rx coil. Cheek pads were placed according to standard PET/MR imaging procedure to limit the motion of the subject. The imaging protocol consisted of pairs of field maps and BOLD EPI sequences with the field map always preceding the BOLD. During the scan the subject was asked to make small motions inside the head coil around a neutral center position when instructed. The subject was asked to remain still throughout the field map collection. Prior to the collection of the sixth and seventh field maps the subject was asked to make a large motion and select a new neutral center position.

The subject motion was estimated using the algorithm outlined in Section 3.3.1.1.1 except that different field map paradigms were used. As the goal of the experiment was to determine the added utility of collecting additional field maps, the motion estimates where all of the individual field maps were used to correct their corresponding EPI datasets was used at the true values. To determine the added benefit provided by distortion correction, the motion estimates from the uncorrected EPI images were also determined. Additionally, the subject motion was estimated using only the field map collected prior to the first BOLD sequence to correct for all subsequent BOLD EPIs.

Typically, the efficacies of field mapping and distortion correction are evaluated based on their improvement of the statistical results from the MR analysis; however, for application to motion correction, it is the motion estimates that are of ultimate interest. A sphere with a radius of 4 mm was constructed at a point 30 mm away from the x-, y-, and z-axes (i.e. center at [30mm, 30mm, 30mm]). The different motion estimates were applied to the sphere and the percent volume overlap was calculated for the true measurements and the alternative measurements (e.g. field map corrected versus uncorrected images).




Figure 50 Effect of Alternative Field Mapping Scheme on Motion Estimates. The motion of the patient as recorded over eight EPI sequences with field maps preceding each EPI sequence as measured using SPM8's rigid-body registration algorithm with distortion correction (top). Percent overlap between an off-center 4 mm ROI repositioned using the reference motion and either the estimated motion using no distortion correction (No DC; red) or only the first field map for distortion correction (First Map DC; blue).

Eight sets of field map and EPI pairs were collected over the scan time of approximately 1800 seconds. One field map had to be repeated. Motion was successfully estimated over the scan time of 1722 seconds for the three field map cases (all field maps, uncorrected/no field maps, and only one field map). The results can be found in Figure 50. Using only one field map

to correct all subsequent EPI sequences gave more consistent results in the individual scans than using only the uncorrected data. Similarly, the estimates from using only one map performed better than no field map correction in a majority of cases. The uncorrected EPI performed better when there was significant motion during the scan. One possible explanation for this outcome is the failure of the individual EPI volumes to accurately register to the mean EPI volume as the simultaneous registration and unwarping is performed in two passes. Additionally, prior to the third EPI acquisition the subject moved during the initial shimming between the end of the field map collection and the start of the EPI series which may have contributed to observed difference.

The error in motion estimation caused by ignoring distortion correction for the EPI images lead to a minimum percent overlap of 90.6% for the uncorrected data, and 90.4% for the data corrected using only the first field map. However, there was marginal improvement in the overlap with the true values when the data were corrected using only the first field map versus uncorrected, with a maximum difference in the overlap of 2.6%. Using only the first field map showed less stochastic noise in the estimates. These results suggest that there is a detectable difference in the motion estimates when the geometric distortions are considered in the registration algorithm. Rather expectedly, repeated analysis showed that smaller ROIs further from the center of the field of view are more sensitive to this error than larger ROIs and those closer to the center. This suggests that while acquiring field maps before each EPI is ideal, the addition of a single field map for each EPI sequence type can be used to improve the estimates.

3.3.1.1.2 *Cloverleaf Navigators*

A higher temporal frequency method of motion tracking is to use cloverleaf navigator (CLN) “real-time” tracking methods. The earliest techniques involved acquiring single lines of k-space, performing a fast-Fourier transform, and comparing it to some reference projection [200]. While these 1D navigator techniques were sensitive to motion, they could only decode it along a single axis and could not easily decode translations and rotations. Subsequent techniques sought to minimize navigator time by choosing optimized k-space trajectories which could provide intrascan motion estimates and to provide information in all three spatial dimensions [201]. Masamune is capable of decoding the CLNs from the MR data for estimation of the bulk head motion.

3.3.1.2 *MR-based Motion Estimation Algorithm for Multiple Sequences*

A typical MR imaging study consists of multiple MR sequences (MPRAGE, BOLD, DTI, etc), many of which can be used to provide information about the subject’s motion. A retrospective analysis of 10 AD subjects showed that the motion estimates using MR-based methods could be derived for an average of ~60% of the study duration. As the study duration may include multiple sequences with different contrasts, linear-least squares registration may not be suitable to derive the motion. Furthermore, as some sequences like DTI have contrasts that vary throughout the sequence, a better registration can be obtained using alternative registration schemes. Instead, the relative motion is estimated within each sequence and then the absolute motion is estimated.

To estimate the motion over each individual sequence, the mosaic DICOM images are first converted to a 3D Nifti image with the orientation data corrected for the offset between the MR and PET scanners. Once all the images related to a single sequence are generated, they are registered using a two pass registration, first to the initial image in the series, then a mean image of the registered images is generated and the images are registered to the mean image. The choice of registration scheme was found to depend on the data and was most influenced by the change in contrast over an individual sequence. For BOLD and ASL imaging, where the contrast remains relatively constant over the sequence, linear least-squares fitting was performed as it was faster. For sequences where the contrast did significantly change over the imaging duration, e.g. DTI, the data were registered using a mutual information algorithm. At this point the images are registered to the initial image. If the sequence was the first sequence and occurred during the radiotracer injection, the relative motion estimates are the absolute motion estimates. However, for subsequent sequences the absolute motion estimates were determined using the mean images from the individual sequences.

Consider the 4x4 affine transformation matrix that registers an image to another image given by T_n such that

$$\mathbf{X}' = \mathbf{T}\mathbf{X}$$

where X are the points of the original image and X' are the points after registration. This equation will hold true for the registration of the n^{th} image of an image series to the first image in that series. Additionally it can describe the registration of the mean image of that series to the mean image of another series. Determining absolute motion for the n^{th} image of the i^{th}

series is simply the transform that would register the n^{th} image to the first image of the first series,

$$\mathbf{X}' = \mathbf{T}_{i1} \bar{\mathbf{T}}_{n1} \mathbf{X}$$

where the barred T is the relative motion of the n^{th} frame and the unbarred T is the absolute motion of the first image in the i^{th} series. However, as the registration was done in two passes, if there is no relative motion between the first frame of the series and the mean, then the mean images that have higher SNR can be used. Determining the order of the sequences was done automatically by reading the DICOM metadata, and the offset between the MR and PET clocks were then determined as per Section 2.4.2. A representative plot of the absolute motion estimates over a dynamic FDG-PET/MR scan of a healthy volunteer can be found in Figure 51.

Once the motion was estimated for the duration of the scan, the MPRAGE was registered to the mean of the first sequence, placing it in the motion corrected position. All MR-derived PET-based corrections (e.g. Automatic ROI Generation for Partial Volume Effect Correction, Attenuation Correction, etc) were taken relative to this registered MPRAGE.




Figure 51 Representative MR-based Motion Estimates for FDG-PET/MR Study. The motion of a healthy volunteer was tracked over simultaneous PET/MR imaging using the acquired EPI MR data. As multiple MR sequences may be used over imaging with different contrasts, mutual information based schemes as well as least-squares algorithms are used to derive the estimates. The estimates are recorded relative to a reference time which is taken as the bolus arrival time.

3.3.1.3 PET-based Motion Estimation Method

While the use of EPI- and CLN-based sequences allows for the tracking of the head motion, it is not feasible to implement these methods for all MR sequence. Similarly, no tracking is performed between sequences or during the automated preparation and coil/field adjustments. However, PET data are continuously acquired and motion can be significant during these moments between sequences. For instance, staff may talk to the patients to get confirmation that they are okay. Similarly, as the PET camera does not make as much noise as the MR, the subject may be unaware that data are being collected and may take the opportunity to move to a more comfortable position. To address the problem of motion when MR is unavailable, we explored the use of short-frame PET images to estimate the subject motion.

The method is similar to the concept of PACE imaging, where the estimation of the motion is derived by comparing the rapidly acquired images. Due to the significant differences between the origin and nature of the MR and the PET signals, the PET data must fulfill additional requirements for the method to succeed. Most important for registration is the extent of the signal; in most EPI images the brain has a suitable contrast-to-noise where it can be isolated from background. The signal in PET has a strong dependence on the specificity of the radioligand and therefore may not have adequate contrast-to-noise to isolate the entire brain. Thus this PET-based method requires ubiquitous uptake in the head, limiting its application to tracers that have a significant nonspecific binding or uptake. If there is sufficient nonspecific binding, then the shape of the head can be determined and the relative position of the head estimated. The optimal duration of the short frames is dependent on a number of factors, including activity, current activity distribution, and uptake pattern of the radiotracer, requiring that the optimal duration not only vary over the scan time, but also with the tracer.

An outline of the method is shown in Figure 52. First, the PET data are framed into short (e.g. 30-second) intervals and images are reconstructed accounting for the detector normalization and coil attenuation. Since the estimation is based on the coregistration of images, the smoothness of the image has a stronger impact on the cost function than the bias in the image. As such the PET data are reconstructed with 32 iterations and 1 subset. Since a frame invariant (i.e. static) attenuation map would introduce bias into any software-based coregistration scheme due to mismatch with the emission data, the individual normalization-corrected PET images are used to automatically derive the corresponding attenuation maps for further processing. The attenuation map is generated from a binary mask of the outline of the

head. The volumes were projected onto the x-, y-, and z- planes and thresholding is performed along with morphologic functions to improve the accuracy of the created mask. The 2-D projections are then stacked to span the image volume according to their planar orientation, and the mask volumes are multiplied together. The resulting volume mask is set to the linear attenuation coefficient of water (0.096 cm^{-1}) to yield the PET-based attenuation map. This attenuation map is combined with the attenuation map of the MR coil and used to generate an attenuation and scatter corrected image. The resulting images are then coregistered to an MPAGE image using the normalized mutual information algorithm in SPM. As these registration algorithms are sensitive to local minima, the MPAGE was corrected for the offset between the MR and PET fields of view to improve the initial overlap. The registration of each frame to the MPAGE was run as a multi-resolution two pass function [212], where the first pass sampled the data every 2 mm and the second pass sampled the data every 1 mm. Similarly, the motion estimates of the previous frame are used as the initial offsets for the subsequent image if available.

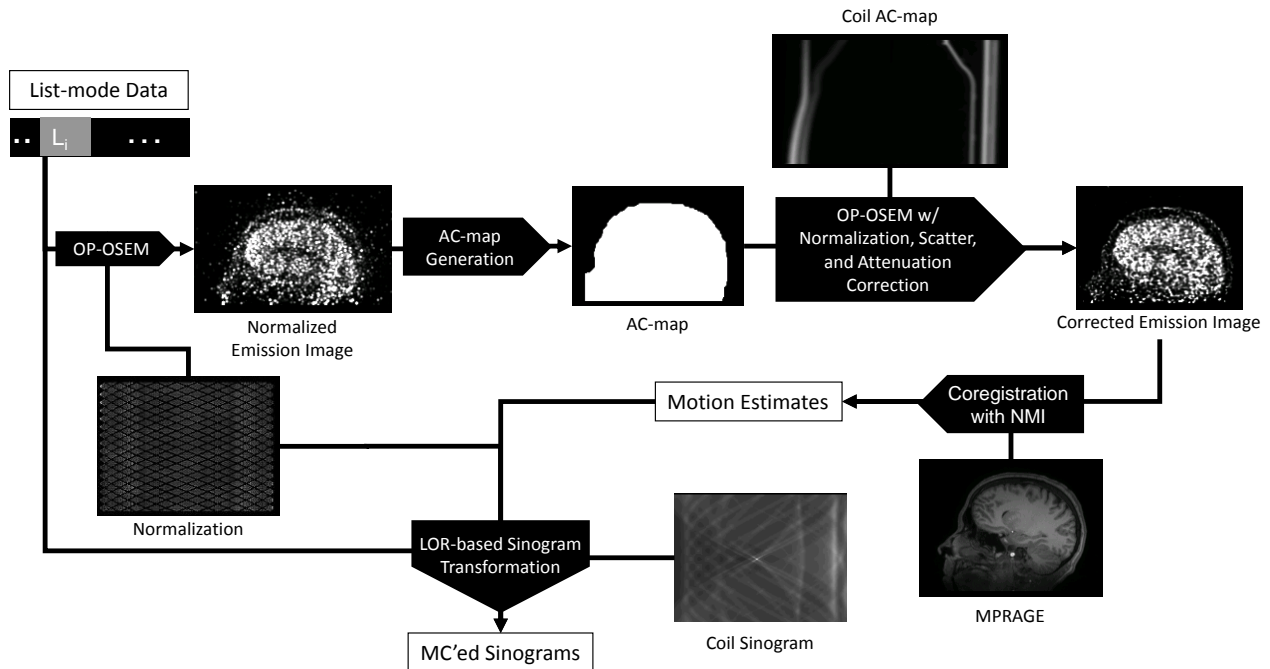


Figure 52 PET-based Motion Correction Algorithm. In this PET-based estimation of the subject motion, short frames are initially reconstructed without attenuation correction. Mu maps are derived from these reconstructions and the data are again reconstructed, but with attenuation and scatter correction. The resulting images are registered to a reference image and the relative motion is determined. The individual mu maps prevent bias due to mismatch between the emission and attenuation image from being introduced into the estimates.

The coregistration algorithms seek to maximize the shared information between the two images at a voxel intensity level. One potential source of misregistration comes from inconsistencies between the tissue intensity ranges. For instance, while grey matter has similar intensity across the cortical and subcortical structures, selective tracers (e.g. NNC112 which targets D1 receptors) can lead to a broad range of PET values in grey matter voxels. To address this potential source of error, once the motion estimates are determined, they are applied to the images, and a mean image from the short frame PET images is generated and the registration step is performed again, however the mean image is used rather than the MPRAGE.

At this point the images have all been registered to the MRPAGE, but the goal is to estimate the motion over the PET imaging duration, i.e. the motion relative to the beginning of

the scan. If the 4x4 affine transformation matrix that registers image n to the MPRAGE is given by T_n such that

$$\mathbf{X}' = \mathbf{T}_n \mathbf{X} \quad (3.11)$$

where X are the points of the original image and X' are the points after registration, then there exists an transform such that

$$\mathbf{X} = \mathbf{T}_n' \mathbf{X}'. \quad (3.12)$$

It can then be shown that T' is the inverse of T_n ,

$$\begin{aligned} \mathbf{X} &= \mathbf{T}_n' \mathbf{X}' \\ \mathbf{X} &= \mathbf{T}_n' (\mathbf{T}_n \mathbf{X}) \\ \mathbf{X} &= (\mathbf{T}_n' \mathbf{T}_n) \mathbf{X} \\ \mathbf{T}_n' \mathbf{T}_n &= \mathbf{I} = \mathbf{T}_n^{-1} \mathbf{T}_n \end{aligned} \quad (3.13)$$

where I is the identity matrix. If T_n is the transformation that aligns the PET to the MPRAGE and T_n^{-1} aligns an image registered to the MPRAGE to the n^{th} PET image, then the matrix M , which aligns the n^{th} PET image to the first pet image is

$$\mathbf{M}_n = \mathbf{T}_1^{-1} \mathbf{T}_n \quad (3.14)$$

Using this relationship the motion of each of the short frames relative to the first frame is calculated. The PET-based motion estimates for the same subject depicted in Figure 51 can be found in Figure 53.




Figure 53 PET-based MC Estimates for Dynamic FDG-PET/MR Study. The motion profile of the patient was determined using PET-based methods with a frame length of 30 seconds. The study used to generate this profile was the same study as presented in Figure 51, where MR-based motion estimation methods are shown.

3.3.1.4 Unified Motion Estimation

As mentioned, there are times when MR-based estimation of the subject motion is impossible, when no MR is being performed or between sequences. While the total time of these moments can be minimized through careful planning of the study, they cannot be removed completely. As such, PET-based estimates will always have a role in motion correction, albeit perhaps a minor one. It is important to note that both motion estimation schemes determine relative head motion and not the motion with regard to some ideal head position. In MR-based methods the estimates are determined relative to the MR image corresponding with the first motion estimate, while in PET-based approaches it is to the first PET frame. This may result in an offset between the PET- and MR-based estimates.

In order to correct for the difference in the reference volumes, the same math as described in Section 3.3.1.2 can be used and the offset applied to each frame. This requires the

registration of PET data to the first MR dataset. As with the MR-based motion correction, the mean of the first sequence is used. To maximize the anatomic detail, a single PET image whose duration spans the entire imaging time is created from the PET data using the PET-derived motion estimates. In this way, motion estimates can be provided throughout the entire study. This method, wherein the MR-based motion correction is extended using PET-based estimates derived is the PET-Assisted Continuation of MR Motion Alignment and Normalization (PAC-MMAN).

3.3.2 Rigid-body Motion Correction

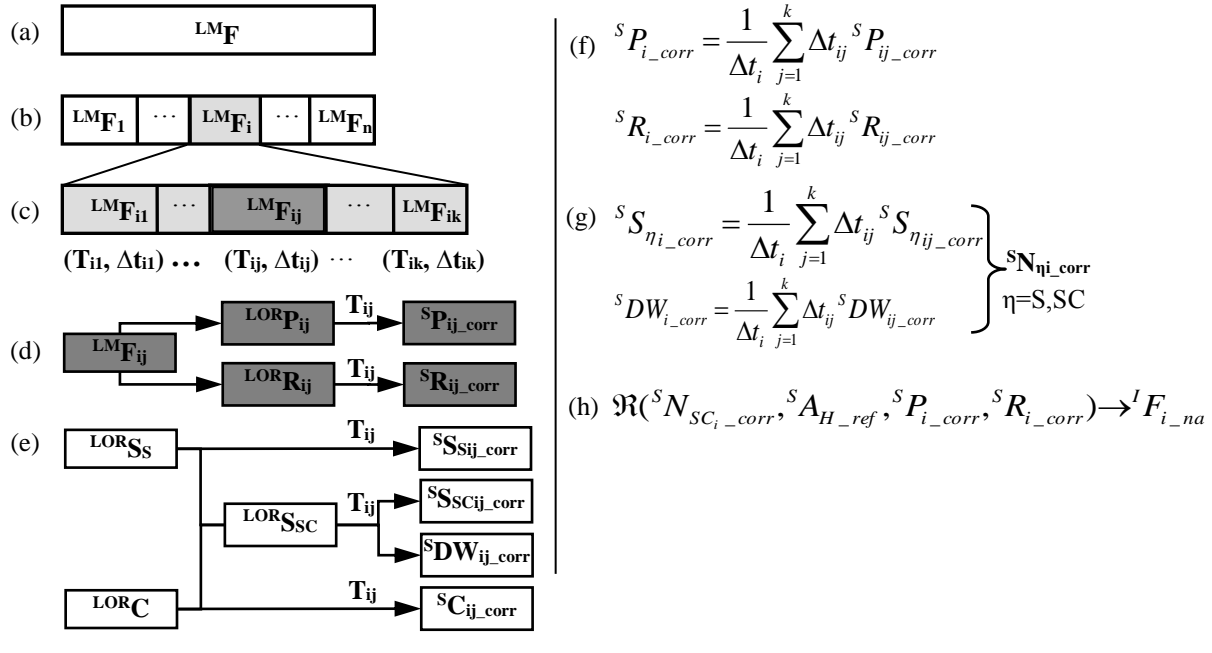
Once the rigid-body motion of the head has been measured it can be introduced into the reconstruction. Motion correction can be performed either after the reconstruction (image-space) or it can be included in the reconstruction (LOR-space). Including the motion correction in the reconstruction is typically preferred for a number of reasons. If the correction is performed in image-space using an iterative algorithm, each subframe must have sufficient counts to be reconstructed; with LOR-based motion correction the temporal resolution of the subframes is only limited to the temporal resolution of the estimates and the temporal resolution of the camera. Furthermore, with LOR-space motion correction a single mu map can be used to reconstruct all the data, while image-space correction requires the mu map be reoriented to each frame prior to reconstruction. Finally, in LOR-space motion correction, the PET data does not need to undergo any additional interpolation after reconstruction. The data sorting algorithm allows for the addition of translation and rotation parameters during the binning of the LOR data. The following subsections will explain the motion correction algorithm

used for the BrainPET and characterize the systematic effects that motion correction introduces through the OP-OSEM reconstruction algorithm.

3.3.2.1 Motion Correction Algorithm

A graphical outline of the algorithm used to correct the PET data for motion can be found in Figure 54. PET data are collected in list-mode format as a series of energy-qualified coincidence events with time-marks recorded every 200 μs . The list-mode file which spans the imaging duration can be divided into frames ($F_i, i=1, n$) of variable duration (Δt) based on the desired dynamic framing protocol. An individual frame of interest can be composed of multiple subframes where each subframe is represented by a single set of motion estimates. The list-mode data from a subframe is then sorted into LORs, discarding the timing information, and then further compressed into a sinogram for fast and efficient processing. The conversion of the LOR data to sinograms is performed using a lookup table that corresponds to the geometry of the scanner. In the event of detected motion, the subject is treated as stationary and the scanner is instead “moved”. As the LOR file format for the BrainPET currently does not support virtual crystals (i.e. the gaps are not represented in the LOR format), reassigning the events to LORs is not possible directly from list-mode as some data would be expected to move into the gaps of the system. Rather this is performed in the conversion to sinograms, creating motion corrected prompt and random sinograms for each subframe. To generate the final prompt and random sinograms for a given frame, a time weighted average of the subframes is performed. Since the normalization accounts for the sensitivity of each LOR pair (a characteristic of the

scanner) and any hardware like the MR coils are fixed to the scanner, for the accurate generation of images they must also be “moved” with the scanner.



- (i) ${}^{LOR}Scat_{low-rez} = SSM({}^IF_{i_na}, {}^IA_{H_ref})$
- (j) ${}^{LOR}Scat_{low-rez} \xrightarrow{T_{ij}^{-1}} {}^{LOR}Scat_{ij_low-rez} \xrightarrow{\text{expand}} {}^SScat_{ij_corr}$
- (k) ${}^SScat_i = \frac{1}{\Delta t_i} \sum_{j=1}^k \Delta t_{ij} {}^SScat_{ij_corr}$
- (l) ${}^SA_{HC_ref} = {}^SA_{H_ref} \times \frac{1}{\Delta t_i} \sum_{j=1}^k \Delta t_{ij} {}^SC_{ij_corr}$
- (m) $\Re({}^SN_{S_i_corr}, {}^SA_{HC_ref}, {}^SP_{i_corr}, {}^SR_{i_corr}, {}^SScat_i) \rightarrow {}^IF_{i_nas}$

Figure 54 Motion Correction Algorithm. F=frame, S=sensitivity, R=randoms, T=transformations, t=time, P=prompts, DW=dwel, N=normalization, Scat=scatter, and A=attenuation. The superscript values refer to the data format with LM=list-mode, LOR=line-of-response, S=sinogram, and I=image.

A result of each frame being broken down into smaller subframes that are individually corrected for motion in LOR space and recombined in sinogram space is that count rate corrections can no longer be applied in image space. The decay and deadtime corrections are applied to the individual prompt and delay coincidence subframe sinograms before they are recombined. The correction of subframes is important as neglecting it may introduce artifacts

in the reconstructed image for frames where the activity changes significantly between subframes. Prior to applying the correction factors to the delay and prompt sinograms, the effects from the dwell must be removed. The dwell accounts for the non-unique mapping of LORs into sinogram space (i.e. multiple LORs, ranging from 2 to 36, are mapped to a single sinogram bin). The dwell can be formatted as a sinogram where the count in each bin represents the number of LORs mapped to that bin. As mentioned in Equation (1.4), the true count rate for the BrainPET is given by,

$$T = [(P - D) \times N - S] \times A \quad (1.4)$$

where T is the true sinogram, P is the prompts, D is the randoms, S is the scatter, A is the attenuation, and N is the normalization. This is of importance here as the algorithm treats the subject as not moving, which in turn means the mu map associated with the subject remains stationary. In the case of simultaneous MR imaging the mu map includes the MR coils, which are fixed relative to the scanner. In the absence of scatter, the true equation reduces to

$$T = (P - D) \times N \times A \quad (3.15)$$

where the correction to the difference between the prompts and randoms sinogram can be considered as the product of the normalization and the attenuation. In this case the coil attenuation can be combined with the normalization and moved with the scanner. To generate the scatter corrected image, the normalization and the attenuation can no longer be considered together due to the scatter term. To address this issue, the coil is considered with the normalization for the generation of the attenuation corrected image and as part of the mu map for the final fully corrected image. In order to accomplish this, the coil attenuation map is converted into its LOR representation, split evenly among all LORs that are mapped to the

sinogram bin. As the LOR format does not have virtual crystals for the gaps, part of the coil attenuation sinogram information is lost; however that data is also set to zero in the normalization and thus ignored.

To generate the normalizations that do and do not include the MR hardware, the average sensitivity of a sinogram bin and the dwell must be accounted for and averaged independently. This separation is performed during the sorting from LORs to sinograms using the transformation estimates. After a time-weighted averaging of the sensitivity and the dwell of the subframes, they can be recombined to form the normalization. The same procedure is followed to generate the normalization that includes the hardware, except that the LOR of the sensitivity is first divided by the coil attenuation in LOR format.

To generate the attenuation corrected image, the motion corrected prompts and randoms sinograms are used. In this case, the product of the sinograms is corrected using the normalization that includes the subject's and MR coil μ map. This attenuation corrected image is used to estimate the scatter profile of the image using a single-scatter model. Briefly, a low-resolution model of the scatter is first generated as a sinogram and converted into LORs; next a higher resolution model of the scatter is generated in LOR space; finally, the scatter is converted back into a sinogram and scaled to fit the tails of the motion corrected PET frame.

As the algorithm predominantly operates in LOR-space rather than sinograms-space, the gap structure is enforced resulting in empty sinogram bins, which may not correspond to the gap structure in the motion corrected prompt and random sinograms. To address this fact a scatter sinogram is estimated for each subframe and combined before it is scaled to fit the tails of the motion corrected data. As each subframe may not have sufficient coincidence events for

an accurate estimate of the scatter, the complete attenuation corrected image is used. In the sorting from sinogram space to LOR space the scanner is moved in the opposite direction as the motion correction, thereby placing the subject in the position occupied in real space during that subframe. The expansion is then performed with the subject in that position. The expanded scatter LORs of the subframes are then converted back into sinograms with the lookup table which accounts for their motion, moving the data back to the corrected position. The time weighted average of the scatter over the subframes is then determined and the data are scaled accordingly, accounting for the MR hardware in the normalization, and the resulting scatter sinogram is used to reconstruct the attenuation and scatter corrected image.

To generate the attenuation sinogram that includes the MR hardware for the final reconstruction of the attenuation and scatter corrected image, the attenuation sinogram of the subject is multiplied by the time weighted average of the subframe coil sinograms. An alternative would be to generate a time weighted average of the coil in image space and then forward project that into an attenuation sinogram; however, that would be inaccurate as the absence of the gap structure will allow parts of the coil to contribute to attenuation effects for LORs of a frame that may have corresponded to gaps in the individual subframe.

3.3.2.2 Characterization of Reconstruction and Systematic Effects of Motion Correction

This section is focused on characterizing the 3D OP-OSEM reconstruction algorithm used on the BrainPET in regards to its operation with motion correction. Affine transformations are briefly discussed as they are the form in which the motion estimates must be presented to the

reconstruction algorithm. Simulations of motion are then performed and the effect of losing LORs as they are placed outside of the field of view is examined for motions up to 1 cm and 10 degrees. Finally the generation of the affine transformation is studied: the effects of differing conventions, how to properly average multiple estimates of motion, and how to filter the estimates. A novel L2-norm based filtering is suggested and compared against a simple median filter of the parameters independently and the limits on the motion magnitude under which the latter are sufficient are discussed.

3.3.2.2.1 *Affine Transformation*

The rigid-body motion of an object in real space can be thought of as a series of rotations and translations or more broadly as a linear transformation. There are a number of mathematical formalisms that have been defined for the representation and properties of the rigid-body rotation of an object: quaternion, axis-angle, and rotation matrices. One advantage of the rotation matrix representation is that it can be represented in an affine transformation matrix, which can account for the translation and rotation of an object, as well as changes in skew and scale.

In the Euclidean space with a Cartesian coordinate system, the translation of a point \mathbf{x} can be represented in vector form by $\mathbf{T}\mathbf{x}$ where

$$\mathbf{T}(\alpha, \beta, \gamma) = \begin{bmatrix} 1 & 0 & 0 & \alpha \\ 0 & 1 & 0 & \beta \\ 0 & 0 & 1 & \gamma \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.16)$$

and where \mathbf{x} is given by the vector of coordinates along the principle axes,

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (3.17)$$

Rotations of a point are performed along three orthogonal axes, the x, y, and z axes.

The rotation transformations are given by \mathbf{R}_x where

$$\begin{aligned} \mathbf{R}_x(\theta) &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) & 0 \\ 0 & \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ \mathbf{R}_y(\varphi) &= \begin{bmatrix} \cos(\varphi) & 0 & \sin(\varphi) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\varphi) & 0 & \cos(\varphi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ \mathbf{R}_z(\phi) &= \begin{bmatrix} \cos(\phi) & -\sin(\phi) & 0 & 0 \\ \sin(\phi) & \cos(\phi) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{aligned} \quad (3.18)$$

Transformations involving a series of rotations are given by the left multiplication of the individual rotation matrices, i.e. $\mathbf{R}=\mathbf{R}_3\mathbf{R}_2\mathbf{R}_1$. The commutative property does not hold for rotations and thus the order in which rotation along the x-, y-, and z-axes are performed is relevant. The standard convention of $\mathbf{R}=\mathbf{R}_x\mathbf{R}_y\mathbf{R}_z$ is used in this work. The coregistration and reconstruction software may present or require the transformation estimates as either the individual angles (e.g. SPM) or the rotation matrix (FLIRT). If the rotation matrix is provided or required then the convention is trivial as the transformation matrix is unique, however if only the angles are presented then the convention must be known to generate the matrix.

3.3.2.2.2 Motion Correction Reconstruction Resolution Effects

As the position of the PET detectors does not actually change, performing the rebinning of the LORs for motion correction can lead to events being discarded as they would be placed in LORs that are not accounted for in the model. Discarding events can lead to a reduction in image contrast and signal to noise ratio. This effect is independent of errors due to inaccurate estimation of the motion fields; however, it is worthwhile to quantify since unlike inaccurate motion estimates, which can be overcome through better registration algorithms or finer modeling, these reconstruction-based errors are inherent to the motion correction algorithm.




Figure 55 MC Reconstruction Effect. To quantify the effect of discarded events, motion was simulated using a high-resolution MPRAGE image (A). To estimate distortions due to reconstruction, the motion corrected images were compared to a version of the original image that had been reconstructed with the sinogram bins corresponding to the gaps set to zero (B).

To estimate the error, a high-resolution MR MEMPRAGE image was used. To account for the effect that the gaps between detectors has on the data, the image was forward projected into sinogram space, the gap structure was applied to the image as a binary mask, and then the image was reconstructed using OSEM. The results are illustrated in Figure 55. To quantify the distortion introduced as a result of discarding counts, patient motion was simulated, motion correction was performed in LOR space, and the resulting motion-corrected image was then

registered to the MPRAGE image which had not been moved, but had been reconstructed with the sinogram gap-structure (Figure 55B). To determine the magnitude of motion, six random numbers were generated between 0 and 1 and multiplied by 10. These numbers were then used as the values for translations (in mm) and rotations (in deg). The results can be found in Figure 56. The noise spike was due to an abnormally noisy voxel that appeared in the reconstructed image (1000 times higher than the maxima of the other replicate images). The six offsets incurred from reconstruction can be found in Table 5 both with the one noisy frame and without (dagger).




Figure 56 MC Reconstruction Simulation. (Left) 100 random sets of translations and rotations on the scale of 0-10 (mm or deg) were applied to the phantom in image space and then corrected in LOR space. The resulting MC images were coregistered to original phantom that had been corrected for the gap structure of the scanner. Of the 100, one result showed poor results from software coregistration, however that was due to a set of voxels outside the head that had a high value due to the OSEM reconstruction. (right) Results of the remaining 99 sets.

Table 5 Mean Offsets for MC Reconstruction Simulation

Replicates	Translations (mm)			Rotations (deg)		
	x-axis	y-axis	z-axis	x-axis	y-axis	z-axis
100	0.1008±0.0760	-0.0234±0.0692	0.0146±0.0680	-0.0824±0.5089	0.0378±0.0373	0.0050±0.1633
99†	0.1064±0.0518	-0.019±0.054	-0.0211±0.0231	-0.0315±0.0463	0.0394±0.0338	0.0109±0.0386

The outlier replicate highlights a shortcoming of the SPM coregistration algorithm that was used and is relied on throughout Masamune. SPM uses a version of the information theory-based approach proposed by Collignon et al. [213] optimized to give a smoother cost

function. Essentially a two dimensional joint histogram of intensities may be produced from a pair of images (e.g. two MR images, an MR and PET image, etc). When the images are perfectly registered, the joint histogram has a specific pattern. If the two images are the same image, this pattern will be a line whose projection into the x-y plane will be of slope one. As the images are brought further out of alignment the pattern becomes more diffuse. This concept is represented mathematically using mutual information, the difference between the joint entropy and the sum of the Shannon-Weiner entropies of the two images. The registration algorithm seeks to maximize the mutual information, effectively maximizing the statistical dependence of the two datasets. It then follows that the accuracy of registration by mutual information can be influenced by the presence of a large volume of background or the presence of noisy voxels outside the object of interest [214]. An alternative registration method would be to use the contours of the anatomy with the principle that form follows function, i.e. boundaries in the functional data will match boundaries in the structural data. While FreeSurfer's boundary-based cost function registration (bregister) was shown to have a more accurate performance than SPM for the registration of structural MR and fMRI [174], it requires an accurate segmentation of the grey/white matter boundary which limits the subject populations to those that FreeSurfer can accurately segment and requires ubiquitous uptake across the individual tissue classes to specify the boundary.

3.3.2.2.3 *Effects of Estimation and Reconstruction Convention Mismatch*

While it may be favorable to determine affine transformations in terms of pure rotations followed by pure translations, there is no standard convention for the order in which the rotations are applied. Changing the order of in which the transformations are applied may

yield different results for the resulting motion. This is of most concern when integrating motion estimates from different sources (e.g. FSL-FLIRT, SPM, custom software) where conventions may be different. To explore the effect that this would lead to in the reconstructed image the 100 randomly calculated motion estimates were used to reconstruct images and register them to the static image as previously described in this section; however, the motion estimates were determined as translations followed by rotations (T*R), but the LOR rebinning followed a separate convention, translations followed by rotations (R*T). The offsets for the improperly reconstructed images can be found in Figure 57. To determine the practical implications that this would have on analysis, a spherical ROI of 1cm radius was created at position 1cm off of each of the principal axes [1cm,1cm,1cm]. The measured errors in coregistration were then applied to the sphere and the percent volume overlap was then determined between the resulting sphere and the original sphere. As a comparison the volume overlap was also calculated for the coregistration errors determined when the transformation convention was consistent throughout, i.e. rotations followed by translations (T*R).

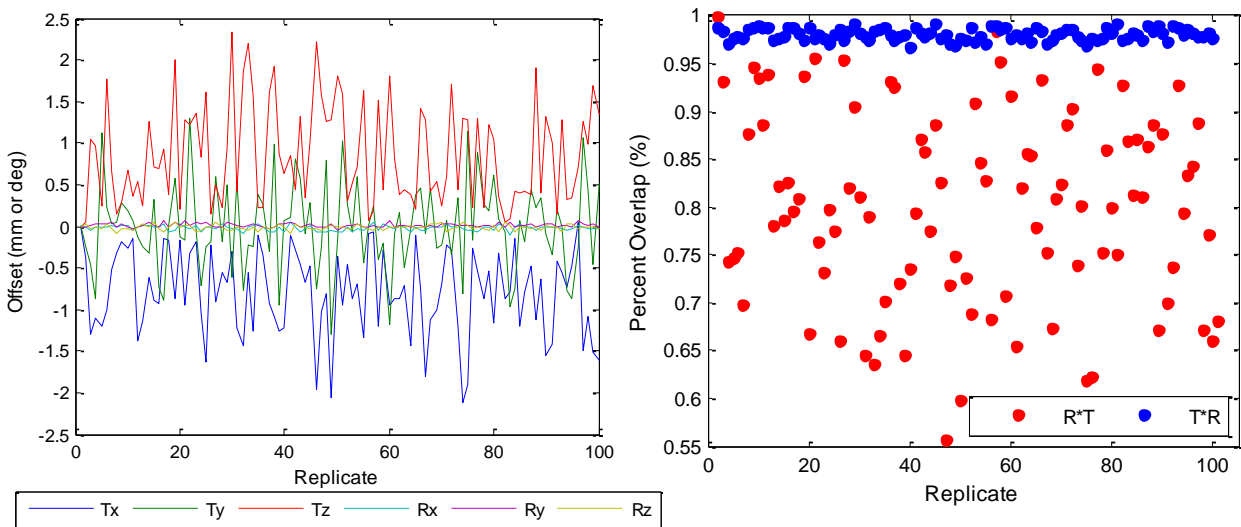


Figure 57 Effects of different affine conventions. The percent volume overlap of a 1cm sphere with center [1cm,1cm,1cm] was found to degrade from $98.09 \pm 0.63\%$ to $80.02 \pm 10.22\%$ when the order of translations and rotations are reversed. Empirically, to achieve the same level in error as a 3 mm spherical ROI with proper orientation, the ROI for the improper orientation would have to be 50 mm in radius.

The mean spherical overlap when the transformations were consistent was found to be $98.09 \pm 0.63\%$ while the overlap when there was an inconsistency was found to be $80.02 \pm 10.22\%$. This significant drop in mean demonstrates that inconsistencies in convention can become a considerable source of bias, especially in small structures. Another aspect is the 17-fold rise in the standard deviation of the overlap. If each motion estimate represented the transformations of a sinogram, then the cumulative effect would be a loss of spatial resolution in the final reconstructed image. The significant variability in overlap resulting from mismatch between the estimation and reconstruction convention illustrate that considerable attention needs to be paid in regards understanding the estimation algorithm (e.g. FLIRT, SPM register) as the mismatch in the images may be difficult to evaluate qualitatively.

3.3.2.3 Averaging Motion Estimates

Several MR and PET-based methods that we have implemented for rigid-body motion estimation were discussed in Section 3.3.1. The temporal frequency of the estimations varies greatly from method to method, from on the order of 20 ms (cloverleaf navigators) to 10 minutes (late frame intraframe PET image). Some of these motion estimates, specifically those collected at higher temporal frequencies, are not individually useful for the PET reconstruction. Determining the temporal frequency that is most appropriate for a particular experiment requires careful consideration of the activity, its distribution, and the desired duration of the image frame to name just a few factors. To this end, it may be desirable to construct a set of average motion estimates over a given period of time; however, each estimate is composed of a series of translations and rotations and must be taken as a set as they are not independent (rather together they define the orientation of the head). An example of such can be seen in Figure 58.




Figure 58 Averaging Motion Estimates. From Left to right: Original object (1), two rotations (2 & 3), proper average of the two rotations using polar decomposition (4) and improper averaging using the arithmetic mean (5), results from taking the average of each rotation separately.

Curtis *et al* showed that the arithmetic mean of a set of a pair of rotation estimates can be performed using either the quaternion representation or rotation matrices [215]. They also note that only the method they derive using rotation matrices can easily be expanded to the

mean of sets greater than two, so that method was used. While Curtis' approach relied on the single value decomposition, a more recent derivation uses the formalism of the polar decomposition [216]. In brief, the arithmetic mean of a set of rotation vectors, $\bar{\mathbf{R}}_F$ is given by

$$\mathbf{R}_{\text{arith}} = \bar{\mathbf{R}}_F \mathbf{S} \quad (3.19)$$

where $\mathbf{R}_{\text{arith}}$ is given by the arithmetic mean of the elements of the individual rotation matrices,

$$\mathbf{R}_{\text{arith}} = \frac{1}{N} \sum_{i=1}^N \mathbf{R}_i \quad (3.20)$$

and \mathbf{S} is a symmetric positive definite matrix. As long as the order in which the orthogonal rotations are applied remains constant, transforming from the estimates to the rotation matrix is trivial. As the affine transformation formulation is rotations followed by translations, the rotations can be averaged first using polar decomposition and then the arithmetic mean of the translations can be computed.

3.3.2.4 Median Filtering and Errors Introduced from Independent

Averaging

While the arithmetic and geometric means of a series of motion estimates have been defined [217], morphological functions like the median are poorly defined as the motion estimates (or at least the rotational element) must be taken as a set. Morphologic functions, like the median, can be useful in signal processing as they can preserve the detail of the signal more than a moving average and they are more robust to outliers. Here we suggest two approaches, one which relies on the L2-norm and one which treats each of the rotations as independent.

3.3.2.4.1 *L2-norm based Median Filtering*

As the median is only defined for data that can be represented as one-dimensional, it can be used on the motion estimates if the three rotations can be represented as a single value. To transform the estimates from a set of three numbers to a single value we use the L2-norm. To begin, we define an arbitrary set of points, for the sake of simplicity a cube that is centered at the origin and of length 2.0 cm was used. The vertices of the cube can then be transformed using the motion estimates and the Euclidian distance between the corresponding original points and the transformed points can be measured. The sum of the displacements can then be used as the single value on which to perform the median filtering. After the median within the kernel is determined, the rotations corresponding to that displacement can be recalled and used as the measurement at that time. One important aspect to consider is that the transformation from rotations to the sum of norms is not unique, as multiple transformations can lead to the same sum of displacements. To account for the degeneracy, all transformations that exist in the kernel that could produce the median value are averaged together using the approach outlined in Section 3.3.2.3. In the case of an even kernel, where the median may fall between two values, all the transformations corresponding to the nearest sum of norms below and above the median can be averaged together.

3.3.2.4.2 *Independent Median Filtering*

The naive approach to taking the arithmetic mean,

$$\bar{\boldsymbol{\theta}} = \left[\frac{1}{N} \sum_{i=1}^N \theta_{xi}, \frac{1}{N} \sum_{i=1}^N \theta_{yi}, \frac{1}{N} \sum_{i=1}^N \theta_{zi} \right] \quad (3.21)$$

is appealing since it can then be argued that since the rotations can be treated as independent, then median filters can be applied independently for the rotations. While this method was shown to be inaccurate in Section 3.3.2.3, there are some situations where this approximation might hold. One argument for this approach is that patient motion is largely stochastic and not continuous. If this is true and the motion estimates are determined frequently enough, then they can be seen as independent measurements of the same position of the head. A center of mass calculation of the estimates would provide a true measurement of the rigid-body rotation, which in turn is the independent averaging of the individual components. Another argument is the practical aspects of motion in the head, and more so when the MR coil is considered. The arithmetic mean of a set of rotation vectors has been shown to be a second-order approximation of the Euclidean mean in the case of small deviations [217], however the point at which the approximation no longer holds, and more so at which point it would lead to detectable errors in PET imaging, must be considered.

A number of metrics exist to determine the effect of motion, and in this case the error in motion estimation methods. For example, the sum of the Euclidian distance of a set of points (e.g. points on a sphere or cube) using one set of motion estimates and the same set of points using another set of motion estimates can be considered. One drawback to this metric are that the number does not have a meaningful interpretation. An alternative metric is to look at the volume overlap between the two objects rather than the Euclidian distances between point pairs. This is equivalent to the Dice coefficient as the two objects have the same shape. There are a number of benefits and drawbacks to such a metric. Firstly, it is very common in PET analysis to analyze the data using a ROI-based approach. In this case it is not the absolute

orientation that is necessarily important, but that the ROI is accurately placed over the structure of interest. Here the percent overlap of the two methods is meaningful; a high overlap means the motion estimation methods provide similar results, and the error is quantifiable in a term related to the biological analysis. Similarly, it is also acceptable in PET analysis to present variations on the simple mean of an ROI, for example studies present the max value an ROI or the mean of a percentage of the ROI (e.g. mean of the top 75% of voxels). As such measuring the net Euclidian distance may be too strict for evaluation of the errors for practical PET imaging.

To determine if the arithmetic mean averaging was valid and set limits on the motion parameters wherein the use of the arithmetic mean of the rotation components approximated the mean through polar decomposition, the degree of error was determined over the space of possible rotations (0-180° rotations along all three axes). Spheres whose centers were displaced 30 mm from the origin along all three principle axes were constructed. The centers of the spheres to be compared were determined by applying an average transformation to the center of the original sphere centered at [30mm, 30mm, 30mm]. The spheres were then constructed and the volume overlap between these two cases was calculated. The average transformations were determined between a vector of no rotations [0°,0°,0°] and the set of rotation vectors, either using the proper method of averaging presented in 3.3.2.3 or by taking the arithmetic mean of the individual rotation components. In this way the results of a vector [10°,10°,10°] would represent the effect of averaging the head at rest with a sudden rotation of 10° around each axis. A volume overlap of 1 would correspond to a set of rotations where the proper mean was equivalent to the arithmetic mean of the rotations. As the volume overlap is dependent on

the radius of the sphere, four different sized spheres were used which corresponded to the volume of four FreeSurfer derived structures ranging from the smallest used to the largest: sulcus intermedius primus (of Jensen) [4mm radius], left choroid plexus [7mm], left caudate [10 mm], and left cerebellum [25mm]. The spherical overlap was calculated for all combinations of rotations that ranged between -90° and 90° along the x-, y-, and z-axes. A contour plot for the case of no y rotation and a constant y rotation of 10° can be found in Figure 59 for the previously mentioned spheres. The results show that there is at least a 90% overlap in all spherical ROIs if the arithmetic mean of the rotations is used rather than the proper average if the rotations are less than 15° .

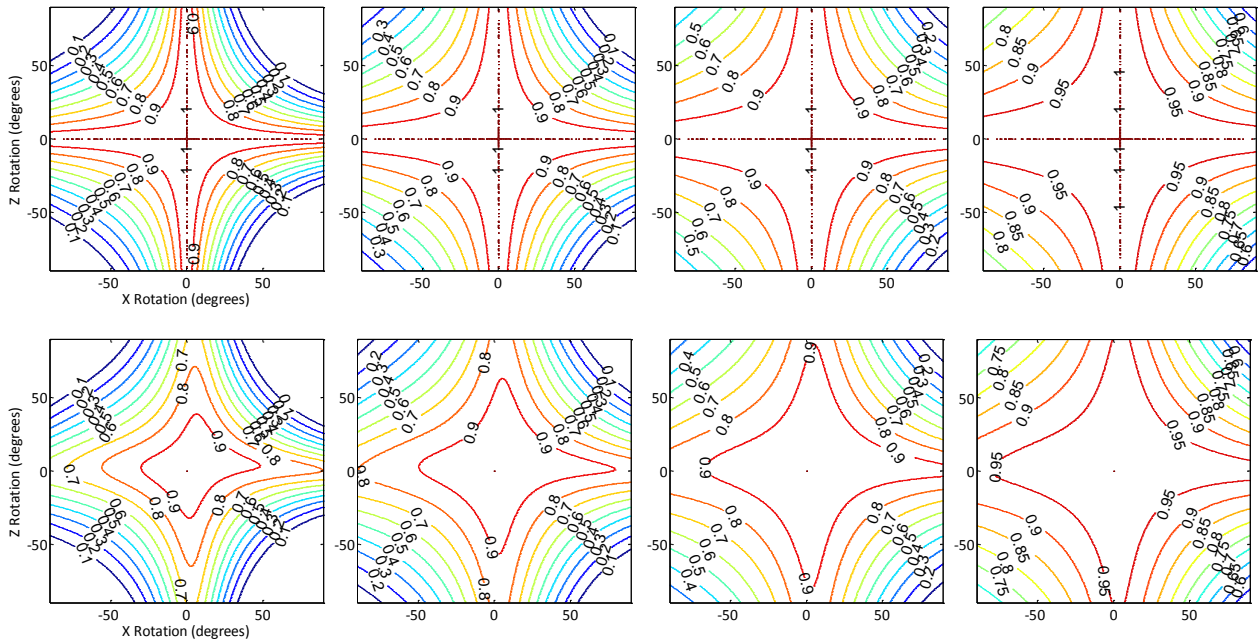


Figure 59 Percent volume overlap between spheres using proper and arithmetic averages. (Top-from left to right) Spheres of 4, 7, 10, and 25 mm radii for fixed y rotation of 0° . (Bottom-from left to right) Same spheres, but for a fixed y rotation of 10° .

Chapter 4 ASSESS THE ADDED VALUE OF IMPROVED PET DATA QUANTIFICATION FOR NEUROLOGICAL APPLICATIONS

In this section, three projects are described to highlight how the integrated MR-BrainPET scanner and the tools proposed in the previous chapters could be used for performing advanced neurological PET/MR studies. The first example was aimed at studying the interference between the two modalities from a physiological point of view. It required static and dynamic PET data processing and the integration of advanced MR-based data analysis tools. In the second example, the effects of combined MR-assisted motion and partial volume effects corrections on the estimation of binding potential were investigated. Finally, the preliminary work toward developing a non-invasive method for deriving the radiotracer input function is presented. This last project benefits from the flexibility provided by the automated PET data processing and reconstruction pipeline and requires all the advanced method for increasing the reliability of the PET estimates.

4.1 Analysis of MR-derived Interference Effects on FDG-PET

Brain Quantification in Human Studies

Simultaneous PET/MR systems are a result of considerable advances in both technologies. While significant work has been done to address the hardware integration and to minimize the electromagnetic interference, there are other aspects that have to be considered in an integrated scanner. For example, the potential “physiologic interference” in a combined system – the influence of one modality on the biological process measured by the other – has not been thoroughly investigated. While the effects of static and varying magnetic fields on the resting brain’s metabolic activity have been studied [218, 219], others sources remain and must be understood for the validation of simultaneous PET/MR as a clinical tool. One potential source, for example, is the MR gradient-induced acoustic noise stimulating the auditory pathways of the brain.

The PET/MR environment is significantly different from the traditional PET environment; for routine clinical FDG-PET studies, it is common practice to isolate the patient in a quiet environment for the initial uptake phase [220]. Because acquiring MR data during this phase would shorten the total PET/MR scan duration, and could directly benefit PET (e.g. motion correction, etc.), it is desirable to administer the radiotracer on the table and perform MR immediately after radiotracer administration. However, the repetitive acoustic noise characteristic to the MR environment can reach as high as 131 dB [221], which causes neuronal stimulation of the primary auditory cortex (PAC), leading to focal increase in blood flow and glucose consumption [222]. Using fMRI techniques, MR acoustic noise has been shown to

directly stimulate the PAC, requiring the development of specialized paradigms to minimize signal artifacts when studying the PAC [223-225]. Unlike fMRI, whose time course is on the order of seconds for the measurement of a stimulus, and thus allows for special data collection paradigms, FDG has an uptake on the order of tens of minutes and is influenced by any changes experienced over this timeframe. As FDG-PET non-specifically images glucose uptake, the acoustic noise could potentially lead to artifacts in the form of focal uptake of FDG. These artifacts could, in turn, significantly impact the FDG-PET metabolic patterns observed in PET studies.

This pilot study sought to characterize any regional changes in FDG uptake resulting from MR acoustic noise and set an upper limit on its effect. As a first step, phantom experiments were performed to assess the hardware-related effects (i.e. electromagnetic interference and MR eddy-current induced heating) on the PET data acquisition using our standard MR protocol. Next, advanced data analysis tools were used to compare the FDG-PET data acquired with and without running MR sequences during the FDG uptake phase in healthy volunteers that underwent two PET/MR imaging sessions. While particular attention was paid to the PAC and associated areas, an explorative analysis of the entire brain was also performed.

4.1.1 Materials and Methods

4.1.1.1 Human Volunteer Studies

4.1.1.1.1 Subject Selection

Ten healthy volunteers (10 males, mean age 38.6 ± 11.2 years) were enrolled in this study. The mean height and weight among subjects was 181.4 ± 5.1 cm and 88.9 ± 22.1 kg. All subjects gave written informed consent and were studied in accordance with a protocol approved by the affiliated Institutional Review Board. Subjects were asked to follow a similar routine for the 24 hours prior to each scan to minimize metabolic changes prior to imaging which included fasting for at least 6 hours prior to the injection. Immediately prior to imaging, subjects' blood glucose levels (BGL) were measured. The mean BGL for the control and stimulation environment was 103.9 ± 6.57 g/dl and 103.9 ± 9.29 g/dl, respectively; mean difference between visits one and two was 0.00 ± 5.79 g/dl.

4.1.1.1.2 Data Acquisition Protocol

Due to the relatively long half-life of FDG (110 minutes), in order to minimize residual uptake, subjects underwent simultaneous FDG-PET/MR imaging on two consecutive days, alternating between an MR "quiet" (control) and "noisy" (stimulation) environment. Prior to PET data collection in either environment, an MR localizer was run to assess subject position in the field of view. Next, after PET data collection was initiated, ~ 185 MBq of FDG was administered intravenously allowing for dynamic imaging. In the control condition (i.e. no MR acoustic noise during the FDG uptake phase), no MR sequences were run for the first 40

minutes of the PET data acquisition. Subsequently, the ME-MPRAGE and DUTE sequences were run for deriving ROIs and the attenuation correction map, respectively. For the auditory stimulation environment, the MR sequences that make up our standard brain tumor imaging protocol were run immediately after FDG injection (i.e. MR acoustic noise was present during the FDG uptake phase). MR imaging was performed using the CP and 8-channel concentric head coils. In this study, we focused on the MR sequences that are part of our standard brain tumor imaging protocol. In addition to the localizer, these included: T2-Sampling Perfection with Application optimized Contrasts using different flip angle Evolution [**T2-SPACE**] (0.9 mm isotropic, TR/TE=3200/487 ms, TA=4:43 min); FLuid Attenuated Inversion Recovery [**FLAIR**] (1.1×0.9×5 mm³, TR/TE/TI=10000/70/2500 ms, FA=150°, TA=3:02 min); **DUTE** (1.67 mm isotropic, TR/TE1/TE2 =200/0.07/2.24 ms, flip angle=10°, TA=3:20 min); Blood Oxygen Level Dependent effect [**BOLD**] (3.4×3.4×8 mm³, TR/TE=2000/19, FA=90°, TA=14:06 min); Time-of-Flight [**TOF**] (0.7×0.5×0.7 mm³, TR/TE=24/3.68 ms, FA=18°, TA=6:16 min); Multi-echo Fast Low Angle Shot [**MEFLASH**] (2 mm isotropic, TR/TE1/TE2/TE3/TE4 =2.46/4.92/7.38/9.84 ms, TA=2:28 min); **T1 Mapping** (1.8 mm isotropic, TR/TE=7.3/4.41 ms, FA=15°, TA=00:13 min); Dynamic Contrast Enhanced imaging [**DCE**] (2.6×1.8×2.1 mm³, TR/TE1/TE2=6.8/2.6/3.89 ms, FA=10°,TA=5:59 min); Diffusion Tensor Imaging [**DTI**] (1.9 mm isotropic, TR/TE=7990/84 ms, TA=6:57); and Multi-Echo Magnetization Prepared Rapid Gradient-Echo [**ME-MPRAGE**] (1 mm isotropic, TR/TI/TE 2530/1200/1.64 ms, flip angle=7°, TA=4:56 min). As per required protocol, standard earplugs with a noise reduction rating of 29 dB were used; however, the acoustic noise of the MR scanner remained clearly audible to the subjects.

4.1.1.1.3 *Data Analysis*

Voxel-wise and ROI-based analyses were performed to identify potential effects of the MR on the FDG uptake in the brain by evaluating:

1. TACs derived from dynamic (4-D) PET data with a framing of 8×60 sec, 2×150 sec, 2×180 sec, and 8×300 sec;
2. LCMRGlC parametric images derived from the first 40 minutes of dynamic data using the method presented by Wu [226] [227];
3. static images reconstructed from the data acquired 40-60 minutes post-injection (similar to a static FDG-PET scan). SUVs corrected for lean body mass (SUV_{lbm}) were calculated in each case.

In addition to voxel-wise statistical analysis using SPM8 [228], FreeSurfer [229] was used to segment ROIs corresponding to cortical and subcortical structures. These ROIs were used for the PET data analysis as described in the following subsections. Global whole brain, grey and white matter mean values were also computed. A diagram of the experimental design, processing, and subsequent analysis can be found in Figure 60.

4.1.1.1.3.1 Time Activity Curves Analysis

For analysis of the PET dynamic data, TACs were defined as the mean within the ROI at each time point with the standard deviation defined as the inter-subject standard deviation. To determine the effect of gradient acoustic noise on the PAC (superior transverse temporal gyrus) the left and right PACs were combined to yield a single ROI. For each time point across subjects, the brain regions of interest were normalized to whole brain and then averaged over

the subject population. The mean uptake for the control and stimulation environment for each time point were then compared to one another using a paired, two-sample Student's t-test.

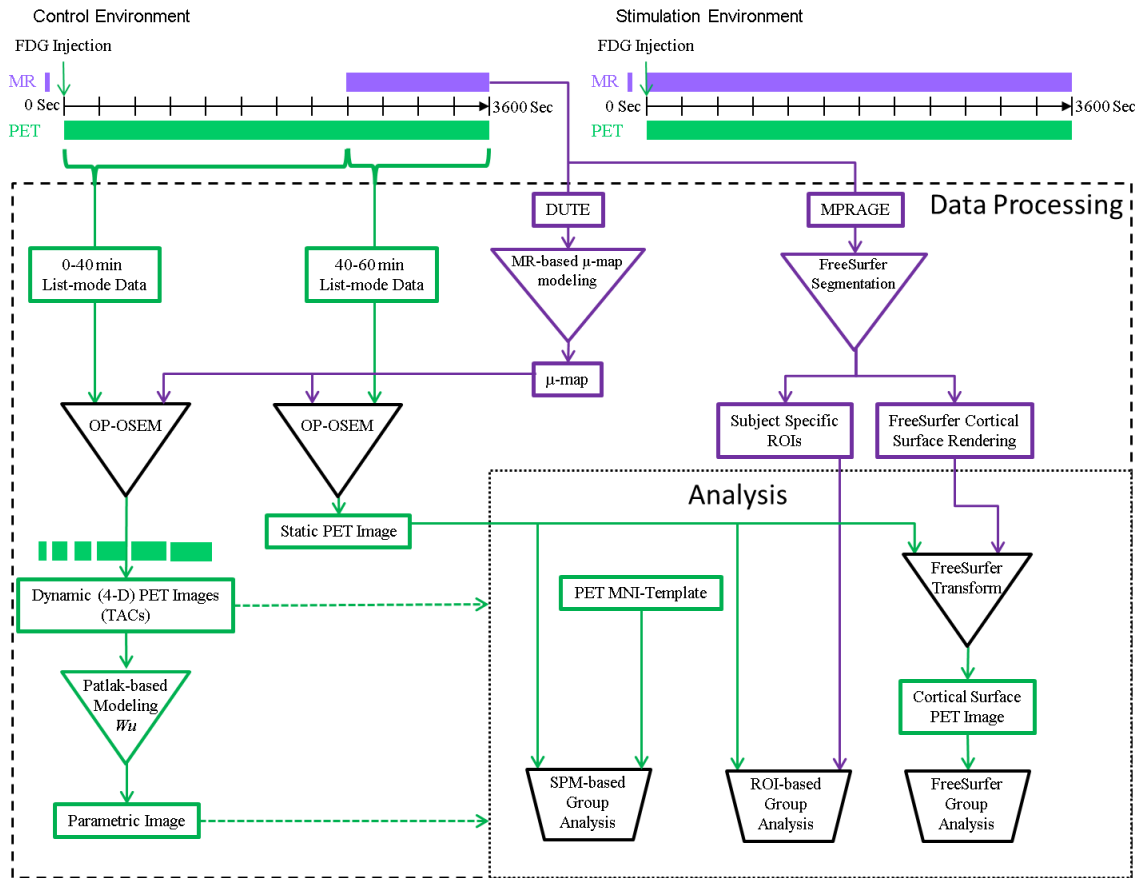


Figure 60 Design of paradigm, processing, and analysis for simultaneous FDG-PET/MR experiment to illicit MR-based physiologic effects on FDG-PET uptake. Subjects were scanned on two separate days in either a quiet control environment or a stimulation environment, where MR acoustic noise was present during the FDG uptake phase. The data processing steps were identical for both control and stimulation environments, but for simplicity they are only shown for the control environment. For the same reason, the analysis is only shown for the Static PET image. Squares represent images, where green are PET and purple are MR. Triangles represent processes (e.g. PET reconstruction using OP-OSEM) where: green means only PET data was used in that process; purple means only MR data is used; and black where both MR and PET data were used as inputs. The dotted arrows connecting the Parametric Image to the Analysis is meant to show that it was processed identically to the Static PET image. Only ROI-based analysis was performed on the TACs.

4.1.1.1.3.2 SPM-based Analysis

As SPM8 does not provide subject specific definition of cortical and subcortical structures, group analysis requires that the subjects be first registered to a common space. For the SPM8 analysis, the common space to which the images were normalized to was the MNI

PET contrast template (MNI-normalized). To derive the transformation fields to the common space, a smoothed (3 mm isotropic Gaussian kernel) version of the individual PET images was fit using an 12-parameter affine transformation followed by a non-linear deformation using SPM8 and the normalized image was cropped to a volume of 79×95×68 voxels measuring 2 mm isotropic. Prior to statistical analysis, the MNI-normalized PET images were smoothed (3 mm isotropic Gaussian kernel) using SPM8. A paired t-test was performed on the data with a cluster size requirement of greater than 20 voxels and both an uncorrected $p < 0.001$ and a familywise-error (FEW) of 0.05 for significance. Anatomic correlation of the PAC, defined as Herschel's gyrus for SPM8 based analyses, and more general cluster location was determined using a Talairach Daemon atlas with the xjView toolbox for SPM8 (<http://www.alivelearn.net/xjview8/>) and is described in detail elsewhere [230, 231].

4.1.1.1.3.3 FreeSurfer ROI-based Analysis

In the FreeSurfer output, each voxel is given an integer label corresponding to the segmented region [135]. For all FreeSurfer based analyses, the PAC was determined as the transverse temporal gyrus from the cortical parcellation. Binary masks were generated for each segmented region using Matlab and the voxel addresses of the ROIs were determined and used to create a lookup table which was then applied to the PET image. To account for the differences in spatial resolution and the spatial mismatch between the two scanners, a downsampling operation and an affine transformation were applied to the segmentation maps using nearest-neighbor interpolation, and the subsequent analysis was performed in the PET space. To account for interscan motion, the skull-stripped MPRAGE image was coregistered to

the PET image using a mutual information algorithm to derive any additional affine transformations, which in turn were applied to the segmentation map.

For analysis of the parametric and static images, paired two-sample Student's t-test between the control and stimulation environments were performed across subjects for each cortical and subcortical brain region segmented in FreeSurfer. To normalize the static images across subjects, the regions were divided by the whole brain activity (WB-normalized). As the goal of this work is to identify any potential regions of significant change, to be conservative, no Bonferroni corrections were included in determining statistical significance ($p < 0.05$) across the brain regions.

4.1.1.1.3.4 FreeSurfer GLM-based Analysis

For group analysis, FreeSurfer uses a space obtained from "inflating" the ME-MPRAGE-derived brain of the individual subject. The same transformations were applied to the PET data and a surface was derived where each vertex value was the mean value along the normal path through the cortical ribbon. The subject specific changes were determined by subtracting the stimulation and control PET images and a study averaged image was derived by taking the mean change across subjects. The resulting image was smoothed (10 mm isotropic Gaussian kernel) and a generalized linear model (GLM) analysis was performed. For conservative evaluation, data were visualized with a $p < 0.01$ which is equivalent to the default threshold of $-\log(p) = 2$ used in FreeSurfer and FSL.

4.1.2 Results

4.1.2.1 Human Volunteer Studies




Figure 61 Creation of high-resolution ROIs for PET imaging . The high-resolution ME-MPRAGE (A) is used to generate a labeled mask volume (B) with FreeSurfer, both of which were subsequently resampled to the PET geometry. The labeled mask was then be used as regions of interest for PET analysis and applied to the 4-D PET images to derive TACs as well as both the parametric [0-40 minute] images (C), and static [40-60 minute] images (D). Unlike the static images, where MPRAGE and UTE were acquired during PET acquisition in a similar manner in both the stimulation and control environments, no MR sequences were run in the control environment during the timeframe used to generate the parametric images.

4.1.2.1.1 Time Activity Curves

A representative subject's FreeSurfer segmentation, along with the ME-MPRAGE image, can be found in Figure 61. The combined left and right PAC ROI had a mean volumetric change of 0.157 ± 0.196 ml between the two visits ($p= 0.21$). A representative subject's TACs for the PAC, motor cortex (M1), and white matter for the control and stimulation environments can be found in Figure 62. As previously mentioned, for TAC analysis, the PAC and white matter time points were normalized to whole brain and averaged across patients to create a group TAC, which can be found in Figure 62 along with the whole brain TACs. A paired t-tests for each cross-subject time point yielded no significant points ($p<0.05$) for either whole brain, white matter, or the PAC for the nonWB-normalized images; however, a paired t-test of the WB-

normalized PAC showed significant differences at all time points after 270 seconds.

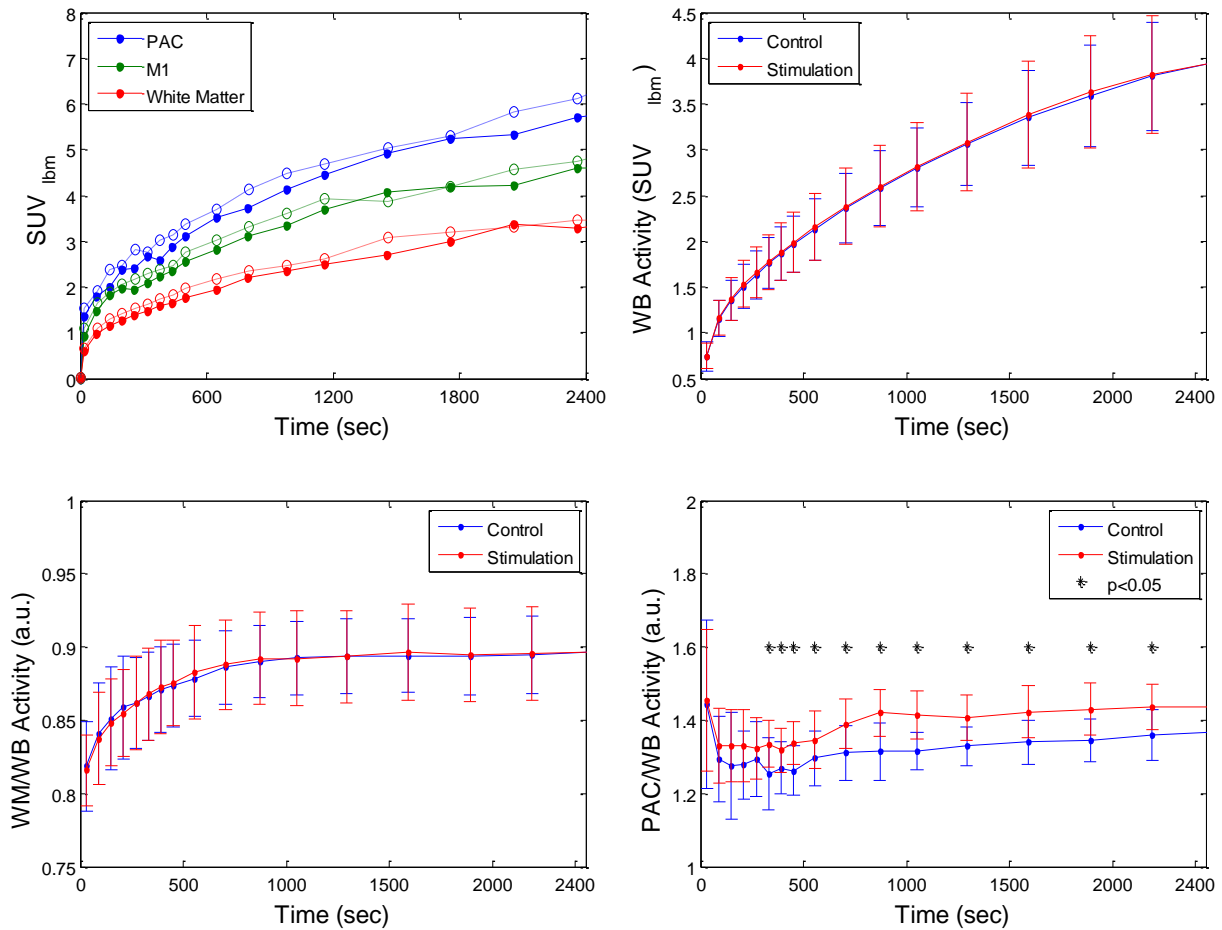


Figure 62 TACs. (A) Representative subject's TAC for both the control (solid line) and stimulation (hashed line) environments. In addition to the PAC (blue), the M1 (green) and white matter (red) are shown for reference. Solid line represents the control environment (no MR performed during first 40 minutes post injection) while the dashed lines represent stimulation environment. (B) Whole Brain TAC; (C) White Matter normalized to WB; (D) PAC TAC normalized to white matter. Each point represents the average of the four subjects with error bars representing one standard deviation (* represents $p < 0.05$).

4.1.2.1.2 Parametric Images

The mean uptake in the whole brain for the stimulation and control cases was $22.993 \pm 0.005 \mu\text{mol}/\text{min}/100\text{g}$ and $22.997 \pm 0.005 \mu\text{mol}/\text{min}/100\text{g}$, respectively with a mean change of $-0.004 \pm 0.006 \mu\text{mol}/\text{min}/100\text{g}$ ($p=0.0643$). The mean uptake in the gray matter for the stimulation and control cases was $26.577 \pm 0.492 \mu\text{mol}/\text{min}/100\text{g}$ and 26.515 ± 0.425

$\mu\text{mol}/\text{min}/100\text{g}$, respectively, with a mean change of $0.0625 \pm 0.4182 \mu\text{mol}/\text{min}/100\text{g}$ ($p=0.648$).

The mean uptake in the white matter for the stimulation and control cases was $20.896 \pm 0.766 \mu\text{mol}/\text{min}/100\text{g}$ and $20.877 \pm 0.615 \mu\text{mol}/\text{min}/100\text{g}$, respectively, with a mean change of $0.019 \pm 0.500 \mu\text{mol}/\text{min}/100\text{g}$ ($p=0.907$). When tested with a paired t-test, none of these regions were found to have $p\text{-values} < 0.05$.

4.1.2.1.2.1 SPM-based Analysis

SPM8 analysis using a FWE of 0.05 and minimum cluster size of 20 voxels yielded no significant activation clusters between the stimulation and control environments. SPM analysis using a $p < 0.001$ resolved focal bilateral activation clusters which extended into Herschel's gyrus as defined by an MNI atlas and can be found in Figure 63. Both regions were found to have increased uptake in the stimulation as compared to the control environment. No regions with significantly decreased uptake were detected. The left and right cluster sizes were 45 and 41 voxels with a change between stimulation and control of $2.46 \pm 0.51 \mu\text{mol}/\text{min}/100\text{g}$ ($8.37 \pm 1.55\%$) and $2.66 \pm 0.55 \mu\text{mol}/\text{min}/100\text{g}$ ($8.70 \pm 1.17\%$), respectively.

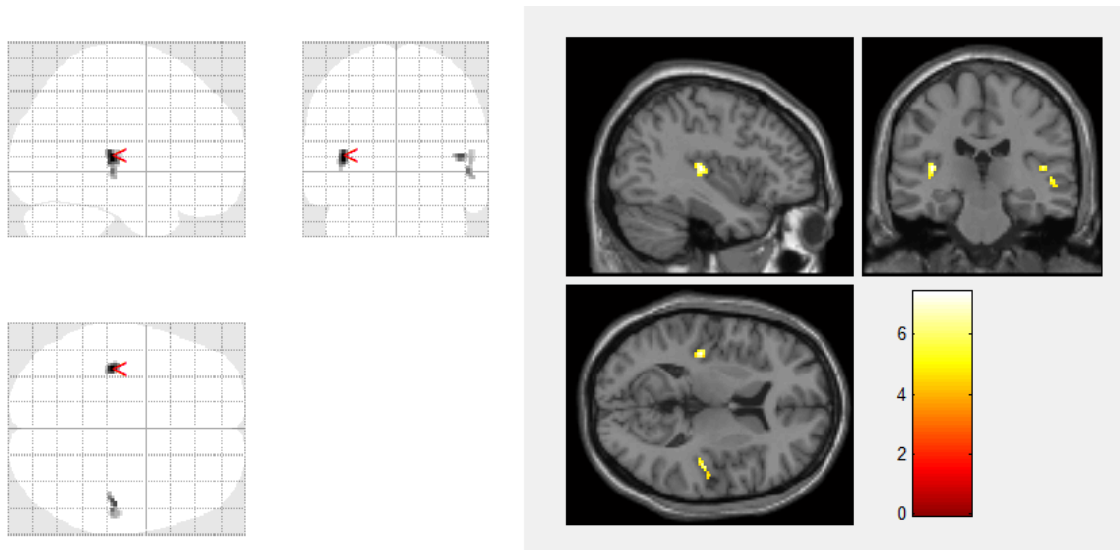


Figure 63 SPM8 Glass Brain of parametric image . Cluster analysis performed with $p < 0.001$ (uncorrected) and minimum cluster size of 20 voxels (left). Activation plotted on three orthogonal slices of single subject's MRI from MNI database (right).

4.1.2.1.2.2 FreeSurfer ROI-based Analysis

The parametric images showed a bilateral increase in PAC LCMRglu in six subjects with opposing changes in the left and right PAC uptake in the remaining four. The mean LCMRglu of the control environment for the left and right PAC glucose uptake was found to be $31.631 \pm 1.895 \mu\text{mol}/\text{min}/100\text{g}$ and $31.734 \pm 1.263 \mu\text{mol}/\text{min}/100\text{g}$, respectively, and $32.952 \pm 2.642 \mu\text{mol}/\text{min}/100\text{g}$ and $34.615 \pm 1.683 \mu\text{mol}/\text{min}/100\text{g}$ for the stimulation environment. No statistically significant change was found in the left PAC ($1.321 \pm 2.927 \mu\text{mol}/\text{min}/100\text{g}$, $4.42 \pm 9.27\%$; $p = 0.1872$), however a significant increase in the right PAC was noted ($2.88 \pm 1.46 \mu\text{mol}/\text{min}/100\text{g}$, $9.13 \pm 4.73\%$; $p = 0.0002$). When the left and right PAC were treated as a single ROI the mean change was found to be $2.10 \pm 1.52 \mu\text{mol}/\text{min}/100\text{g}$ ($6.67 \pm 4.91\%$; $p = 0.0018$).

Statistically significant increases in LCMRglu were found in five other regions: the long gyrus of the left insula ($0.79 \pm 0.89 \mu\text{mol}/\text{min}/100\text{g}$, $3.48 \pm 4.08\%$; $p = 0.0204$), the left superior

temporal plane ($1.19 \pm 1.16 \mu\text{mol}/\text{min}/100\text{g}$, $4.52 \pm 4.33\%$; $p=0.0097$), the left middle temporal gyrus ($0.70 \pm 0.94 \mu\text{mol}/\text{min}/100\text{g}$, $3.14 \pm 4.10\%$; $p=0.0429$), the right superior temporal plane ($1.38 \pm 1.66 \mu\text{mol}/\text{min}/100\text{g}$, $4.94 \pm 5.73\%$; $p=0.0279$), and the right transverse temporal sulci ($2.92 \pm 3.52 \mu\text{mol}/\text{min}/100\text{g}$, $10.03 \pm 12.32\%$; $p=0.0277$). A statistically significant decrease in the LCMRglu was found in the right choroid plexus ($-1.01 \pm 1.27 \mu\text{mol}/\text{min}/100\text{g}$, $-6.66 \pm 8.09\%$; $p=0.0336$).

4.1.2.1.2.3 FreeSurfer GLM-based Analysis

The results of the FreeSurfer GLM-based analysis can be found in Figure 64. Surface-based analysis displayed increased metabolism in the left hemisphere's frontal inferior triangular gyrus, superior temporal lateral gyrus extending into the temporal superior plane, and inferior circular sulci of the insula extending into the superior transverse temporal gyrus. Additionally, in the left hemisphere small foci of decreased activation were observed in the postcentral sulci, occipital middle gyrus, and inferior angular parietal gyrus. The right hemisphere showed increased uptake in the superior transverse temporal gyrus, superior transverse temporal sulci, temporal superior plane, opercular part of the inferior frontal gyrus, as well as the H-shaped orbital sulcus and orbital gyrus. The right hemisphere showed a small focus of decreased uptake in the inferior angular parietal gyrus.




Figure 64 Cortical FreeSurfer analysis of parametric image . Segmented cortical regions are shown in color. Activation clusters are shown in a hot and cold color scheme corresponding to Z score reflecting statistical increase or decrease of FDG compared to no change in uptake, respectively. P-values<0.01 are shown.

4.1.2.1.3 *Static Image*

The mean uptake in the whole brain for the stimulation and control cases was 4.127 ± 0.823 S_{lbm} and 4.149 ± 0.624 SUV_{lbm} respectively with a mean change of -0.022 ± 0.856 SUV_{lbm} ($p=0.937$). The mean uptake in the gray matter for the stimulation and control cases was 4.870 ± 0.954 SUV_{lbm} and 4.90 ± 0.758 SUV_{lbm} respectively with a mean change of -0.0357 ± 0.980 SUV_{lbm} ($p=0.911$). The mean uptake in the white matter for the stimulation and control cases was 3.617 ± 0.650 SUV_{lbm} and 3.636 ± 0.548 SUV_{lbm} respectively with a mean change of -0.019 ± 0.700 SUV_{lbm} ($p=0.933$). The WB-normalized mean uptake in the gray matter for the stimulation and control cases was 1.180 ± 0.023 and 1.18 ± 0.021 respectively with a mean change of -0.001 ± 0.011 ($p=0.780$). The WB-normalized mean uptake in the white matter for the stimulation and control cases was 0.880 ± 0.037 and 0.877 ± 0.004 respectively with a mean change of 0.004 ± 0.017 ($p=0.476$).

4.1.2.1.3.1 SPM-based Analysis

SPM8 analysis using a FWE of 0.05 and minimum cluster size of 20 voxels yielded no significant activation clusters between the stimulation and control environments. SPM analysis using a $p < 0.001$, which can be found in Figure 65, resolved four distinct clusters: one within the left Heschel's gyrus, one in the right superior temporal gyrus on the border of Brodman's 22 and 42, and two in the frontal white matter. The mean change in the left Heschel's gyrus was 0.100 ± 0.041 ($8.68 \pm 3.89\%$) and 0.082 ± 0.024 ($6.75 \pm 1.94\%$) in the superior temporal gyrus.

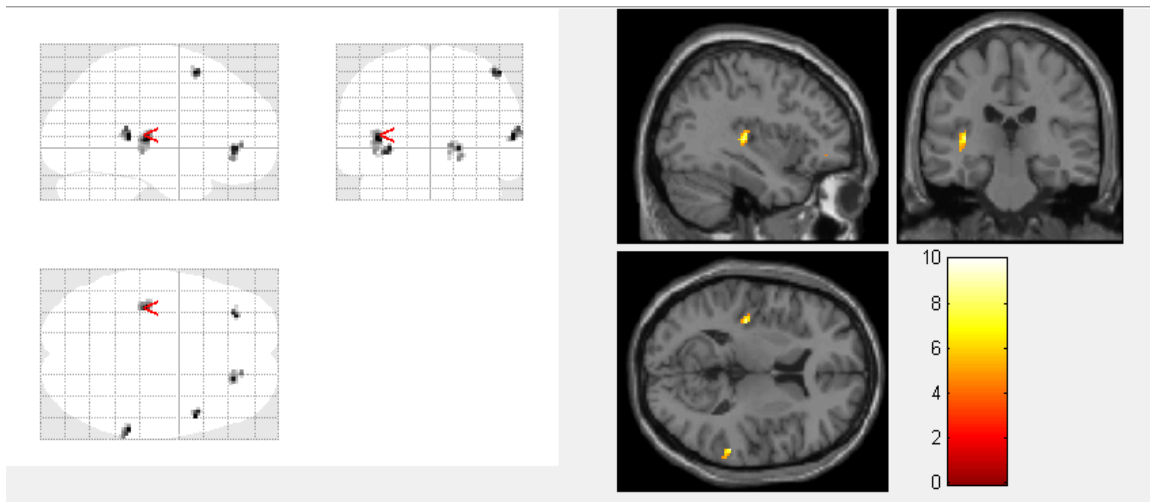


Figure 65 SPM8 Glass Brain of static image . Cluster analysis performed with $p < 0.001$ (uncorrected) and minimum cluster size of 20 voxels (left). Activation plotted on three orthogonal slices of single subject's MRI from MNI database (right).

4.1.2.1.3.2 FreeSurfer ROI-based Analysis

The analysis of the static frames revealed a bilateral increase in PAC SUV_{lbm} in four subjects, bilateral decrease in one, and opposing changes in the left and right PAC in five subjects. The mean FDG uptake for the left and right PAC glucose uptake was found to be 5.660 ± 0.953 SUV_{lbm} and 5.723 ± 0.837 SUV_{lbm} , respectively for the control environment, and 5.766 ± 1.048 SUV_{lbm} 5.931 ± 1.187 SUV_{lbm} for the stimulation environment. After normalizing the

subjects to their whole-brain activity, seven subjects showed bilateral increases in FDG uptake while the remaining three showed a decrease in the left PAC. The WB-normalized mean uptake or the left and right PAC glucose uptake was found to be 1.363 ± 0.084 and 1.381 ± 0.063 , respectively for the control environment, and 1.40 ± 0.083 1.438 ± 0.065 for the stimulation environment. The change between the stimulation and control environments for the left and right PACs was found to be 0.0398 ± 0.969 ($3.17\% \pm 7.28\%$; $p=0.202$) and 0.0572 ± 0.0397 ($4.18 \pm 2.87\%$; $p=0.001$), respectively. When the left and right PAC were treated as a single ROI the mean change was found to be 0.0485 ± 0.0437 ($3.61 \pm 3.34\%$; $p=0.007$).

The remaining 103 cortical and subcortical regions of the brain were evaluated across subjects using a paired, two-sample Student's t-test between the control and stimulation environments. Statistically significant regional changes in the normalized brains were found in the midposterior of the corpus callosum (-0.020 ± 0.027 , $-4.00 \pm 5.23\%$; $p=0.0422$), the left long gyrus and central sulcus of the insula (0.019 ± 0.021 , $1.94 \pm 2.10\%$; $p=0.0186$), the left temporal superior plane (0.032 ± 0.036 , $2.80 \pm 3.00\%$; $p=0.0188$), left H-shaped orbital sulcus and orbital gyrus (0.028 ± 0.034 , $2.11 \pm 2.53\%$; $p=0.027$), right choroid plexus (-0.033 ± 0.038 , $-5.21 \pm 5.78\%$; $p=0.022$), right opercular part of the inferior frontal gyrus (0.024 ± 0.019 , $1.97 \pm 1.69\%$ $p=0.003$), and the right superior transverse temporal gyrus (0.057 ± 0.040 , $4.18 \pm 2.87\%$; $p=0.001$).

4.1.2.1.3.3 FreeSurfer GLM-based Analysis

The results of the FreeSurfer GLM-based analysis can be found in Figure 66. Surface-based analysis displayed increased metabolism in the left hemisphere's front inferior triangular gyrus, superior temporal lateral gyrus extending into the temporal superior plane, and inferior circular sulci of the insula extending into the superior transverse temporal gyrus. Additionally

in the left hemisphere small foci of decreased activation were observed in the occipital middle gyrus. The right hemisphere showed activation in the superior transverse temporal gyrus, superior transverse temporal sulci, the temporal superior plane, opercular part of the inferior frontal gyrus, as well as the inferior angular parietal gyrus. Foci of decreased uptake were noted in the middle occipital gyrus and the superior parietal gyrus.




Figure 66 Cortical FreeSurfer analysis of static image . Segmented cortical regions are shown in color. Activation clusters are shown in a hot and cold color scheme corresponding to Z score reflecting statistical increase or decrease of FDG compared to no change in uptake, respectively. P-values<0.01 are shown.

4.1.3 Discussion

The primary purpose of this study was to explore MR physiologic interference on FDG uptake. Although likely not an issue for most applications, the observed MR-induced increase in FDG uptake in the PAC could have implications for certain studies. For example, care should be taken when evaluating progression in lesions that lie near the auditory cortex if MR sequences are run during the uptake phase of the FDG. It is worth noting that both the static and parametric images showed a statically significant decrease in the right choroid plexus uptake. While the change in FDG uptake of the choroid plexus is likely an artifact due to its small size, which makes it susceptible to artifacts arising from motion, studies in rats has reported that

intense noise can lead to cell damage [232]; however, given the limited body of existing work concerning this potential phenomena, a more extensive study is necessary before a correlation should be suggested.

The unintended physiological effects of combining other medical devices with PET have been previously explored. For instance, transcranial magnetic stimulation (rTMS) has been combined with PET to explore *in vivo* brain connectivity [233-235]. Siebner et al. found that repetitive rTMS of the sensorimotor hand area (SM1) lead to increased FDG uptake in the PAC as a result of acoustic noise produced by the rTMS hardware [236]. Stimulating the left SM1, they noted a change in the left and right PAC of 7.2% and 6.6%, respectively using an ROI-based analysis. While the changes reported here of $8.37 \pm 1.55\%$ and $8.70 \pm 1.17\%$ are slightly larger, they are consistent with those presented by Siebner *et al.* especially considering the higher resolution of the BrainPET compared to the Siemens 951 R/31 PET scanner and the smaller smoothing kernel used in this study. Furthermore, Siebner *et al.* used 2 cm circular ROIs centered on the peak activation in regions using registered T1-weighted MR images for anatomic landmarks [237].

Direct analysis of the TACs, the parametric 0-40 minute image, and the static 40-60 minute image illustrate three different situations which are of importance for evaluating potential physiological interference.

First, by distinguishing significant changes in uptake on a frame-by-frame basis, the TAC analysis is sensitive to specific MR sequences. The dynamics of FDG uptake in unstimulated brain tissue has a predictable time course; any deflections from its smooth projected trajectory would suggest prior stimulation and spikes would suggest possible RF interference. The

absence of spikes in the whole brain activity TAC (Figure 62A) and WB-normalized white matter (Figure 62B), as determined by serial t-tests, suggests lack of RF interference (at least for the framing and reconstruction used in our study). This in turn suggests that any changes in FDG uptake are linked to physiologic changes induced from sources other than hardware interference. This result is further confirmed by the ROI analysis of PAC TAC where a statistically significant change between the control and stimulation environment is maintained across subsequent time points.

Second, the parametric images, in addition to being sensitive to the integral of the activity and accounting for the FDG transport and trapping components, allow for the detection of the cumulative effects of noise in the uptake phase. In this study we used a blood-free approach to estimate LCMRGlu using the method presented by Wu [226]. An appealing aspect of this method is that the whole brain value of each patient is normalized to a constant, which has been used in other group studies on the effects of MR on LCMRGlu [218, 219]. Similarly, we only used the first 40 minutes rather than 60. Monden et al. demonstrated that using a shorter duration could lead to an increase in K_i on the order of 3-5% [227]; however, in this study we were interested in the relative change between the control and stimulation environments so the error derived from this reduced duration should be present in both control and stimulation datasets.

The SPM8/MNI-normalized parametric images showed a bilateral statistically significant increase in LCMRGlu within PAC suggesting that this region was being activated over the first 40 minutes in the stimulation case relative to the control. Activation of the PAC is also supported by the FreeSurfer ROI-based analysis of the LCMRGlu image, where a statistically significant

increase was observed in an ROI spanning both the left and right PACs. When taken separately, the left PAC did not show significance; however, upon closer observation of the FreeSurfer GLM image, this could be due to the small size of the region with significant uptake with respect to the ROI. In addition to a statistically significant increase in LCMRGlu in the right PAC, the ROI analysis showed statistically significant activation in other regions known to have involvement in auditory stimulation: the left and right superior temporal planes, right transverse temporal sulci, and left middle temporal gyrus [238].

Third, the semi-quantitative static analysis is most likely to be used clinically with the early adopters of simultaneous PET/MR, and thus any interference as a result of adopting our protocol (i.e. administer the FDG in the scanner while running MR sequences) should be determined. In this study the simultaneously acquired MR sequences during the static frame (i.e. minutes 40-60 post injection) were the same in both the control and stimulation environments; as a result, differences between the static images of the stimulation and control environments should only reflect changes in uptake that occurred between injection and minute 40. The static images showed a glucose uptake pattern which had some similarities to the parametric images – increased uptake in the PAC in all three analytic methods (SPM8, FreeSurfer ROI-, FreeSurfer GLM-based). An ROI spanning the combined left and right PAC showed a significant increase in FDG uptake, however only the right PAC was found to be significantly different between the control and stimulation environments. The GLM-based analysis explains this difference – a region of increased uptake is apparent just anterior to the PAC. While this activation may be physiologic, it could also be a result of stochastic errors.

The two FreeSurfer-based methods used for analyzing the PET data in this work provided complementary information concerning activation. The minimum requirement for significance in the ROI-based analysis is that there must be either a small focal change (or a number of smaller foci) of rather significant intensity (when compared to the ROI size), or a less intense change distributed over a large portion of the ROI. The ROI-based method places no requirements on the connectedness of the regions of increased or decreased activation within an ROI, rather it is sensitive to the average change in the region likely detecting whether or not the PAC has been activated. The FreeSurfer GLM-based analysis, on the other hand, focuses on clusters of activation allowing them to traverse multiple brain regions. This method may be preferred to detect focal changes within a brain region, for example tonotopic activation [239, 240].

Aside from evaluating the effects of MR acoustic noise on cerebral FDG uptake, we have shown that FreeSurfer, in comparison to SPM8, can provide statistical descriptions of cerebral FDG uptake. In the case of non-simultaneous acquisition, where spatial registration of the two modalities may be inaccurate, the more conservative analytic technique is to transform the PET data directly to a conformed space. Accurate transformation to this conformed space usually requires that the subject data be significantly smoothed as it must have smoothness similar to that of the template. Smoothing the data can obscure small and/or less intense clusters of uptake. In an integrated system, where the spatial transformation between the MR and PET images is known, the MR data can provide all the necessary transformations to a conformed space reducing the need for additional smoothing of the PET data. In FreeSurfer, where this conformed space is derived directly from the subject specific brain, there is a better inter-

subject alignment of cortical regions which can improve the localization or contrast of statistically significant changes in the brain [136]. One example of this improvement in contrast is the conservation of the uptake patterns between the parametric and static images. As the latter occurred over a time period where the glucose uptake is small relative to earlier time points it should have a metabolic pattern which resembles the parametric image which is sensitive to the 0-40 minute uptake phase. Comparing the similarities of the metabolic patterns between the static and parametric frames provides a metric for evaluating SPM- and FreeSurfer-based analyses. The PAC cluster in the right hemisphere is not conserved between the parametric and static images using SPM8; however, they, in addition to the clusters in the temporal superior plane are conserved in the FreeSurfer-based analysis. Another example is the conservation of activation of the left opercular part of the inferior frontal gyrus which is associated with Broca's Area.

FreeSurfer does not entirely replace the need for SPM-based analyses. A drawback to FreeSurfer is the processing time required to segment and inflate the brain. While the transformation to the MNI space with SPM8 can be computed in minutes, processing of the MR data can take on the order of 20 hours with FreeSurfer making SPM appealing for realtime clinical targeted applications [241]. For studies where the activation area is expected to be larger or where ROIs can be determined prior to analysis, SPM8 can provide faster results. Additionally, FreeSurfer does not segment some of the brainstem structures, like the inferior colliculus, a region which has been shown to play a role in auditory processing [242, 243].

One aspect which was not controlled in this experiment was the placement of the earplugs. The placement was not standardized, subjects were given the option to place the ear

plugs in themselves or have a technologist place them for them, which is our center's standard protocol for hearing protection in the MR environment. Subjects were asked for confirmation that the earplugs were adequately placed and whether or not they were equal bilaterally prior to initiating MR scanning. To ensure maximum efficacy from the earplugs, we could have tested the subjects' hearing, however this would bias our results as this is not done routinely for MR subjects and it is not likely to occur in a clinical setting.

4.1.4 Conclusion

This study sought to determine the effects of MR on brain FDG uptake when MR sequences are run during the FDG uptake phase in an integrated PET/MR system. Group analyses of parametric images derived from the uptake phase, along with static images reconstructed from the data acquired 40-60 minutes post radiotracer administration showed foci of uptake in cortical areas associated with auditory processing likely a result of acoustic noise produced by the MR gradients. The relative increase in the PAC glucose uptake ranged from 3-9% depending on the image type and method. Using SPM8 and distorting the data to the MNI normalized space showed clusters of statistical significance only with an uncorrected $p < 0.001$, however no clusters were observed with a FWE of 0.05. ROI analysis showed similar results of increased uptake in the PAC with some ROIs significant and some trending towards significance, however they were insensitive to focal activations in large ROIs and activations which spanned the border of multiple ROIs. Statistical maps of the parametric and static images derived from FreeSurfer were qualitatively more similar than those derived from SPM8 with many of the clusters conserved, suggesting that in an integrated system where software

coregistration is unnecessary, an improvement in group analysis can be attained using FreeSurfer and the simultaneously acquired MR data. This suggests that while the PET/MR environment is considerably different than the traditional PET environment, its impact on patient physiology is minimal and can thus still produce comparable results. Our results also demonstrate that the performance of the PET scanner is virtually unaffected by the MR data acquisition and highly reproducible PET data can be obtained. Still, for certain experiments specific knowledge of the effects of integrated multimodal imaging is required.

4.2 Impact of PET Data Optimization on NNC112 Binding

Potential Estimation

In addition to using the PET and MR data to make inferences about one another or further characterize the system, the degree of the co-localization of MR and PET signals is often of interest. In this study, the PET imaging was performed using ^{11}C -NNC112, a dopamine (D1) receptor antagonist, while working-memory task-based fMRI activation data were collected. While fMRI can interrogate functional regions of the brain with a temporal resolution of seconds and allows the use of multiple and different stimuli over an imaging session, the kinetic modeling of the PET data is restricted to providing a snapshot over the entire imaging duration. The considerable length of scan time required to collect the information necessary for the kinetic modeling of the system necessitates the addition of motion correction as motion is likely to occur and patients may need to leave the scanner and return [244]. Similarly, as the spatial localization of the signal(s) is important, voxel-wise partial volume effect correction must be

applied. The goal of this section is to characterize the combined effect of MR-assisted motion correction and partial volume effects correction on the estimation of ^{11}C -NNC112 binding potential (BP) in healthy volunteers.

4.2.1 Materials and Methods

4.2.1.1 Data Acquisition

29 healthy volunteers (18 males, 11 females, mean age 32.7 ± 10.7 years) were scanned on the BrainPET as part of a larger study. All subjects gave written informed consent and were studied in accordance with a protocol approved by the affiliated Institutional Review Board. As ^{11}C -NNC112 is known to have non-specific binding to central 5HT2a receptors [245], subjects were administered an oral dose of 2 mg of resperidone, a 5HT2a antagonist, 90 minutes prior to radiotracer injection. PET data were acquired in list-mode format for 90-minutes following the intravenous administration of 9.59 ± 1.65 mCi (range 5.07-11.96mCi) of ^{11}C -NNC112 and dynamic PET images were generated. MR data acquired with the CP and 8-channel concentric head coils included sequences routinely used for brain studies at our center (e.g. AAScout, MPAGE, UTE, TOF, DTI, BOLD).

4.2.1.2 Data Analysis

To correct for stochastic motion over the duration of the scan, head motion estimates were derived from the MR data for all the EPI-based acquisitions and used to correct the PET data in LOR space before image reconstruction as described In Section 3.3. When MR data were unavailable to provide motion estimates, PAC-MMAN (Section 3.3.1.4) was used using a frame

length of 60 seconds for the first 40 minutes of data and 120 seconds for the remainder of the scan. For comparison, the data were also reconstructed without motion correction.

The head attenuation map was generated from the MPRAGE data using the pseudo-CT method (Section 3.1.1.3.3). The PET data were sorted into sinograms using a framing of: 8×10s, 3×20s, 2×30s, 1×60s, 1×120s, 1×180s, 8×300s, 4×600s and reconstructed using OSEM. Post-reconstruction partial volume effects correction was applied to the motion corrected PET images using the RBV method (described in Section 3.2). The ROIs used for partial volume effects correction were derived in part from the FreeSurfer segmentation. To the FreeSurfer segmentation, additional ROIs were added accounting for functional activation areas obtained from fMRI.

BPnd for each ROI was estimated in PMOD using the simplified reference tissue kinetic model (SRTM) and the cerebellum as a reference tissue. To quantify the improvement provided by motion and partial volume effect correction, the relative change (RC) for structures of interest were calculated, where the RC is given by

$$RC_{Corr} = \frac{BPnd_{Corr} - BPnd_{Uncorr}}{BPnd_{Uncorr}}. \quad (4.1)$$

To determine the RC for motion correction, the uncorrected BPnds are those derived from the uncorrected data; for the RC due to partial volume effects, $BPnd_{Uncorr}$ are the values derived from the data only corrected for motion while $BPnd_{Corr}$ are the values that are corrected for both motion and partial volume effects.

4.2.2 Results

4.2.2.1 Motion Correction

Of the 29 volunteers, three left the scanner for a period of time over the duration of the scan and had to be repositioned in the scanner. The motion of all subjects was successfully estimated over the scanning duration. Maximum translations of up to 9 mm and rotations of up to 12 degrees were observed. The 6-parameter motion estimates converted into a single value representing the displacement of a cube of length 2 cm with a center of [5 cm, 5 cm, 5 cm] revealed that all subjects moved over the scan duration. A plot of the subjects displacement values over the can be found in Figure 67.




Figure 67 NNC112 Motion Profiles. Each color represents an individual subject's displacement over the duration of the scan. The duration was calculated as the sum of the displacements of the vertices of a cube 2 cm in length and displaced from the origin by 5 cm along the principle axes.

The PET data were reconstructed with and without motion correction and TACs were determined for the selected brain regions. Significant deviations in the shape of the TACs were observed between motion corrected and nonmotion corrected data. TACs of the left and right caudate and cerebellar cortices before and after motion correction can be found in Figure 68.




Figure 68 Effect of Motion Correction on TACs . TACs without motion correction (left) show significant deviations from TACs when motion correction is implemented (right). Motion correction significantly changes the contrast between the caudate (red,blue) and cerebellar cortices (pink,navy) which in turn will have lead to a contrast in their estimated BPnd.

4.2.2.2 Partial Volume Effect Correction




Figure 69 Partial Volume Effect Correction for a representative NNC112 subject. From left to right: The skull stripped MPRAGE data which was used to derive the anatomic labels using FreeSurfer; a NNC112-PET frame acquired during the scanning duration; the PET frame corrected for partial volume effects using the GTM method which required the FreeSurfer labels; the final, fully corrected PET image using the RBV method which requires the GTM-PVEC image as an input.

Partial volume effect correction with resolution modeling was applied to all reconstructed motion corrected frames using the methods described in Section 3.2. All structures outside the brain were neglected from the GTM. The MR, motion corrected PET, and partial volume corrected image can be found in Figure 69.

4.2.2.3 Effect of Corrections

The percentage changes in BPnd after motion correction and partial volume effects correction for several ROIs in a representative subject are shown in Figure 70, revealing both under- and overestimation before correction. The cumulative effect exceeds 100% for some of the structures analyzed.




Figure 70 Effects of Motion and Partial Volume Effect Correction on a Single Subject

The mean BPnds across subjects, with the corresponding RCs can be found in Figure 71.

The RC of the BPnd suggested that the improvement due to partial volume effects was on average considerably larger than the improvement due to motion correction when considering all subjects as a single group. From the motion profiles, some subjects had considerably more motion than other, suggesting that the benefit from motion correction may not be uniform across all subjects.




Figure 71 Relative Change of BPnd Due to Corrections

To determine if there was a dependence of the RCs from motion correction, the subjects were separated into three groups based on the integral of their motion profiles: low, medium and high amplitude motion. The BPnds and RCs for the three groups can be found in Figure 72 and Figure 73 respectively.




Figure 72 BpNd by Motion Amplitude

The RC due to motion correction showed to be the greatest across all regions in the high amplitude motion group. The RC due to motion for the hippocampus (HPC), amygdala (AMYG), and pallidum (PAL) was greater in the low amplitude motion group than in the medium amplitude motion group, however more regions showed a change in the medium amplitude group. The HPC, AMYG, and PAL were some of the smaller regions and had lower uptake. RC due to partial volume effects was relatively stable across groups.




Figure 73 BNhd Relative Change By Motion Amplitude

4.2.3 Discussion

Motion and partial volume effects were shown to have considerable effects on the parametric data. The improvement due to motion correction can be the result of a number of phenomena. The first is the repositioning of the subject's head between sequences. This can be overt, as in the cases where the subject had to exit the scanner to urinate, but can also be due to more subtle causes, such as the patient shaking their head in response to questions from the scanner operator between MR scans. An example of such motion is shown in Figure 74.

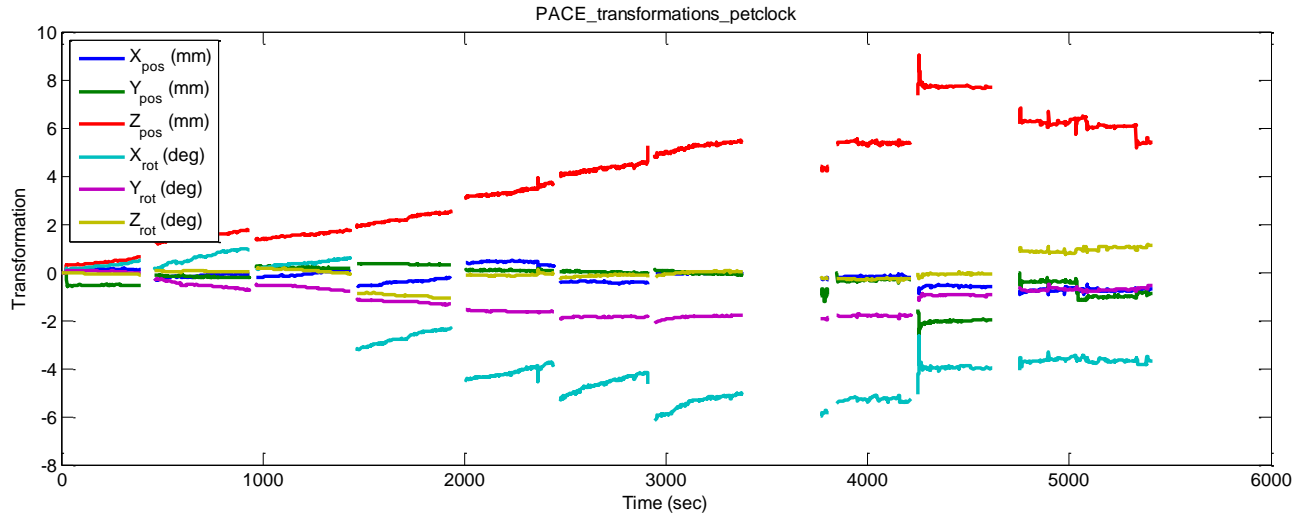


Figure 74 Intrasequence Motion for a Single Subject Derived from MR Estimates. Between each sequence the subject was questioned by study staff using the MR's intercom. The rotation about the x-axis observed at the start of each sequence between 1500 and 3500 seconds appears to be the patient lifting their head to hear the operator's question.

In addition to ensuring that the statistically defined ROI always corresponds to the corresponding anatomy, properly repositioning the data prevents bias due to mismatch between the emission data and the attenuation map. Additionally, motion correction fills in the gaps of the emission sinograms (Section 2.1) providing more complete information for OSEM to estimate the distribution from. Finally, as the motion correction is performed on a frame-by-frame basis, the data is better localized and the contrast is improved.

The change in the estimated BPNs due to the corrections was not uniform across structures, nor was it necessarily intuitive. When there are only two structures of interest and the tracer distribution is well characterized, as is the case with FDG where after the uptake phase the distribution in gray matter to white matter is on the order of 4:1, it is simple to contextualize the effect of partial volume effects on data; however, in a dynamic PET scan with many regions and no *a priori* knowledge of the expected tracer distribution, the problem becomes significantly more challenging. Predictably, those subjects who were determined to

have the most motion also had the largest benefit from motion correction; however the relative benefit of motion correction to partial volume effect correction was not intuitively obvious. Similarly, some ROIs showed large RCs as a result of motion correction but small changes with the addition of partial volume effects, while other ROIs exhibited the opposite effect.

Significant motion and partial volume effects were observed in all the subjects. Although MR-based attenuation correction is usually thought to have the largest influence on quantification in integrated PET/MRI scanners, less than 5% bias compared to the CT-based approaches was reported using recently developed methods (Section 3.1.1.3.3). This means the bias introduced by motion and partial volume effects becomes the dominant source of errors in the PET estimates. More importantly, the combined effect is more difficult to predict, depending on the size and location of the structure of interest and patient compliance. Without addressing these issues of motion and partial volume effects, the value of the BPnDs derived from these data is questionable.

4.3 MR-guided Image-derived PET Radiotracer Input

Function Estimation

The possibility of performing full kinetic modeling, opened by dynamic PET imaging, requires a measurement of the PET radiotracer input function. The arterial radiotracer input function is traditionally measured through sampling of the radial artery. The numerous contraindications for arterial sampling in addition to subject discomfort, requirements of additional staff, and subjecting staff to additional radiation exposure, make arterial sampling

non-ideal for routine practice. A number of image-derived input functions (IDIF) estimation methods for use in neurological PET studies have been proposed to address this problem (see [51] for an extensive review). These methods can broadly be separated into those that explicitly segment the voxels corresponding to vessels and those that choose the input function based on the time course of the signal. The latter methods are strongly reliant on the kinetic model and thus may not work with all tracers. Direct segmentation, however, is also subject to its challenges. Segmentation based on early phase PET images can be difficult due to image noise. Segmentation based on CT and MR is highly dependent on accurate coregistration, so much so that it has been suggested that the two carotid vessels be registered to the PET independently of one another [246]. Furthermore, MR data acquisition for MR-based segmentation can be time consuming.

In this section we present a novel IDIF estimation method that takes advantage of the advanced methods proposed in the previous sections for addressing most of these challenges mentioned above. Specifically, the vasculature is delineated from the simultaneously acquired MR data using multiple MR sequences. The deformation of the vessels is similar in the two datasets because of the simultaneous data acquisition. MR-based motion and partial volume effects correction eliminate two of the sources of bias in the resulting PET estimates of the average activity in the vessels (i.e. head motion and spill-in and spill-out effects, respectively). While the partial volume effect correction is applied using a late time-point blood sampling technique, the method could potentially yield adequate estimates without blood sampling using the rPSF of the scanner and an image-based partial volume effect correction algorithm like GTM.

4.3.1 Materials and Methods

4.3.1.1 MR-based Segmentation of Arterial Vasculature

4.3.1.1.1 MR Data Acquisition

The proposed IDIF method relies on two MR sequences to generate the arterial mask, **TOF** MR angiography (0.7×0.5×0.7 mm³, TR/TE=24/3.68 ms, FA=18°, TA=6:16 min) to capture the cerebral arteries as well as **MPRAGE** (1 mm isotropic, TR/TI/TE 2530/1200/1.64 ms, flip angle=7°, TA=4:56 min) to capture the internal carotid arteries.

4.3.1.1.2 Automatic Seeded Region-growing Based Extraction of Vasculature

TOF, the gold standard MR angiography technique, is time consuming and offers limited additional anatomic information. Rather than using TOF to cover the whole subject head, we propose a two-part method which includes the universally acquired MPRAGE images. Arterial blood has been noted to appear bright in MPRAGE images, even in the absence of gadolinium contrast, resulting from fully magnetized blood flowing into the field of view during the delay period after the inversion pulse leading to an inflow enhancement effect [247]. The possibility of deriving useful MRA-like information, especially at higher field strengths, from MPRAGE data has been described [248-250]. However, it was noted that TOF still had a higher vessel-to-background tissue signal [249]. Considering this, we propose a two-part method which uses the TOF to derive the arterial structure of the smaller vessels and uses the MPRAGE to determine the structure of the larger vessels. In this way the TOF field of view can be restricted to the brain, minimizing the need for additional scan time.

The mask generation involves two steps: the derivation of the fine and intermediate vasculature of the head using TOF followed by the extension of the ROI using the MPRAGE image. First, the TOF image is registered to the MPRAGE to account for inter-sequence motion, using SPM8 [228]. To enhance the contrast of the arteries, as well as to reduce non-uniform flip angle and coil sensitivity artifacts, the TOF image is filtered using New Segment (SPM8). As New Segment performs a registration to the MNI152 template as part of the bias correction, the inverse transformation is applied to an MNI template eye mask to create a subject specific eye mask which is used to remove its signal. Once the eyes have been removed, the contrast of the vasculature is enhanced using morphologic opening functions. The TOF image is opened separately with a 2D ball-shaped structuring element (15 voxel radius) as well as a 3D sphere (5 voxel radius) and the resulting images are multiplied together (Figure 75). Additionally the signal from outside the head is removed using a mask generated from skull stripping the coregistered MPRAGE data.

Step1: TOF-MRA filtering to extract arteries





Figure 75 TOF Morphological Filtering. Prior to the initial segmentation of the arteries from the TOF data, the image is filtered to enhance the contribution from the vessels. The eyes are masked from the data during the bias field correction step as they tend to enhance and can lead to artifacts during the extraction algorithm.

After the TOF data is filtered, the arteries are segmented using a seeded-region growing clustering algorithm. To determine the seed, a high threshold is used on the voxels that exist within a region defined by a 50×50 square in the center of each axial slice. Seed points that form objects of less than 5 voxels using a 3D 26-connected neighborhood are removed. An example of the initial seed can be found in Figure 76. The region is then grown by adding the highest intensity neighboring voxels to the ROI structure until the mean of the ROI voxels drop below a threshold value.

The second phase of the arterial ROI generation is performed using the MPRAGE image and the arterial ROI defined in the first step as the seed. As contrast can cause the veins to enhance (which is usually suppressed in non contrast-enhanced MPRAGE), this method is restricted to non contrast-enhanced MPRAGE images. To minimize spreading the mask into the surrounding tissue, the voxels that overlap with the acquired TOF are masked out, save for the bottom five planes of voxels to ensure there is some overlap between the seed and the arteries on the MPRAGE. The seed is then grown through the MPRAGE image using an algorithm similar to that of the first step. The steps are shown in Figure 76.



Initial Seed via Threshold

Seed->TOF

TOF->MPRAGE

Figure 76 Multi-step Seeded Region Growing of Arterial Mask. A Seed region is first generated from applying a threshold to the time-of-flight data. The region is then grown in two steps using a region growing algorithm and the time-of-flight (first) and MPRAGE (second) data.

4.3.1.2 Image Derived Input Function (IDIF) Generation

In order to accurately sample the peak, and to account for inherent variability in injection speed and style between hand injections, it is necessary to track the bolus and frame the PET list-mode accordingly. The bolus arrival time was determined using the four-part method described in Section 2.5.2.2. Next, the PET data were separated into early time-points and late time-points. For bolus studies the radiotracer is injected in a relatively short period of time (on the order of seconds to tens of seconds). In the time that immediately follows the injection the radiotracer concentration can change significantly, which requires high frequency sampling.

4.3.1.2.1 Early Time-point Reconstruction

4.3.1.2.1.1 CPR Reconstruction

To accurately estimate the arterial input function, the first pass and peak value of the radiotracer concentration in the blood must be accurately sampled. For fast bolus injections

this requires sampling on a very short time-scale (e.g. seconds). Standard iterative reconstruction methods are not ideal for reconstructing PET images on this time scale for two major reasons. First, common iterative reconstruction methods have a non-negativity constraint, which results in more positively biased images. Secondly, the low count data can make scatter estimation methods unreliable. As such, the first minute of the data, post-injection, was reconstructed using Complementary Reconstruction (CPR) [251]. Briefly, the short image frames are produced as the subtraction of two images in the following way:

$$\begin{aligned}
 D_T &= \sum_i D_i \\
 x_i^s &= R(D_i) \\
 x_i^c &= R(D_T) - R(D_T - D_i)
 \end{aligned}
 \tag{4.2}$$

where D_i is the sinogram or list-mode data in the i -th dynamic frame, D_T is the total data during the total scan period, and $R(\dots)$ is an iterative single-frame reconstruction operator, e.g. OP-OSEM, such that x_i^s and x_i^c are the reconstructed images using the iterative reconstruction operator and CPR, respectively. In this case, the reconstruction operator, $R(\dots)$, is the OSEM reconstruction algorithm. For the FDG study images were reconstructed using CPR at 5 second intervals where the total scan period was considered to be the first 60 seconds of data (i.e. $T=60$ sec and $i=1,2,3,\dots,60$ sec post-injection).

4.3.1.2.1.2 Optimization of Kernel and Frame Length

As the CPR method requires a subtraction of images longer in duration than the frame of interest, there is the potential for the introduction of bias in the measurement from the length of the kernel. To determine the optimal kernel length, CPR reconstructions were

performed with kernel sizes of 60, 120, and 180 seconds. Additionally the optimal frame length was determined by reconstructing the data at 1, 2, 3, 4, and 5 second intervals.

4.3.1.2.2 Late Time-point Reconstruction

Individual frame-by-frame reconstruction was performed on the PET list-mode data to generate a dynamic time series of images. The late time point PET data, the time points following the first minute after the radiotracer injection, were reconstructed with the standard 3D OP-OSEM algorithm using both prompt and variance-reduced random coincidence events [121] as well as normalization, scatter [34], and attenuation sinograms. Motion correction was performed on the PET images as outlined in Section 3.3. The attenuation sinograms were derived using the pseudo-CT method (Section 3.1.1.3.3).

4.3.1.2.3 Partial Volume Effect Correction of IDIF

Due to the relatively small size of the vessel as compared to the surrounding tissue, partial volume effects are significant. In the early time points, the partial volume effects are dominated by spill-out from the vessel as its activity is significantly higher than that of the surrounding tissue. At later time points however, the partial volume effects of the vessel are dominated by spill in. In order to accurately estimate the activity in the arterial mask, partial volume effects correction must be performed. This section examines a method that requires sparse late-time point blood samples.

4.3.1.2.3.1 Subsegmentation of the Arterial Mask

The PVC method presented is ROI-based and make the assumption that the activity inside the ROI is uniform. To allow for regional variability in the arterial mask, and as the contrast between the vessels and the surrounding tissue is not expected to be homogenous throughout the image—and thus the local spill-out and recovery coefficients change over the arterial mask—the arterial mask is broken down into subsegments. To automate the subsegmentation, a k-means cluster approach using squared Euclidean distance or a global geodesic distance metric were tested. As we are interested in determining an average radiotracer input function and not regional input functions, only the left and right internal carotid arteries were used for validation.

4.3.1.2.3.2 Chen's Method

The presented method, which requires late-time point blood samples, is a modified version of that presented by Chen *et al.* [49]. This method relies on the fact the spill-in and spill-out effects are purely geometric effects and are thus time invariant. As such, the image-measured radioactivity of the artery, c_p^m is the linear combination of the true radioactivity in the artery, c_p , and the radioactivity of the surrounding tissue, c_m , and are all related in the following way:

$$c_p^m(t) = rc \times c_p(t) + m_{tb} \times c_t(t) \quad (4.3)$$

where rc and m_{tb} are the recovery coefficient and the spill-over coefficient, respectively. The recovery and spill-over coefficients were estimated using both the linear least square method and measurements of c_p^m , c_p , and c_m at three separate time points (20, 40, 60 min post

injection for FDG). Similar to Chen *et al.*, regions where the recovery and spill-over coefficients were greater than one were excluded from the arterial mask. Although for these studies c_p was derived from arterial measurements (as venous blood samples were not available), we discuss below how venous measurements could be used as an alternative. The blood samples are then linearly interpolated and the activity corresponding to the midpoint of the corresponding PET frame is estimated. The rc and m_{tb} are determined for each subsegment of the arterial mask and the final IDIF is merely the sum of the individual subsegment corrected input functions, weighted by the percent contribution of the subsegment ROI to the total arterial ROI.

4.3.1.3 Validation

4.3.1.3.1 Validation of IDIF against Manual Arterial Sampling

To validate the IDIF for FDG, ten healthy volunteers (9 males and 1 female, mean age 41.7 ± 14.6 years) were enrolled in this study. The mean height and weight among subjects was 177 ± 10 cm and 81.6 ± 13.1 kg. All the subjects gave written informed consent in accordance with a protocol approved by the affiliated Institutional Review Board.

For all associated studies the subjects were injected with 5.11 ± 0.31 mCi of FDG while inside the PET/MR scanner in a head-first supine position while both MR and PET data acquisitions were being performed. All radiotracer injections were bolus injections performed by hand, followed by a 10 cc saline flush.

To compare the IDIF to the true AIF, arterial sampling was performed in all subjects. A licensed anesthesiologist placed a catheter in the radial artery proximal to the wrist under local anesthesia (lidocaine) with 8" of extension tubing and the a-line was accessed by an

experienced research nurse. Samples were drawn by hand with a vacutainer system into heparinized evacuated blood collection tubes. Prior to radiotracer injection, a 6 cc discard was drawn to eliminate dilution of the blood with saline. Eighteen 3 cc samples were drawn by hand ten seconds apart over the first 180 seconds post-injection. No saline flushes were performed between the ten second samples, however a 5-10 cc saline flush was performed after collection of the eighteenth sample. For all later time points there were (in chronological order): a 6 cc discard, a 3 cc sample, and a 5-10 cc saline flush. For FDG blood analysis, a 300 μ L sample of whole blood was pipetted from each collection tube and the activity was counted using a gamma counter and all counts were decay corrected to the time of injection. The remaining blood was centrifuged and a 300 μ L sample of plasma was also counted.

4.3.1.3.1.1 AUC Analysis and Kinetic Parameter Estimation

The area-under-the-curve (AUC) was calculated for both the AIF and the IDIF. To compute the AUC, piecewise cubic hermite interpolation was performed on the data points. Similarly, cerebral metabolic rate of glucose (CMRGlu) was calculated using both the AIF and IDIF.

4.3.1.3.2 *Comparison of Late Arterial Sampling to Venous Measurements*

The estimation of the recovery and spill-over coefficients is time independent, i.e. the blood samples can be taken at any time throughout the scan. This means that the blood samples can be isolated to late time points where the system has reached a quasi-steady state and there is minimal difference between the arterial and venous tracer values, and venous blood samples can be used instead of arterial sampling, obviating the need for invasive arterial

catheterization. Although venous sampling could be performed through an additional catheter, that means additional discomfort for the patient and another site for potential complications. An alternative to using an additional catheter is to sample the blood directly from the catheter used to inject the radiotracer into the patient [252-255]. The main concern with using the same catheter is the possibility of overestimation due to residual FDG sticking to the catheter walls, however with proper flushing of the line and pre-sample discards, the mean overestimation has been found to be small, on the order of 2% for FDG [253, 254].

To evaluate the possibility of using the injection catheter for venous sampling, a subset of the FDG subjects had additional venous samples drawn at 20, 40, and 60 minutes post injection. Following radiotracer injection, the injection catheter is flushed with 10 cc in the spare port of the dual port system. Additional, before each blood draw, the spare port was flushed with another 10 cc, followed by a 6 cc discard and a 3 cc venous blood sample. Analysis of the venous blood data was identical to the arterial blood.

While the method proposed here is theoretically capable of providing an individual IDIF for each subsegment of the arterial ROI given accurate registration and reconstruction, this work sought to validate its use in replicating the AIF derived from the radial artery. To this end, rather than using the entire subsegmentation, segments spanning the petrous, lacerum, and cavernous segments of the left and right internal carotid arteries were used.

4.3.2 Results

4.3.2.1 MR-based Segmentation

The multi-sequence seeded-region growing segmentation of the vasculature was implemented and was capable of running in an automated fashion without the need for additional user input. The segmented vasculature, along with the MPRAGE and TOF MR data for a representative subject can be found in Figure 77. For cases where no TOF data was collected, it is possible to segment the larger vessels of the head using only the MPRAGE and user defined points.




Figure 77 Multi-MR-Sequence Derivation of Arterial Mask. A seeded region growing is applied to the TOF MR image (A). The final arterial mask (B) is produced from a secondary seeded region growing algorithm, where the results of the TOF segmentation are used as the seed and the image is the MEMPRAGE (C; shown here as a maximum intensity projection).

4.3.2.2 Image-Derived Input Function

The early and late frames were reconstructed using CPR and traditional reconstruction schemes, respectively. The early frames were found to most accurately model the input function and have the least amount of noise between successive measurements when

reconstructed with a CPR kernel size of 60 seconds and a frame length of 5 seconds. Motion throughout the duration of the scan was estimated using PET-based and MR-based methods and used in the reconstruction of the images. The generated arterial ROIs were then segmented into 12 subsets and applied to the data using Chen’s method. The IDIFs were successfully determined for the ten healthy controls. A representative IDIF can be found in Figure 78. The mean IDIF AUC was found to be $101.3\% \pm 4.7\%$ of the AIF AUC with an absolute difference of $4.1\% \pm 2.1\%$.




Figure 78 Representative Image-Derived Input Function. The IDIF is shown on a linear scale (top) as well as a log scale (bottom).

4.3.2.3 Kinetic Parameter Estimation

To quantify the error introduced by using the IDIF as compared to the gold standard AIF, TACs were generated for 116 ROIs determined using FreeSurfer for each patient. The CMRGIc was then calculated for each ROI using the AIF and the IDIF and the results were plotted against

each other. A line was fit to the data which represented the accuracy of the IDIF, where a slope of one would result if the AIF and IDIF were exactly the same. The average slope for the 10 subjects was found to be $98.3\% \pm 5.1\%$. A representative plot for one of the subjects can be found in Figure 79.




Figure 79 Kinetic Parameter Estimation Using AIF and IDIF for a Representative Subject. 116 anatomic ROIs were generated using the subject-specific MPRAGE MR image and FreeSurfer. PET data were reconstructed using the reconstruction algorithms mentioned throughout this work. An IDIF was reconstructed using the steps outline throughout Section 4.3. CMR_{Glc} values were derived for each ROI using the manually sampled AIF and the IDIF.

4.3.2.4 Comparison of Late Arterial Sampling to Venous Measurements

The results of three trials where arterial and venous blood were sampled simultaneously can be found in Figure 80. Of the ten healthy volunteers, venous blood from the injection catheter was drawn on nine of them and all three time points were only drawn on eight of them. The second and third blood samples always showed a high degree of similarity between the arterial and venous samples, however the first draw was found to be less reproducible and was significantly higher in five of the eight cases. This is believed to be the result of the limited experience in how to perform the sampling from the injection catheter (including insufficient flushing of the injection catheter) as evident with the improvement in the reproducibility and similarity of the first measurement as more trials were performed.




Figure 80 Comparison of Late Time-Point Blood Sampling Using the Injection Catheter to Arterial Line Measurements. Comparisons between the arterial activity as measured from an arterial line and the venous activity as measured from the injection catheter are shown for three separate subjects.

4.3.2.5 Application of IDIF to Other Tracers

At no point does the presented method rely on any assumptions on the characteristics of the radioisotope or the tracer, rather it can be thought of as an estimation of the radioactivity distribution in the arterial compartment. The partial volume corrupted IDIF is determined from the dynamic data through the use of an arterial mask so the only assumption is that the free radiotracer signal can be measured from the blood. To consider the IDIF an estimate of the free radiotracer distribution, the assumption must be made that the signal from the intra-arterial compartment is dominated by the free radiotracer; however, even in cases where that may not be true (i.e. active metabolites in the intra-arterial compartment), the method provides an estimate of the intra-arterial compartment. To determine if this procedure could be applied to other tracers as a measurement of the intra-arterial compartment where radioactive metabolites were present, the proposed TOF+MPRAGE derived arterial ROIs were then applied to ^{11}C -NNC112 and ^{11}C -PBR28 datasets. In the case of NNC112 and PBR28, there are radioactive metabolites present in the blood as the tracer does not become trapped in a

tissue. Since the blood samples used to perform the partial volume effect correction also contain parent and metabolite compounds the resulting IDIF is the cumulative signal of the free parent radiotracer and the metabolites. IDIFs derived using Chen's method with the ROIs can be found in Figure 81. While the proposed method does not directly derive the true free-radiotracer input function for tracers with metabolites in the intra-arterial compartment, assumptions can be made about the metabolite correction or the required blood samples could be used for metabolite correction.

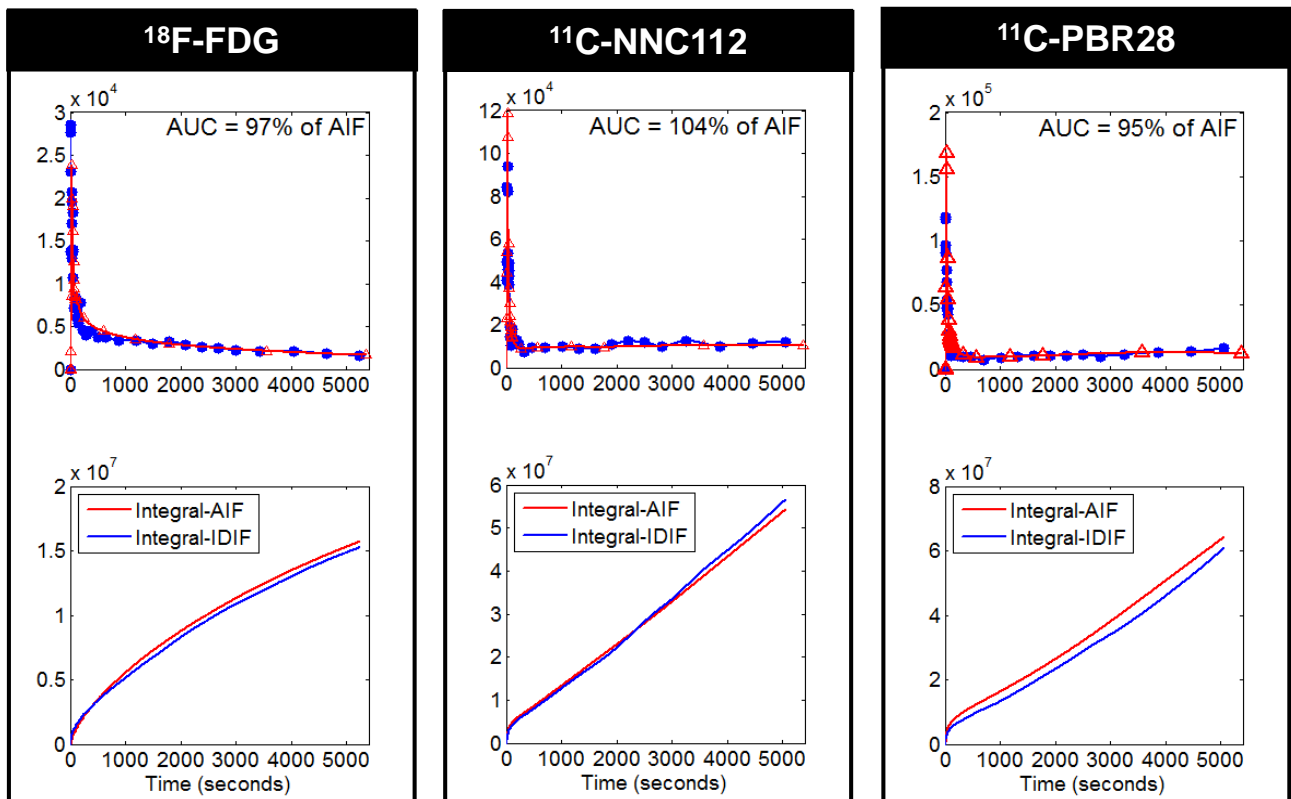


Figure 81 Representative TOF+MPRAGE PET IDIF Results for FDG, NNC112, and PBR28. The method presented, the acquisition of MR data to derive an arterial mask followed by late time-point blood sampling to scale the arterial ROI measurements and correct for partial volume effects, holds for non- ^{18}F tracers (NNC112) as well as tracers that do not obey the two compartment model (PBR28). At no point are assumptions made on the kinetics of the tracer, its injection, or the decay properties of the radioisotope.

4.3.3 Discussion

A minimally invasive method to improve the estimation of the radiotracer AIF has been presented. The arterial ROI was segmented in an automated fashion using a combination of MPRAGE and TOF MR datasets. By using the bright-blood phenomena of the MPRAGE, the TOF field of view can be restricted, reducing the total scan time. The frames used to estimate the bolus were reconstructed using CPR. Comparing kernel sizes and frame lengths showed that the best results were achieved with a kernel of 60 seconds and a frame length of 5 seconds. The IDIF estimated from the dynamic PET data using Chen's method showed AUC values for the IDIF comparable to the AIF, and the kinetic parameters estimated from them correlated well with those derived from arterial sampling. The method was also shown to have the potential to be applied to other tracers, however this must be confirmed by evaluating more datasets.

A major drawback to the presented method for IDIF estimation is the partial volume effect correction using Chen's blood-based method. The major benefit of a blood-less method for partial volume effect correction and scaling like GTM is that it does not require any blood sampling. However, it requires that the PSF of the scanner be known. As such, while it can be implemented on the BrainPET, where the rPSF has been modeled, it cannot easily be implemented on scanners where the PSF is unknown. Additionally, for image-based partial volume effect correction methods to achieve accurate results in larger carotids, those which lie outside the brain and adjacent to noncerebral tissues (e.g. muscles, soft tissue), accurate segmentation of adjacent extra-cerebral tissues is required, which is difficult to automate and time-consuming to perform manually. Similarly, non-blood derived input functions have been shown to be sensitive to patient motion [256], though through implementing MR-based motion

correction methods this issue can be minimized. On the other hand, while blood-based methods have been shown to be more resistant to patient motion [257], they require additional blood sampling. While we have demonstrated that with adequate flushing of the injection catheter, it can be used to sample blood in FDG studies, obviating the need to place an additional catheter for sampling, this may not hold true for other tracers that may impregnate or stick to the injection catheter material. Additional studies are required to determine if the injection catheter can be used for sampling tracers other than FDG. Furthermore, for tracers that undergo metabolism (e.g. ^{11}C -PBR28, ^{11}C -NNC112), unless a population-based metabolism curve is used, blood samples cannot be avoided.

This study sought to compare our IDIF to the “imperfect gold standard” of manual arterial sampling. To obtain the AIF, hand-sampling was performed, which can be inaccurate at estimating high-frequency changes (e.g. the peak and first pass). This in turn may enhance the bias presented in this study. While numerous devices exist for the aid and automation of blood draws and analysis, many such devices are not MR-compatible, so at this time the use of hand-drawn samples is not unreasonable. Similarly, it is not a reasonable expectation for clinical centers to invest in automated sampling equipment, so hand-drawn sampling would be the more likely method to translate to the clinic. To address this, a comparison of our IDIF to samples collected and counted using an MR-compatible automated blood sampler [177] should be performed.

Another potential source of systematic bias in the measurement could be inaccurate attenuation (and scatter) correction. While the effects of inaccurate bone estimation in the head attenuation map have been minimized using the methods discussed in Chapter 3, a

systematic study of the effects that attenuation maps can have on the extracerebral areas is warranted as portions of the carotid run close to the bone, a region where misclassification and errors in attenuation estimation are highest.

Chapter 5 CONCLUSION

This work explored techniques for improved PET data quantification in a simultaneous PET/MR system for neurological imaging. First, the integrated brain PET/MR system was characterized and the PET data processing and image reconstruction were optimized. The PET and MR scanners were found to minimally interfere with each other, permitting the simultaneous collection of data. Algorithms were then presented for the spatial and temporal alignment of the MR and PET data and for the calibration of the PET scanner. Using the simultaneously collected MR data, we have proposed a unified processing pipeline for the reconstruction and quantification of PET data. As demonstrated in this work, MRI can not only provide functional and anatomic information, but it can also be used to improve the quality of the PET data. Specifically, MR-derived data can be used for PET attenuation, motion and partial volume effects corrections as well to inform an algorithm for estimating the PET radiotracer input function directly from the PET images.

To aid in the implementation and application of advanced processing and analysis techniques, Masamune, a graphical user interface was constructed, allowing the streamlined and automated handling of the PET data. To date, Masamune has been used on over 500 PET/MR scans of humans and 225 PET/MR scans of nonhuman primates and other animals. The Matlab environment in which Masamune was implemented provides an expandable

platform that allows the integration of other specialized neuroimaging software (e.g. SPM, Comkat, FreeSurfer) in an automated fashion.

Simultaneous PET/MR offers a unique opportunity to address some of the challenges that limit the implementation of quantitative PET into routine clinical imaging. While PET/MR provides the opportunity to improve our understanding of the mechanisms underlying disease, attention should be paid to how it can improve routine imaging. Moving forward, it is likely that the scanning period will no longer be thought of as sequential bed positions or sequences that are put into context later, but rather as a continuous flow of information. Providing a platform that interprets this continuous flow of information will have a strong impact on the quality of the resulting data and the simplification through automation can make it practical for the additional challenges of clinical imaging.

BIBLIOGRAPHY

1. Beyer T, Pichler B. A decade of combined imaging: from a PET attached to a CT to a PET inside an MR. *Eur J Nucl Med Mol Imaging*. 2009;36 Suppl 1:S1-2.
2. Wahl RL, Wagner HN, Beanlands RS. Principles and practice of PET and PET/CT: Lippincott Williams & Wilkins Philadelphia, PA;; 2009.
3. Smith SM, Jenkinson M, Woolrich MW, et al. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*. 2004;23, Supplement 1:S208-S19.
4. Bailey DL, Townsend DW, Valk PE, et al. Positron emission tomography: Springer; 2005.
5. Wernick MN, Aarsvold JN. Emission tomography: the fundamentals of PET and SPECT: Academic Press; 2004.
6. Caldeira L, Scheins JJ, Almeida P, et al. Maximum a Posteriori Reconstruction using PRESTO and PET/MR data acquired Simultaneously with the 3TMR-BrainPET. Nuclear Science Symposium Conference Record (NSS/MIC), 2010 IEEE; 2010. p. 2879-84.
7. Akbarzadeh A, Ay MR, Ahmadian A, et al. Impact of using different tissue classes on the accuracy of MR-based attenuation correction in PET-MRI. Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE; 2011. p. 2524-30.
8. Hofmann M, Bezrukov I, Mantlik F, et al. MRI-based attenuation correction for whole-body PET/MRI: quantitative evaluation of segmentation- and atlas-based methods. *J Nucl Med*. 2011;52:1392-9.
9. Kinahan PE, Townsend DW, Beyer T, et al. Attenuation correction for a combined 3D PET/CT scanner. *Medical Physics*. 1998;25:2046-53.
10. Nakamoto Y, Osman M, Cohade C, et al. PET/CT: Comparison of Quantitative Tracer Uptake Between Germanium and CT Transmission Attenuation-Corrected Images. *Journal of Nuclear Medicine*. 2002;43:1137-43.
11. Alvarez RE, Macovski A. Energy-Selective Reconstructions in X-Ray Computerized Tomography. *Physics in Medicine and Biology*. 1976;21:733-44.
12. Beyer T, Kinahan PE, Townsend DW, et al. The use of X-ray CT for attenuation correction of PET data. Nuclear Science Symposium and Medical Imaging Conference, 1994, 1994 IEEE Conference Record; 1994. p. 1573-7 vol.4.
13. Lacroix KJ, Tsui BMW, Hasegawa BH, et al. Investigation of the Use of X-Ray Ct Images for Attenuation Compensation in Spect. *IEEE Transactions on Nuclear Science*. 1994;41:2793-9.
14. Robinson PJ, Kreel L. Pulmonary Tissue Attenuation with Computed-Tomography - Comparison of Inspiration and Expiration Scans. *J Comput Assist Tomo*. 1979;3:740-8.
15. Hubbell JH. Photon Mass Attenuation and Energy-Absorption Coefficients from 1 Kev to 20 Mev. *Int J Appl Radiat Is*. 1982;33:1269-90.

16. Mollet P, Keereman V, Bini J, et al. Improvement of Attenuation Correction in Time-of-Flight PET/MR Imaging with a Positron-Emitting Source. *J Nucl Med.* 2014;55:329-36.
17. Zaidi H, Montandon ML, Slosman DO. Magnetic resonance imaging-guided attenuation and scatter corrections in three-dimensional brain positron emission tomography. *Medical Physics.* 2003;30:937-48.
18. Gatehouse PD, Bydder GM. Magnetic Resonance Imaging of Short T2 Components in Tissue. *Clinical Radiology.* 2003;58:1-19.
19. Gold GE, Thedens DR, Pauly JM, et al. MR imaging of articular cartilage of the knee: new methods using ultrashort TEs. *AJR Am J Roentgenol.* 1998;170:1223-6.
20. Catana C, Wu Y, Judenhofer MS, et al. Simultaneous Acquisition of Multislice PET and MR Images: Initial Results with a MR-Compatible PET Scanner. *J Nucl Med.* 2006;47:1968-76.
21. Keereman V, Fierens Y, Broux T, et al. MRI-based attenuation correction for PET/MRI using ultrashort echo time sequences. *J Nucl Med.* 2010;51:812-8.
22. Dogdas B, Shattuck DW, Leahy RM. Segmentation of skull and scalp in 3-D human MRI using mathematical morphology. *Hum Brain Mapp.* 2005;26:273-85.
23. Hofmann M, Steinke F, Scheel V, et al. MRI-based attenuation correction for PET/MRI: a novel approach combining pattern recognition and atlas registration. *J Nucl Med.* 2008;49:1875-83.
24. Kops ER, Herzog H. Alternative methods for attenuation correction for PET images in MR-PET scanners. *Nuclear Science Symposium Conference Record, 2007 NSS '07 IEEE; 2007.* p. 4327-30.
25. Navalpakkam BK, Braun H, Kuwert T, et al. Magnetic resonance-based attenuation correction for PET/MR hybrid imaging using continuous valued attenuation maps. *Invest Radiol.* 2013;48:323-32.
26. Berger MJ, Hubbell J, Seltzer S, et al. XCOM: photon cross sections database.
27. Cherry SR, Sorenson JA, Phelps ME. chapter 6 - Interaction of Radiation with Matter. In: Cherry SR, Sorenson JA, Phelps ME, editors. *Physics in Nuclear Medicine (Fourth Edition).* Philadelphia: W.B. Saunders; 2012:63-85.
28. Evans RD, Noyau A. *The atomic nucleus: McGraw-Hill New York; 1955.*
29. Schmand M, Wienhard K, Casey ME, et al. Performance evaluation of a new LSO high resolution research tomograph-HRRT. *Nuclear Science Symposium, 1999 Conference Record 1999 IEEE; 1999.* p. 1067-71 vol.2.
30. Watson CC, Newport D, Casey ME. A Single Scatter Simulation Technique for Scatter Correction in 3D PET. In: Grangeat P, Amans J-L, editors. *Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine.* Springer Netherlands; 1996:255-68.
31. Accorsi R, Adam L-E, Werner ME, et al. Optimization of a fully 3D single scatter simulation algorithm for 3D PET. *Physics in Medicine and Biology.* 2004;49:2577.
32. Barney JS, Rogers JG, Harrop R, et al. Object shape dependent scatter simulations for PET. *Nuclear Science, IEEE Transactions on.* 1991;38:719-25.
33. Thompson CJ. The problem of scatter correction in positron volume imaging. *Medical Imaging, IEEE Transactions on.* 1993;12:124-32.
34. Watson CC. New, faster, image-based scatter correction for 3D PET. *Nuclear Science, IEEE Transactions on.* 2000;47:1587-94.
35. Watson CC, Newport D, Casey ME, et al. Evaluation of simulation-based scatter correction for 3-D PET cardiac imaging. *Nuclear Science, IEEE Transactions on.* 1997;44:90-7.

36. Daube-Witherspoon ME, Carson RE. Unified deadtime correction model for PET. *Medical Imaging, IEEE Transactions on*. 1991;10:267-75.
37. Haselman MD, Hauck S, Lewellen TK, et al. FPGA-Based Pulse Pileup Correction. *IEEE Nuclear Science Symposium conference record Nuclear Science Symposium*. 2010:3105-12.
38. Vicente E, Herraiz JL, Espana S, et al. Deadtime and pile-up correction method based on the singles to coincidences ratio for PET. *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE; 2011*. p. 2933-5.
39. Kostakoglu L, Goldsmith SJ. 18F-FDG PET Evaluation of the Response to Therapy for Lymphoma and for Breast, Lung, and Colorectal Carcinoma. *Journal of Nuclear Medicine*. 2003;44:224-39.
40. Kim CK, Gupta NC. Dependency of standardized uptake values of fluorine-18 fluorodeoxyglucose on body size: comparison of body surface area correction and lean body mass correction. *Nucl Med Commun*. 1996;17:890-4.
41. Kim CK, Gupta NC, Chandramouli B, et al. Standardized uptake values of FDG: body surface area correction is preferable to body weight correction. *J Nucl Med*. 1994;35:164-7.
42. Schomburg A, Bender H, Reichel C, et al. Standardized uptake values of fluorine-18 fluorodeoxyglucose: the value of different normalization procedures. *Eur J Nucl Med*. 1996;23:571-4.
43. Zasadny KR, Wahl RL. Standardized uptake values of normal tissues at PET with 2-[fluorine-18]-fluoro-2-deoxy-D-glucose: variations with body weight and a method for correction. *Radiology*. 1993;189:847-50.
44. Hoekstra CJ, Paglianiti I, Hoekstra OS, et al. Monitoring response to therapy in cancer using [18F]-2-fluoro-2-deoxy-d-glucose and positron emission tomography: an overview of different analytical methods. *European Journal of Nuclear Medicine*. 2000;27:731-43.
45. Langen KJ, Braun U, Rota Kops E, et al. The influence of plasma glucose levels on fluorine-18-fluorodeoxyglucose uptake in bronchial carcinomas. *J Nucl Med*. 1993;34:355-9.
46. Lindholm P, Minn H, Leskinen-Kallio S, et al. Influence of the blood glucose concentration on FDG uptake in cancer--a PET study. *J Nucl Med*. 1993;34:1-6.
47. Wahl RL, Henry CA, Ethier SP. Serum glucose: effects on tumor and normal tissue accumulation of 2-[F-18]-fluoro-2-deoxy-D-glucose in rodents with mammary carcinoma. *Radiology*. 1992;183:643-7.
48. Sokoloff L, Reivich M, Kennedy C, et al. The [14C]deoxyglucose method for the measurement of local cerebral glucose utilization: theory, procedure, and normal values in the conscious and anesthetized albino rat. *J Neurochem*. 1977;28:897-916.
49. Chen K, Bandy D, Reiman E, et al. Noninvasive quantification of the cerebral metabolic rate for glucose using positron emission tomography, 18F-fluoro-2-deoxyglucose, the Patlak method, and an image-derived input function. *Journal of Cerebral Blood Flow & Metabolism*. 1998;18:716-23.
50. Scheer B, Perel A, Pfeiffer U. Clinical review: complications and risk factors of peripheral arterial catheters used for haemodynamic monitoring in anaesthesia and intensive care medicine. *Critical Care*. 2002;6:199-204.
51. Zanotti-Fregonara P, Chen K, Liow JS, et al. Image-derived input function for brain PET studies: many challenges and few opportunities. *Journal of Cerebral Blood Flow & Metabolism*. 2011.

52. Fung EK, Planeta-Wilson B, Mulnix T, et al. A multimodal approach to image-derived input functions for brain PET. *IEEE*; 2009. p. 2710-4.
53. Gjedde A. High- and low-affinity transport of D-glucose from blood to brain. *J Neurochem*. 1981;36:1463-71.
54. Patlak CS, Blasberg RG. Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data. Generalizations. *J Cereb Blood Flow Metab*. 1985;5:584-90.
55. Patlak CS, Blasberg RG, Fenstermacher JD. Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data. *J Cereb Blood Flow Metab*. 1983;3:1-7.
56. Logan J, Fowler JS, Volkow ND, et al. Graphical analysis of reversible radioligand binding from time-activity measurements applied to [N-11C-methyl]-(-)-cocaine PET studies in human subjects. *J Cereb Blood Flow Metab*. 1990;10:740-7.
57. Levin CS, Hoffman EJ. Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution. *Physics in Medicine and Biology*. 1999:781.
58. Bing B, Laforest R, Smith AM, et al. Evaluation of MAP image reconstruction with positron range modeling for 3D PET. *Nuclear Science Symposium Conference Record, 2005 IEEE*; 2005. p. 2686-9.
59. Hammer BE, Christensen NL, Heil BG. Use of a magnetic field to increase the spatial resolution of positron emission tomography. *Medical Physics*. 1994;21:1917-20.
60. Herzog H, Iida H, Weirich C, et al. Influence from high and ultra-high magnetic field on positron range measured with a 9.4TMR-BrainPET. *Nuclear Science Symposium Conference Record (NSS/MIC), 2010 IEEE*; 2010. p. 3410-3.
61. Iida H, Kanno I, Miura S, et al. A Simulation Study of a Method to Reduce Positron Annihilation Spread Distributions Using a Strong Magnetic Field in Positron Emission Tomography. *Nuclear Science, IEEE Transactions on*. 1986;33:597-600.
62. Raylman RR, Hammer BE, Christensen NL. Combined MRI-PET scanner: a Monte Carlo evaluation of the improvements in PET resolution due to the effects of a static homogeneous magnetic field. *Nuclear Science, IEEE Transactions on*. 1996;43:2406-12.
63. Wirrwar A, Vosberg H, Herzog H, et al. 4.5 tesla magnetic field reduces range of high-energy positrons-potential implications for positron emission tomography. *Nuclear Science, IEEE Transactions on*. 1997;44:184-9.
64. Rahmim A, Qi J, Sossi V. Resolution modeling in PET imaging: Theory, practice, benefits, and pitfalls. *Medical Physics*. 2013;40:064301.
65. Thomas BA, Erlandsson K, Modat M, et al. The importance of appropriate partial volume correction for PET quantification in Alzheimer's disease. *Eur J Nucl Med Mol Imaging*. 2011;38:1104-19.
66. Rousset OG, Ma Y, Wong DF, et al. Chapter 10 - Pixel-Versus Region-Based Partial Volume Correction in PET. In: Carson RE, Daube-Witherspoon ME, Herscovitch P, editors. *Quantitative Functional Brain Imaging with Positron Emission Tomography*. San Diego: Academic Press; 1998:67-75.
67. Montgomery AJ, Thielemans K, Mehta MA, et al. Correction of Head Movement on PET Studies: Comparison of Methods. *J Nucl Med*. 2006;47:1936-44.
68. Goldstein SR, Daube-Witherspoon ME, Green MV, et al. A head motion measurement system suitable for emission computed tomography. *IEEE Trans Med Imaging*. 1997;16:17-27.

69. Lopresti BJ, Russo A, Jones WF, et al. Implementation and performance of an optical motion tracking system for high resolution brain PET imaging. *Nuclear Science, IEEE Transactions on*. 1999;46:2059-67.
70. Picard Y, Thompson CJ. Motion correction of PET images using multiple acquisition frames. *IEEE Trans Med Imaging*. 1997;16:137-44.
71. Rahmim A, Rousset O, Zaidi H. Strategies for Motion Tracking and Correction in PET. *PET Clinics*. 2007;2:251-66.
72. Nihashi T, Dahabreh IJ, Terasawa T. Diagnostic accuracy of PET for recurrent glioma diagnosis: a meta-analysis. *AJNR Am J Neuroradiol*. 2013;34:944-50, S1-11.
73. Grosu AL, Piert M, Weber WA, et al. Positron emission tomography for radiation treatment planning. *Strahlenther Onkol*. 2005;181:483-99.
74. Nihashi T, Dahabreh IJ, Terasawa T. PET in the Clinical Management of Glioma: Evidence Map. *American Journal of Roentgenology*. 2013;200:W654-W60.
75. Zipursky RB, Meyer JH, Verhoeff NP. PET and SPECT imaging in psychiatric disorders. *Canadian Journal of Psychiatry*. 2007;52:146.
76. Saleem A, Brown GD, Brady F, et al. Metabolic activation of temozolomide measured in vivo using positron emission tomography. *Cancer Res*. 2003;63:2409-15.
77. Du ZD, Wang R, Prakash R, et al. Transcranial magnetic stimulation in schizophrenia: the contribution of neuroimaging. *Curr Top Med Chem*. 2012;12:2452-7.
78. Li Y, Rinne JO, Mosconi L, et al. Regional analysis of FDG and PIB-PET images in normal aging, mild cognitive impairment, and Alzheimer's disease. *Eur J Nucl Med Mol Imaging*. 2008;35:2169-81.
79. Bernstein MA, King KF, Zhou XJ. *Handbook of MRI Pulse Sequences*: Elsevier Science; 2004.
80. Haacke EM, Brown RW, Thompson MR, et al. *Magnetic resonance imaging*: Wiley-Liss New York.; 1999.
81. Nishimura DG. *Principles of magnetic resonance imaging*: Stanford University; 1996.
82. Griffiths DJ, Harris EG. *Introduction to quantum mechanics*: Prentice Hall New Jersey; 1995.
83. American College of Radiology (ACR), American Society of Neuroradiology (ASNR), Society for Pediatric Radiology (SPR). *ACR–ASNR–SPR Practice Guideline For the Performance and Interpretation of Magnetic Resonance Imaging (MRI) of the Brain*. Reston, VA: American College of Radiology (ACR); Revised 2013 (Resolution 6).
84. van den Hauwe L, Parizel PM, Van Goethem JW, et al. Clinical usefulness of contrast-enhanced MP-RAGE of the brain. *Neuroradiology*. 1996;38:S14-S9.
85. Frisoni GB, Fox NC, Jack CR, et al. The clinical use of structural MRI in Alzheimer disease. *Nat Rev Neurol*. 2010;6:67-77.
86. Lévy C, Laissy JP, Raveau V, et al. Carotid and vertebral artery dissections: three-dimensional time-of-flight MR angiography and MR imaging versus conventional angiography. *Radiology*. 1994;190:97-103.
87. Dagirmanjian A, Ross JS, Obuchowski N, et al. High resolution, magnetization transfer saturation, variable flip angle, time-of-flight MRA in the detection of intracranial vascular stenoses. *J Comput Assist Tomo*. 1995;19:700-6.
88. Oelerich M, Lentschig MG, Zunker P, et al. Intracranial vascular stenosis and occlusion: comparison of 3D time-of-flight and 3D phase-contrast MR angiography. *Neuroradiology*. 1998;40:567-73.

89. Sailer AM, Wagemans BA, Nelemans PJ, et al. Diagnosing intracranial aneurysms with MR angiography: systematic review and meta-analysis. *Stroke*. 2014;45:119-26.
90. White PM, Teasdale EM, Wardlaw JM, et al. Intracranial aneurysms: CT angiography and MR angiography for detection prospective blinded comparison in a large patient cohort. *Radiology*. 2001;219:739-49.
91. White PM, Wardlaw JM, Easton V. Can Noninvasive Imaging Accurately Depict Intracranial Aneurysms? A Systematic Review 1. *Radiology*. 2000;217:361-70.
92. Binder J. fMRI of Language Systems: Methods and Applications. In: Faro S, Mohamed F, editors. *Functional MRI*. Springer New York; 2006:245-77.
93. Horská A, Tkáč I. Magnetic Resonance Spectroscopy: Clinical Applications. In: Faro SH, Mohamed FB, Law M, et al., editors. *Functional Neuroradiology*. Springer US; 2012:155-94.
94. Lecomte R. Technology challenges in small animal PET imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. 2004;527:157-65.
95. Phelps ME. *PET : molecular imaging and its biological applications*. New York: Springer; 2004.
96. Cherry SR. The 2006 Henry N. Wagner Lecture: Of mice and men (and positrons)--advances in PET imaging technology. *J Nucl Med*. 2006;47:1735-45.
97. Boone JM, Velazquez O, Cherry SR. Small-animal X-ray dose from micro-CT. *Mol Imaging*. 2004;3:149-58.
98. Gould KL, Pan T, Loghin C, et al. Frequent Diagnostic Errors in Cardiac PET/CT Due to Misregistration of CT Attenuation and Emission PET Images: A Definitive Analysis of Causes, Consequences, and Corrections. *J Nucl Med*. 2007;48:1112-21.
99. Martinez-Moller A, Souvatzoglou M, Navab N, et al. Artifacts from Misaligned CT in Cardiac Perfusion PET/CT Studies: Frequency, Effects, and Potential Solutions. *J Nucl Med*. 2007;48:188-93.
100. Pichler BJ, Judenhofer MS, Pfannenbergl C. Multimodal Imaging Approaches: PET/CT and PET/MRI. 2008:109-32.
101. Christensen NL, Hammer BE, Heil BG, et al. Positron Emission Tomography within a Magnetic-Field Using Photomultiplier Tubes and Lightguides. *Physics in Medicine and Biology*. 1995;40:691-7.
102. Lucas AJ, Hawkes RC, Guerra P, et al. Development of a combined microPET®-MR system. *Nuclear Science Symposium Conference Record, 2006 IEEE*; 2006. p. 2345-8.
103. Garlick PB, Marsden PK, Cave AC, et al. PET and NMR dual acquisition (PANDA): applications to isolated, perfused rat hearts. *NMR in Biomedicine*. 1997;10:138-42.
104. Catana C, Procissi D, Wu Y, et al. Simultaneous in vivo positron emission tomography and magnetic resonance imaging. *Proceedings of the National Academy of Sciences*. 2008;105:3705-10.
105. Shao YP, Cherry SR, Farahani K, et al. Simultaneous PET and MR imaging. *Physics in Medicine and Biology*. 1997;42:1965-70.
106. Pichler BJ, Judenhofer MS, Catana C, et al. Performance Test of an LSO-APD Detector in a 7-T MRI Scanner for Simultaneous PET/MRI. *J Nucl Med*. 2006;47:639-47.
107. Marsden PK, Strul D, Keevil SF, et al. Simultaneous PET and NMR. *Br J Radiol*. 2002;75:S53-9.
108. Gilbert KM, et al. Design of field-cycled magnetic resonance systems for small animal imaging. *Physics in Medicine and Biology*. 2006;51:2825.

109. Gilbert KM, Scholl TJ, Handler WB, et al. Evaluation of a positron emission tomography (PET)-compatible field-cycled MRI (FCMRI) scanner. *Magnetic Resonance in Medicine*. 2009;62:1017-25.
110. Liu J, Tian J. Registration of brain MRI/PET images based on adaptive combination of intensity and gradient field mutual information. *International journal of biomedical imaging*. 2007;2007.
111. Guimaraes AR, Melcher JR, Talavage TM, et al. Imaging subcortical auditory activity in humans. *Hum Brain Mapp*. 1998;6:33-41.
112. Cho ZH, Son YD, Kim HK, et al. Observation of glucose metabolism in the thalamic nuclei by fusion PET/MRI. *J Nucl Med*. 2011;52:401-4.
113. Herholz K, Weisenbach S, Zundorf G, et al. In vivo study of acetylcholine esterase in basal forebrain, amygdala, and cortex in mild to moderate Alzheimer disease. *Neuroimage*. 2004;21:136-43.
114. Hirsch F, Sattler B, Sorge I, et al. PET/MR in children. Initial clinical experience in paediatric oncology using an integrated PET/MR scanner. *Pediatr Radiol*. 2013;43:860-75.
115. Schlemmer HP, Pichler BJ, Schmand M, et al. Simultaneous MR/PET imaging of the human brain: feasibility study. *Radiology*. 2008;248:1028-35.
116. Judenhofer MS, Catana C, Swann BK, et al. PET/MR images acquired with a compact MR-compatible PET detector in a 7-T magnet. *Radiology*. 2007;244:807-14.
117. Sander CY, Keil B, Chonde DB, et al. A 31-channel MR brain array coil compatible with positron emission tomography. *Magnetic Resonance in Medicine*. 2014;n/a-n/a.
118. McFarland AR, Newport DF, Atkins B, et al. A CompactPCI Based Event Routing Subsystem for PET and SPECT Data Acquisition. *Nuclear Science Symposium Conference Record, 2006 IEEE*; 2006. p. 3091-3.
119. Atkins BE, Pressley DR, Lenox MW, et al. A Data Acquisition, Event Processing and Coincidence Determination Module for a Distributed Parallel Processing Architecture for PET and SPECT Imaging. *Nuclear Science Symposium Conference Record, 2006 IEEE*; 2006. p. 2439-42.
120. Newport DF, Siegel SB, Swann BK, et al. QuickSilver: A Flexible, Extensible, and High-Speed Architecture for Multi-Modality Imaging. *Nuclear Science Symposium Conference Record, 2006 IEEE*; 2006. p. 2333-4.
121. Byars LG, Sibomana M, Burbar Z, et al. Variance reduction on randoms from coincidence histograms for the HRRT. *Nuclear Science Symposium Conference Record, 2005 IEEE*; 2005. p. 2622-6.
122. Poynton CB, Chen KT, Chonde DB, et al. Probabilistic atlas-based segmentation of combined T1-weighted and DUTE MRI for calculation of head attenuation maps in integrated PET/MRI scanners. *Am J Nucl Med Mol Imaging*. 2014;4:160-71.
123. Catana C, Benner T, van der Kouwe A, et al. MRI-assisted PET motion correction for neurologic studies in an integrated MR-PET scanner. *J Nucl Med*. 2011;52:154-61.
124. Kolb A, Wehrl HF, Hofmann M, et al. Technical performance evaluation of a human brain PET/MRI system. *Eur Radiol*. 2012;22:1776-88.
125. Michel C, Schmand M, Liu X, et al. Reconstruction strategies for the HRRT. *Nuclear Science Symposium Conference Record, 2000 IEEE*; 2000. p. 15/207-15/12 vol.2.
126. Reader AJ, Julyan PJ, Williams H, et al. EM algorithm system modeling by image-space techniques for PET reconstruction. *Nuclear Science, IEEE Transactions on*. 2003;50:1392-7.

127. Bowen SL, Byars LG, Michel CJ, et al. Influence of the partial volume correction method on (18)F-fluorodeoxyglucose brain kinetic modelling from dynamic PET images reconstructed with resolution model based OSEM. *Phys Med Biol.* 2013;58:7081-106.
128. Hudson HM, Larkin RS. Accelerated image reconstruction using ordered subsets of projection data. *Medical Imaging, IEEE Transactions on.* 1994;13:601-9.
129. Kaffanke J, Weirich C, Tellmann L, et al. Investigation of PET count rate reduction during EPI scan on an MR-PET hybrid system. *Proc ISMRM;* 2010. p. 52425.
130. Weirich C, Brenner D, Scheins J, et al. Analysis and correction of count rate reduction during simultaneous MR-PET measurements with the BrainPET scanner. *Medical Imaging, IEEE Transactions on.* 2012;31:1372-80.
131. Singh A, Rathore RKS, Haris M, et al. Improved bolus arrival time and arterial input function estimation for tracer kinetic analysis in DCE - MRI. *Journal of Magnetic Resonance Imaging.* 2009;29:166-76.
132. Meyer E. Simultaneous Correction for Tracer Arrival Delay and Dispersion in CBF Measurements by the H215O Autoradiographic Method and Dynamic PET. *Journal of Nuclear Medicine.* 1989;30:1069-78.
133. Freedman D. SSH/SFTP/SCP For Matlab (v2). *Matlab Central.* 2013.
134. Khan AR, Wang L, Beg MF. FreeSurfer-initiated fully-automated subcortical brain segmentation in MRI using Large Deformation Diffeomorphic Metric Mapping. *Neuroimage.* 2008;41:735-46.
135. Fischl B, Salat DH, Busa E, et al. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron.* 2002;33:341-55.
136. Fischl B, Sereno MI, Tootell RBH, et al. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp.* 1999;8:272-84.
137. Dale AM, Sereno MI. Improved localization of cortical activity by combining eeg and meg with mri cortical surface reconstruction: A linear approach. *Journal of cognitive neuroscience.* 1993;5:162-76.
138. Dale AM, Fischl B, Sereno MI. Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction. *Neuroimage.* 1999;9:179-94.
139. Ségonne F, Dale A, Busa E, et al. A hybrid approach to the skull stripping problem in MRI. *Neuroimage.* 2004;22:1060-75.
140. Fischl B, Liu A, Dale AM. Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. *Medical Imaging, IEEE Transactions on.* 2001;20:70-80.
141. Fischl B, Dale AM. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences.* 2000;97:11050-5.
142. Desikan RS, Ségonne F, Fischl B, et al. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage.* 2006;31:968-80.
143. Destrieux C, Fischl B, Dale A, et al. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage.* 2010;53:1-15.
144. Fischl B, van der Kouwe A, Destrieux C, et al. Automatically parcellating the human cerebral cortex. *Cerebral Cortex.* 2004;14:11-22.
145. Fischl B, Sereno MI, Dale AM. Cortical surface-based analysis: II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage.* 1999;9:195-207.

146. Eggert LD, Sommer J, Jansen A, et al. Accuracy and reliability of automated gray matter segmentation pathways on real and simulated structural magnetic resonance images of the human brain. *PLoS One*. 2012;7:e45081.
147. Klauschen F, Goldman A, Barra V, et al. Evaluation of automated brain MR image segmentation and volumetry methods. *Hum Brain Mapp*. 2009;30:1310-27.
148. Mazziotta J, Toga A, Evans A, et al. A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philos Trans R Soc Lond B Biol Sci*. 2001;356:1293-322.
149. Mazziotta J, Toga A, Evans A, et al. A four-dimensional probabilistic atlas of the human brain. *J Am Med Inform Assoc*. 2001;8:401-30.
150. Mazziotta JC, Toga AW, Evans A, et al. A probabilistic atlas of the human brain: theory and rationale for its development. The International Consortium for Brain Mapping (ICBM). *Neuroimage*. 1995;2:89-101.
151. Tzourio-Mazoyer N, Landeau B, Papathanassiou D, et al. Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain. *Neuroimage*. 2002;15:273-89.
152. Maldjian JA, Laurienti PJ, Kraft RA, et al. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *Neuroimage*. 2003;19:1233-9.
153. Grabner G, Janke AL, Budge MM, et al. Symmetric atlasing and model based segmentation: an application to the hippocampus in older adults. *Med Image Comput Comput Assist Interv*. 2006;9:58-66.
154. Ashburner J, Friston KJ. Unified segmentation. *Neuroimage*. 2005;26:839-51.
155. Diedrichsen J. A spatially unbiased atlas template of the human cerebellum. *Neuroimage*. 2006;33:127-38.
156. Diedrichsen J, Balsters JH, Flavell J, et al. A probabilistic MR atlas of the human cerebellum. *Neuroimage*. 2009;46:39-46.
157. Diedrichsen J, Maderwald S, Kuper M, et al. Imaging the deep cerebellar nuclei: a probabilistic atlas and normalization procedure. *Neuroimage*. 2011;54:1786-94.
158. Eickhoff SB, Stephan KE, Mohlberg H, et al. A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *Neuroimage*. 2005;25:1325-35.
159. Lancaster JL, Woldorff MG, Parsons LM, et al. Automated Talairach atlas labels for functional brain mapping. *Hum Brain Mapp*. 2000;10:120-31.
160. Craddock RC, James GA, Holtzheimer PE, et al. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum Brain Mapp*. 2012;33:1914-28.
161. Dukart J, Mueller K, Horstmann A, et al. Differential effects of global and cerebellar normalization on detection and differentiation of dementia in FDG-PET studies. *Neuroimage*. 2010;49:1490-5.
162. Yakushev I, Landvogt C, Buchholz H-G, et al. Choice of reference area in studies of Alzheimer's disease using positron emission tomography with fluorodeoxyglucose-F18. *Psychiatry Research: Neuroimaging*. 2008;164:143-53.
163. Chonde DB, Abolmaali N, Arabasz G, et al. Effect of MRI acoustic noise on cerebral fludeoxyglucose uptake in simultaneous MR-PET imaging. *Invest Radiol*. 2013;48:302-12.
164. Mandeville JB, Sander CY, Jenkins BG, et al. A receptor-based model for dopamine-induced fMRI signal. *Neuroimage*. 2013;75:46-57.

165. Sander CY, Hooker JM, Catana C, et al. Neurovascular coupling to D2/D3 dopamine receptor occupancy using simultaneous PET/functional MRI. *Proc Natl Acad Sci U S A*. 2013;110:11169-74.
166. Wey H-Y, Catana C, Hooker JM, et al. Simultaneous fMRI-PET of the opioidergic pain system in human brain. *Neuroimage*. 2014;102, Part 2:275-82.
167. Hagler Jr DJ, Riecke L, Sereno MI. Parietal and superior frontal visuospatial maps activated by pointing and saccades. *Neuroimage*. 2007;35:1562-77.
168. Hagler Jr DJ, Saygin AP, Sereno MI. Smoothing and cluster thresholding for cortical surface-based group analysis of fMRI data. *Neuroimage*. 2006;33:1093-103.
169. Becker JA, Hedden T, Carmasin J, et al. Amyloid-beta associated cortical thinning in clinically normal elderly. *Ann Neurol*. 2011;69:1032-42.
170. Greve DN, Svarer C, Fisher PM, et al. Cortical surface-based analysis reduces bias and variance in kinetic modeling of brain PET data. *Neuroimage*. 2014;92:225-36.
171. Kochunov P, Ramage AE, Lancaster JL, et al. Loss of cerebral white matter structural integrity tracks the gray matter metabolic decline in normal aging. *Neuroimage*. 2009;45:17-28.
172. Park HJ, Lee JD, Chun JW, et al. Cortical surface-based analysis of 18F-FDG PET: measured metabolic abnormalities in schizophrenia are affected by cortical structural abnormalities. *Neuroimage*. 2006;31:1434-44.
173. Protas HD, Huang SC, Kepe V, et al. FDDNP binding using MR derived cortical surface maps. *Neuroimage*. 2010;49:240-8.
174. Greve DN, Fischl B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage*. 2009;48:63-72.
175. Feng D, Huang S-C, Wang X. Models for computer simulation studies of input functions for tracer kinetic modeling with positron emission tomography. *International journal of bio-medical computing*. 1993;32:95-110.
176. Feng D, Wang Z. A three-stage parameter estimation algorithm for tracer concentration kinetic modelling with positron emission tomography. *American Control Conference, 1991: IEEE; 1991*. p. 1404-5.
177. Breuer J, Grazioso R, Zhang N, et al. Evaluation of an MR-compatible blood sampler for PET. *Phys Med Biol*. 2010;55:5883-93.
178. Bergstrom M, Bohm C, Ericson K, et al. Corrections for Attenuation, Scattered Radiation, and Random Coincidences in a Ring Detector Positron Emission Transaxial Tomograph. *Nuclear Science, IEEE Transactions on*. 1980;27:549-54.
179. Phelps ME, Hoffman EJ, Mullani NA, et al. Application of Annihilation Coincidence Detection to Transaxial Reconstruction Tomography. *Journal of Nuclear Medicine*. 1975;16:210-24.
180. Bergstrom M, Litton J, Eriksson L, et al. Determination of object contour from projections for attenuation correction in cranial positron emission tomography. *J Comput Assist Tomogr*. 1982;6:365-72.
181. Siegel S, Dahlbom M. Implementation and evaluation of a calculated attenuation correction for PET. *Nuclear Science, IEEE Transactions on*. 1992;39:1117-21.
182. Ciesiński KL, Yang Y, Ay I, et al. Fibrin-targeted PET probes for the detection of thrombi. *Mol Pharm*. 2013;10:1100-10.
183. Burger C, Goerres G, Schoenes S, et al. PET attenuation coefficients from CT images: experimental evaluation of the transformation of CT into PET 511-keV attenuation coefficients. *Eur J Nucl Med Mol Imaging*. 2002;29:922-7.

184. Carney JPJ, Townsend DW, Rappoport V, et al. Method for transforming CT images for attenuation correction in PET/CT imaging. *Medical Physics*. 2006;33:976-83.
185. Bai C, Ling S, Da Silva AJ, et al. A generalized model for the conversion from CT numbers to linear attenuation coefficients. *Nuclear Science, IEEE Transactions on*. 2003;50:1510-5.
186. Catana C, van der Kouwe A, Benner T, et al. Toward implementing an MRI-based PET attenuation-correction method for neurologic studies on the MR-PET brain prototype. *J Nucl Med*. 2010;51:1431-8.
187. Malone IB, Ansorge RE, Williams GB, et al. Attenuation Correction Methods Suitable for Brain Imaging with a PET/MRI Scanner: A Comparison of Tissue Atlas and Template Attenuation Map Approaches. *Journal of Nuclear Medicine*. 2011;52:1142-9.
188. Ashburner J. A fast diffeomorphic image registration algorithm. *Neuroimage*. 2007;38:95-113.
189. Rousset OG, Ma Y, Evans AC. Correction for Partial Volume Effects in PET: Principle and Validation. *Journal of Nuclear Medicine*. 1998;39:904-11.
190. Fulton RR, Meikle SR, Eberl S, et al. Correction for head movements in positron emission tomography using an optical motion-tracking system. *IEEE Transactions on Nuclear Science*. 2002;49:116-23.
191. Bloomfield PM, Spinks TJ, Reed J, et al. The design and implementation of a motion correction scheme for neurological PET. *Physics in Medicine and Biology*. 2003;48:959-78.
192. Schlemmer HP, Schafer J, Pfannenbergl C, et al. Fast whole-body assessment of metastatic disease using a novel magnetic resonance imaging system: initial experiences. *Invest Radiol*. 2005;40:64-71.
193. Chen Q-s, Defrise M, Deconinck F. Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 1994;16:1156-68.
194. Hawkes D, Hill D, Hallpike L, et al. Coregistration of Structural and Functional Images. In: Bailey D, Townsend D, Valk P, et al., editors. *Positron Emission Tomography*. Springer London; 2005:161-77.
195. Eddy WF, Fitzgerald M, Noll DC. Improved image registration by using Fourier interpolation. *Magnetic Resonance in Medicine*. 1996;36:923-31.
196. Friston K, Ashburner J, Frith CD, et al. Spatial registration and normalization of images. *Hum Brain Mapp*. 1995;3:165-89.
197. Friston KJ, Williams S, Howard R, et al. Movement - related effects in fMRI time - series. *Magnetic Resonance in Medicine*. 1996;35:346-55.
198. Maas LC, Renshaw PF. Decoupled automated rotational and translational registration for functional MRI time series data: the DART registration algorithm. *Magnetic Resonance in Medicine*. 1997;37:131-9.
199. Tong R, Cox RW. Rotation of NMR images using the 2D chirp-z transform. *Magnetic Resonance in Medicine*. 1999;41:253-6.
200. Ehman RL, Felmlee JP. Adaptive technique for high-definition MR imaging of moving structures. *Radiology*. 1989;173:255-63.
201. Kouwe AJWvd, Benner T, Dale AM. Real-time rigid body motion correction and shimming using cloverleaf navigators. *Magnetic Resonance in Medicine*. 2006;56:1019-32.
202. Hajnal JV, Myers R, Oatridge A, et al. Artifacts due to stimulus correlated motion in functional imaging of the brain. *Magn Reson Med*. 1994;31:283-91.

203. Power JD, Barnes KA, Snyder AZ, et al. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*. 2012;59:2142-54.
204. Yendiki A, Koldewyn K, Kakunoori S, et al. Spurious group differences due to head motion in a diffusion MRI study. *Neuroimage*. 2014;88:79-90.
205. Thesen S, Heid O, Mueller E, et al. Prospective acquisition correction for head motion with image-based tracking for real-time fMRI. *Magnetic Resonance in Medicine*. 2000;44:457-65.
206. Thesen S. Reduction of Motion Artefacts using Prospective Acquisition Correction for fMRI.
207. Naganawa S, Koshikawa T, Fukatsu H, et al. Whole-brain vascular reactivity measured by fMRI using hyperventilation and breath-holding tasks: efficacy of 3D prospective acquisition correction (3D-PACE) for head motion. *European Radiology*. 2004;14:1484-8.
208. Oakes T, Johnstone T, Ores Walsh K, et al. Comparison of fMRI motion correction software tools. *Neuroimage*. 2005;28:529-43.
209. Maurer Jr CR, Aboutanos GB, Dawant BM, et al. Technical note. Effect of geometrical distortion correction in MR on image registration accuracy. *J Comput Assist Tomo*. 1996;20:666-79.
210. Jezzard P, Balaban RS. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*. 1995;34:65-73.
211. Weisskoff R, Davis T. Correcting gross distortion on echo planar images. 11th Ann Sci Mtg Soc of Magn Reson in Med. 1992:4515.
212. Studholme C, Hill DL, Hawkes DJ. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Medical Physics*. 1997;24:25-35.
213. Collignon A, Maes F, Delaere D, et al. Automated multi-modality image registration based on information theory. *Information processing in medical imaging*; 1995. p. 263-74.
214. Whiteford MH, Whiteford HM, Yee LF, et al. Usefulness of FDG-PET scan in the assessment of suspected metastatic or recurrent adenocarcinoma of the colon and rectum. *Diseases of the colon & rectum*. 2000;43:759-67.
215. Curtis WD, Janin AL, Zikan K. A note on averaging rotations. *Virtual Reality Annual International Symposium*, 1993, 1993 IEEE; 1993. p. 377-85.
216. Moakher M. Means and Averaging in the Group of Rotations. *SIAM J Matrix Anal Appl*. 2002;24:1-16.
217. Sharf I, Wolf A, Rubin MB. Arithmetic and geometric solutions for average rigid-body rotation. *Mechanism and Machine Theory*. 2010;45:1239-51.
218. Volkow ND, Tomasi D, Wang GJ, et al. Effects of low-field magnetic stimulation on brain glucose metabolism. *Neuroimage*. 2010;51:623-8.
219. Volkow ND, Wang GJ, Fowler JS, et al. Resting brain metabolic activity in a 4 tesla magnetic field. *Magn Reson Med*. 2000;44:701-5.
220. Waxman A, Herholz K, Lewis D, et al. Society of Nuclear Medicine Procedure Guideline for FDG PET Brain Imaging Version 1.0. Society of Nuclear Medicine. 2009.
221. Foster JR, Hall DA, Summerfield AQ, et al. Sound-level measurements and calculations of safe noise dosage during EPI at 3 T. *J Magn Reson Imaging*. 2000;12:157-63.
222. Fox PT, Raichle ME, Mintun MA, et al. Nonoxidative glucose consumption during focal physiologic neural activity. *Science*. 1988;241:462-4.

223. Belin P, Zatorre RJ, Hoge R, et al. Event-related fMRI of the auditory cortex. *Neuroimage*. 1999;10:417-29.
224. Cho ZH, Chung SC, Lim DW, et al. Effects of the acoustic noise of the gradient systems on fMRI: a study on auditory, motor, and visual cortices. *Magn Reson Med*. 1998;39:331-5.
225. Schmitter S, Diesch E, Amann M, et al. Silent echo-planar imaging for auditory FMRI. *MAGMA*. 2008;21:317-25.
226. Wu YG. Noninvasive quantification of local cerebral metabolic rate of glucose for clinical application using positron emission tomography and 18F-fluoro-2-deoxy-D-glucose. *Journal of cerebral blood flow and metabolism: official journal of the International Society of Cerebral Blood Flow and Metabolism*. 2008;28:242.
227. Monden T, Kudomi N, Sasakawa Y, et al. Shortening the duration of [18F]FDG PET brain examination for diagnosis of brain glioma. *Mol Imaging Biol*. 2011;13:754-8.
228. Friston KJ, Holmes AP, Worsley KJ, et al. Statistical parametric maps in functional imaging: A general linear approach. *Hum Brain Mapp*. 1994;2:189-210.
229. Fischl B. FreeSurfer. *Neuroimage*. 2012;62:774-81.
230. Maldjian JA, Laurienti PJ, Kraft RA, et al. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *Neuroimage*. 2003;19:1233-9.
231. Yoon H, Park K, Jeong Y, et al. Correlation between neuropsychological tests and hypoperfusion in MCI patients: anatomical labeling using xjView and Talairach Daemon Software. *Annals of Nuclear Medicine*. 1-9.
232. Aydin MD, Ungoren MK, Aydin N, et al. The effects of impulse noise on the epithelial cells of the choroid plexus. *Turk Neurosurg*. 2011;21:191-6.
233. Paus T, Jech R, Thompson CJ, et al. Transcranial magnetic stimulation during positron emission tomography: a new method for studying connectivity of the human cerebral cortex. *J Neurosci*. 1997;17:3178-84.
234. Fox P, Ingham R, George MS, et al. Imaging human intra-cerebral connectivity by PET during TMS. *Neuroreport*. 1997;8:2787-91.
235. Siebner HR, Takano B, Peinemann A, et al. Continuous transcranial magnetic stimulation during positron emission tomography: a suitable tool for imaging regional excitability of the human cortex. *Neuroimage*. 2001;14:883-90.
236. Siebner HR, Peller M, Willoch F, et al. Imaging functional activation of the auditory cortex during focal repetitive transcranial magnetic stimulation of the primary motor cortex in normal subjects. *Neurosci Lett*. 1999;270:37-40.
237. Pietrzyk U, Herholz K, Fink G, et al. An interactive technique for three-dimensional image registration: validation for PET, SPECT, MRI and CT brain studies. *J Nucl Med*. 1994;35:2011-8.
238. Binder JR, Frost JA, Hammeke TA, et al. Human Temporal Lobe Activation by Speech and Nonspeech Sounds. *Cerebral Cortex*. 2000;10:512-28.
239. Lauter JL, Herscovitch P, Formby C, et al. Tonotopic organization in human auditory cortex revealed by positron emission tomography. *Hear Res*. 1985;20:199-205.
240. Wessinger CM, Buonocore MH, Kussmaul CL, et al. Tonotopy in human auditory cortex examined with functional magnetic resonance imaging. *Hum Brain Mapp*. 1997;5:18-25.
241. Signorini M, Paulesu E, Friston K, et al. Rapid Assessment of Regional Cerebral Metabolic Abnormalities in Single Subjects with Quantitative and Nonquantitative [18F]FDG PET: A Clinical Validation of Statistical Parametric Mapping. *Neuroimage*. 1999;9:63-80.

242. Krishna BS, Semple MN. Auditory temporal processing: responses to sinusoidally amplitude-modulated tones in the inferior colliculus. *J Neurophysiol.* 2000;84:255-73.
243. Huffman RF, Henson Jr OW. The descending auditory pathway and acousticomotor systems: connections with the inferior colliculus. *Brain Research Reviews.* 1990;15:295-323.
244. Chonde D, Arabasz G, Catana C. Evaluation of motion correction schemes for simultaneous FDG-PET/MR studies in Alzheimer's disease. *Society of Nuclear Medicine Annual Meeting Abstracts: Soc Nuclear Med;* 2012. p. 208.
245. Slifstein M, Kegeles LS, Gonzales R, et al. [11C] NNC 112 selectivity for dopamine D1 and serotonin 5-HT_{2A} receptors: a PET study in healthy human subjects. *Journal of Cerebral Blood Flow & Metabolism.* 2007;27:1733-41.
246. Fung EK, Planeta-Wilson B, Mulnix T, et al. A Multimodal Approach to Image-Derived Input Functions for Brain PET. *IEEE Nucl Sci Symp Conf Rec (1997).* 2009;2009:2710-4.
247. Mugler JP, Brookeman JR. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). *Magnetic Resonance in Medicine.* 1990;15:152-7.
248. Grinstead J, Rooney W, Laub G. The origins of bright blood MPRAGE at 7 Tesla and a simultaneous method for T1 imaging and non-contrast MRA. *Proc Intl Soc Magn Reson Med;* 2010.
249. Maderwald S, Ladd S, Gizewski E, et al. To TOF or not to TOF: strategies for non-contrast-enhanced intracranial MRA at 7 T. *Magnetic Resonance Materials in Physics, Biology and Medicine.* 2008;21:159-67.
250. Umutlu L, Theysohn N, Maderwald S, et al. 7 Tesla MPRAGE Imaging of the Intracranial Arterial Vasculature: Nonenhanced versus Contrast-Enhanced. *Acad Radiol.* 2013;20:628-34.
251. Hong I, Sanghee C, Casey M, et al. Complementary reconstruction: Improving image quality in dynamic PET studies. *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE;* 2011. p. 4327-8.
252. de Langen AJ, Klabbers B, Lubberink M, et al. Reproducibility of quantitative 18F-3'-deoxy-3'-fluorothymidine measurements using positron emission tomography. *Eur J Nucl Med Mol Imaging.* 2009;36:389-95.
253. Hoekstra CJ, Hoekstra OS, Lammertsma AA. On the use of the injection catheter for venous blood sampling in quantitative FDG PET studies. *Eur J Nucl Med.* 2000;27:1579.
254. Ponto LLB, Md MMG, Richmond JC, et al. Contamination Levels in Blood Samples Drawn from the Injection Intravenous Line. *Molecular Imaging & Biology.* 2002;4:410-4.
255. van der Veldt AA, Hendrikse NH, Smit EF, et al. Biodistribution and radiation dosimetry of 11C-labelled docetaxel in cancer patients. *European Journal of Nuclear Medicine and Molecular Imaging.* 2010;37:1950-8.
256. Mourik JE, Lubberink M, Lammertsma AA, et al. Image derived input functions: effects of motion on tracer kinetic analyses. *Molecular Imaging and Biology.* 2011;13:25-31.
257. Zanolli-Fregonara P, Liow JS, Comtat C, et al. Image-derived input function in PET brain studies: blood-based methods are resistant to motion artifacts. *Nucl Med Commun.* 2012;33:982-9.

APPENDIX

BrainPET Quality Control Experiment Protocols

Perform Spatial Offset Calibration

The calibration can be performed with a fillable Derenzo phantom and should be performed whenever the BrainPET is moved or any service is done to the MR that may alter its isocenter (e.g. whenever the gradients are removed or exposed). This procedure estimates the offset between the two scanners. A relatively long-lived tracer like F-18 is favored to C-11 so the individual PET frames have similar statistics. While only the estimated offsets are saved, at the time they are determined the standard error is also displayed in the desktop workspace. From previous experiments, the translational standard error should be on the order of 0.25mm or less, and the rotational error on the order of radians. To minimize any potential spatial biases, the phantom should be imaged in a number of positions (8 recommended).

4. Fill Phantom with F18
5. Position phantom in CP coil
6. Perform 3-5 minute PET list-mode acquisition with simultaneous MR FLASH scan.
7. Reorient phantom
8. Repeat Steps (2-4)
9. Send Data to Bourget
10. Retrieve MR data using Masamune Organize MR button
11. Remove any FLASH sequences such that there is only one per position
12. Press QC>Alignment to bring up GUI.

Masamune System State Executable Commands

Generating Detector Block Effect File

The block effect file can be generated from scatbe.flor using the following command:

```
>> dos(['scan_sort -F scatbe.flor -S scatbe.n -d -f']);
```

Generating Normalization File

The normalization file can be generated from scatbe.flor using the following command:

```
>>dos(['scan_sort -F C:\bin\norm\psnorm_2014Jul25.flor -N psnorm_2014Jul25.n']);
```

Useful Commands for PET data processing

Sort List-mode Data in LOR/DCmap

Generate Sinogram Data from LOR/DCmap

```
scan_sort -r %5 -L %1.lor.gz -S %1.s -D %1.dcmmap -R %1_ra.s -B dead_blocks.txt
```

Forward Project Image using OSEM

If a volume is forward projected into sinogram space using OSEM, then the volume may be reconstructed using:

```
>>dos(['osem3d_rm -i image.i -W 0 -F 1,0 -l 0 -h image_sinogram.s']);
```

and the reconstructed image can be generated using:

```
>>dos(['osem3d_rm -t image_sinogram.s -P -W 0 -I 32 -N -S 1 -F 1,0 -o output_image.i -
T 8']);
```

Additional arguments may include resolution modeling and normalization if the block effect is added to the data. If the threads are not specified then only 16 subsets may be used. Similarly, OSEM exhibits some instability when run using the maximum 16 threads. Similar stability issues are observed when hyper-threading is active (this may be disabled in the BIOS).

Repairing List-mode Time-Marks

Occasionally a glitch in the scanner during data acquisition may introduce an error in the time-marks. This can be diagnosed by computing the frame length of the time mark using:

```
>>calc_lst_durtation('Frame1.lst')
```

As the scanner begins counting at zero, negative numbers are not permitted and the resulting frame time should be a positive number. If the number is negative, it is likely that a time-mark was miswritten in the data. This can be corrected by reading through each of the time-marks, comparing them to the previous time-mark and making sure they increment by one unit each appearance. This can be performed by running

```
>>fix_lst_t_marks('Frame1.lst')
```