We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

6,900
Open access books available

185,000
International authors and editors

200M
Downloads

154
Countries delivered to

Our authors are among the

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Basic PET Data Analysis Techniques

Karmen K. Yoder

Additional information is available at the end of the chapter

http://dx.doi.org/10.5772/57126

1. Introduction

1.1. Purpose of chapter

In many neuroscience-based PET research labs, procedures for data analyses are developed in-house and passed along as students, staff and post-doctoral fellows transition through training cycles. Although image processing and data analysis techniques are quite similar across many groups, there has not been any formal information available to the general scientific public. This becomes problematic from an instructional standpoint, as the increasingly cross-disciplinary nature of neuroimaging attracts researchers with vastly diverse backgrounds. It is not uncommon to find behavioral pharmacologists, bench neuroscientists, neuropsychologists, and neuroradiologists interested in using neuroimaging techniques for their research. However, often these individuals cannot pursue formal training in PET because of time constraints from other job demands. Although it is easy for seasoned PET researchers to quickly train someone in a laboratory-codified stream of image processing, the "why" of the steps may not get communicated sufficiently, which is a clear disservice to the trainees. This chapter was designed to remedy this problem. The intent of this chapter is to provide a broad foundation of the concepts behind basic PET image processing and data analyses, using data and images from several neuroligands to illustrate key points.

The reader is expected to have a basic working understanding of positron emission, gamma ray generation, and photon detection by the PET scanner.

1.2. Importance of study planning

First and foremost, the scientific question at hand should drive the research process. The first question to answer should be: does your institution have the capability to synthesize or obtain the ligand you need to answer your burning question about neuroscience? If the answer is yes, then the next step is in-depth consultation with the research PET experts at the institution, so

that the study design and data analysis pathway(s) are clearly defined from the outset. The study design, data acquisition protocols, image processing stream, and analysis will differ from study to study, and will depend heavily on both the radioligand and the neurophysiological phenomenon of interest. Types of questions that need to be addressed include (but are not limited to) the following:

What types of data analyses are available/accepted for the tracer? In the clinic, non-quantitative (i.e., visual inspection) of PET images is perfectly acceptable- is the lesion still there? Getting larger? Shrinking? However, in research, there is a requirement for numerical characterization of the dependent variable. For most neuroligand tracers, extensive work has been done to determine what the best and most appropriate approaches are for generating the endpoint of interest. These can range from relatively simple, semi-quantitative methods, to conceptually complex and mathematically rigorous processes that may require additional invasive procedures (arterial cannulation), as well as computational expertise for implementation. Ultimately, the success of a neuroligand PET study will depend on understanding what the field accepts as reasonable outcome measures for a given tracer, and ensuring that the proper infrastructure exists to provide this information.

What type of effect size is expected? This is relevant for determining the number of subjects needed for the study – which, given the great expense of PET, is a nontrivial concern. If possible, it is helpful to know the test-retest reliability of a particular ligand, and to have a general idea of whether your effect of interest is expected to rise above this inherent background noise in the data. In the absence of this, relative variance could be ascertained from previously published data. If no previous documentation exists on the ligand in the species/population of interest, then caution should be used to not over-reach with study design in the beginning. Small pilot studies are very useful at providing initial data on anticipated effect sizes of group, treatment, and/or condition. Study design is a key component of arriving at a sample size: Are the tests to be single measurements between groups (for example, relative receptor availability between healthy normals and a disease condition), or multiple measurements within subjects? Is the tracer known for having either poor or stellar signal-to-noise ratio? All these factors- and others- will affect the ability to detect significant differences.

Group size is not the only consideration- knowledge of the expected spatial extent of the effect is also important. The newer-generation human PET scanners (and most small animal PET scanners) have excellent spatial resolution (1-2 mm^3), but excitement about this technological progress may be mitigated if your hypothesis is restricted to the CA3 region of the hippocampus in humans, or even the whole hippocampus in a mouse. Additionally, the spatial extent of the effect in question will affect the decision to use a region-of-interest based approach versus a voxel-wise analysis (see below).

At this point, hopefully the reader is now familiar with the importance of understanding the type of data that will result from the study, even before the study begins. Although study design is critically important, a thorough discussion of this topic is beyond the scope of this chapter. The remainder of the text will focus on defining concepts and outlining processes for preparing and analyzing neuroligand PET data. Within each subsection, the descriptions will

be presented in a linear fashion. However, the ultimate choices an investigator makes regarding a processing/analysis scheme will depend on multiple factors.

It is important to note that it is not the author's intent to endorse any one particular product or software platform. Examples used here are based primarily on the author's experience, and it is highly likely that many excellent programs are not mentioned. Choices of hardware and software should be made based on investigator preference and availability of individual and/ or institutional licenses.

2. Types of PET data – Definitions and purpose

Dynamic acquisitions: The term "dynamic data" refers to acquiring data in such a manner so that we may observe the long-term behavior of the tracer in tissue. Image acquisition begins immediately upon tracer injection, and the tracer's radioactivity is monitored continuously or near-continuously during the course of the scan. Dynamic data generates "time-activity curves" (TACs) of the tissue concentration of radioactivity (e.g., Bq/mL) over time. Dynamic data acquisition is the only way to obtain truly quantitative measurements of the system of interest. Said another way, "quantitative" measurements are operationally defined by pharmacokinetic and pharmacodynamic properties of a system (for example, B_{max}, K_D). The behavior of the tracer in the system (the TACs) can be described by sets of differential equations; the solutions to these equations yield quantitative outcome parameters. Common parameters of interest include terms such as "volume of distribution" and "binding potential." An excellent review of the definition and derivation of quantitative outcome variables can be found in Innis et al. (2007). In most cases, quantitative outcomes are preferable to semi-quantitative measures (see below). However, quantitative data requires information about the tracer in either arterial plasma of interest, or in a tissue that contains little or no targets of the ligand ("reference region").

Dynamic data can be acquired in two ways. One is by pre-specifying "frame times" for the acquisition, usually of increasing duration (for example, 6 frames at 10 seconds, 12 frames at 20 seconds, 5x60s, 5x120s, 4x300s, 2x600s). The scanner records all the coincidence events that occur during each specified time frame, and the reconstructed image consists of the average amount of radioactivity detected at each voxel during each time frame. The other method is "listmode acquisition", where the scanner records all the coincidence events continuously over time. After acquisition, the investigator specifies how the data should be binned into time frames during reconstruction. Listmode acquisition offers more flexibility for the investigator, especially when the ideal time frame sequence has not been identified. The capability for listmode acquisition varies across scanner platforms.

"Static" Acquisition: In the strictest sense, this refers to specifying one time frame over the course of scan acquisition. The result is a single frame that represents the average amount of radio-activity during the scan period. Only semi-quantitative information can be derived from static acquisitions, the most common of which is Standardized Uptake Value (SUV). SUV is the amount of radioactivity in the tissue (e.g., kBq/mL) divided by the injected dose per body-

weight (e.g., MBq/kg). Static acquisitions are often preceded by a tracer uptake period outside of the scanner environment. Some "static" protocols incorporate a "dynamic" component to facilitate motion-correction (see below), for instance, a 30-minute uptake period followed by five, five-minute frames. Even though multiple time frames are specified, because no image information is captured during the uptake period, this protocol is not considered to be truly "dynamic" data. The intensity values from the acquired frames are typically averaged to generate the mean radioactivity concentration during the scan – the functional equivalent of a "static" scan.

When deciding upon a static versus dynamic protocol, it should be kept in mind that capturing dynamic data leaves open the possibility for quantitative metrics (if the proper methods are available); static acquisition does not. Static images can always be created from dynamic data by calculating the weighted average of radioactivity over a specified set of time frames. However, static data cannot be "undone" into dynamic data.

3. Image processing algorithms – Qualitative description and functions

The advantage of PET imaging is that it provides unique information about the chemistry and physiology of the brain. However, even with high-resolution scanners, PET data often do not contain sufficient neuroanatomic information for identification of specific structures within the brain. The solution to this apparent conundrum is to collect an anatomic Magnetic Resonance Image (MRI) sequence (often an "T1-weighted" sequence) in the same subjects that underwent PET imaging. Having the MRI data confers many advantages to the PET image processing stream, as will be evident below. However, the PET images and MRI images, fresh off the scanner and reconstruction queues, will not be automatically matched up in image space. This is caused by many factors, but the biggest one is differences in final voxel dimensions and final image volume. One of the major objectives of post-processing of PET images is to move the PET and MRI images from the same subject into the same three-dimensional space.

A second main objective of post-processing is motion correction of the PET data. PET acquisitions typically require the subject to try and lay still for as little as 15 minutes, or for up to 90 minutes at a time. It is not uncommon for subjects to move their heads- from coughing, talking, or falling asleep (singing subjects have also been observed). In some protocols, subjects are allowed to get up for a break during the scan acquisition – which automatically means that the PET data will not be in the same exact place in the scanner. Some institutions have developed sophisticated motion-detection and correction systems that work at the level of the reconstruction; however, most investigators do not have access to this technology. Here, we describe a post-hoc method for motion correction after the image has been generated. Because the brain is encased by the skull, there is little concern about movement of the brain within its external bony boundaries. Therefore, the concept of using temporal gating to correct for organ motion, which is a major concern for cardiac and pulmonary imaging, will not be addressed here.

Finally, certain types of data analysis – specifically, voxel-wise analyses – require that all subjects brains be in the same coordinate space. We describe the process that "spatially normalizes" MRI images so that data can be sampled objectively and equivalently across subjects. The processing stream described herein has the goal of translating the PET image(s) into MRI space, so that spatial normalization parameters derived for the MRI likewise can be applied to the PET data.

A wealth of literature and scholarly work has been published on the mathematical basis for algorithms that shift, realign, warp, and reslice three-dimensional images from different modalities so they align correctly. The purpose of this section is to provide a basic, qualitative description of some of these algorithms in context of why they are useful for PET data.

Note: some image processing programs use the terms like "realign" and "co-register" to designate a very specific series of algorithm implementations. To avoid confusion, we will use these terms generically, without attaching any algorithmic meaning to either. We leave it to the reader to investigate the semantics and procedural implementations of a particular program.
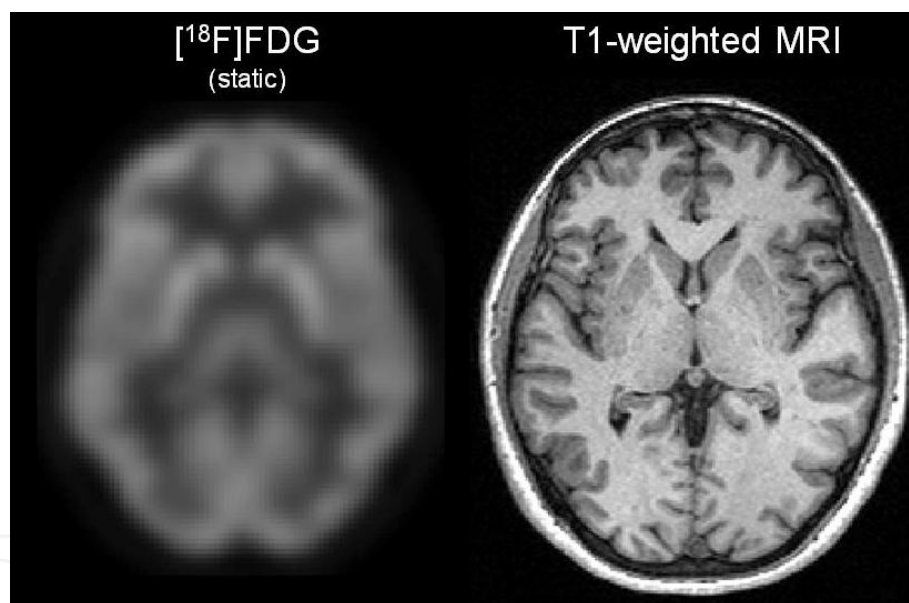

Figure 1. Representative examples of spatially normalized, co-registered images from a healthy subject. Images are axial slices at the level of the striatum and thalamus. Left, a "static" PET image of [18F]flurodeoxyglucose (FDG) Right, the corresponding anatomic T1-weighted MRI. Note that the FDG image contains a high degree of anatomic information that is shared with the MRI (cohesive brain outline and subcortical structure delineation).

Rigid body transformations. Algorithms that perform rigid body transformations are based on the assumption that the rigid bodies (in our case, the PET and MRI image volumes of the same brain) are roughly the same size and geometry. Rigid body transforms only perform translations within object space, they do not allow for "stretching" or "shrinking". To move one object in space to match another's orientation, six parameters are required. Three translations are made along the x, y, and z axes (typically considered right-left, superior-inferior, and anterior-

posterior axes, respectively). Rotations are also made around the three axes; these are called pitch, roll, and yaw.

Rigid body algorithms typically "converge" (that is, come to the final, ostensibly correct answer) fairly quickly, and most co-registrations of PET and MRI are successful. However, the algorithms typically rely on the PET and MRI to share a sufficient amount of contrast and outline among anatomic structures for the alignment to work. In cases in which the PET data "looks" sufficiently similar to the structural MR (Figure 1), the registration process is straight-forward, and the PET and MR can be aligned without additional steps. However, in the case of dynamic data, the tracer distribution and resulting structural information changes significantly over time. Additionally, different tracers will provide varying degrees of structural information (Figure 2). Because of the lack of similarity to the MRI, attempts to co-register individual early or late-time images will likely fail. Here, an intermediate strategy is often successful: create a PET image that shares sufficient features with *both* the MRI and *all* dynamic PET images so that the co-registration algorithm is successful. The general idea is as follows: (1) Create an average or summed image of early PET frames that will share properties with the MR (structural outlines/contrasts) and early and late-time PET data (Figure 3). At this time, performing an alignment or co-registration of the selected subset of PET frames to the first frame is helpful for eliminating spatial variance introduced by motion. The final balance of images to include will be unique to each tracer and frame sequence. Empirical testing is the best way to determine what will be an acceptable combination. For all tracers, early time images are dominated by "blood flow kinetics", or the extraction of tracer from the blood into the tissue. Thus, the early images will trace the general outline of the brain. In the case of tracers like [^{11}C]raclopride and [^{18}F]fallypride, mid-and late time images will be dominated by binding in the striatum, and the brain outline becomes diffuse (e.g., Figure 2a). Inclusion of too many of these striatal images will skew the registration process and should be avoided. For tracers that may not necessarily have a lot of tracer retention (e.g., amyloid in the example of [^{11}C]PiB, inflammation for [^{11}C]PBR28), the entire set of dynamic images may be needed to generate a sufficiently robust brain PET image. (2) Co-register the mean PET to the native-space structural MRI. Make sure the transformation parameters have been saved in the header files of the resliced PET. (3) Co-register all the dynamic PET frames to the co-registered mean PET (which is now in native MR space). Because all the frames are being registered to the same target, this step has the convenient function of also providing a robust method for motion correction. Additional refinements for motion-correction may be needed in cases in which the motion may be too severe to be corrected by a rigid-body algorithm alone. Occasionally, "manual" repositioning of a timeframe (meaning, the user specifies the translations to change the orientation) can be used to provide the registration algorithm with a better "initial guess." It is our experience that manually adjusting the position of a poorly aligned timeframe and re-running the algorithm can result in a successful motion correction. Representative time-activity curves from before and after manual manipulation of two errant time frames are given in Figure 4.
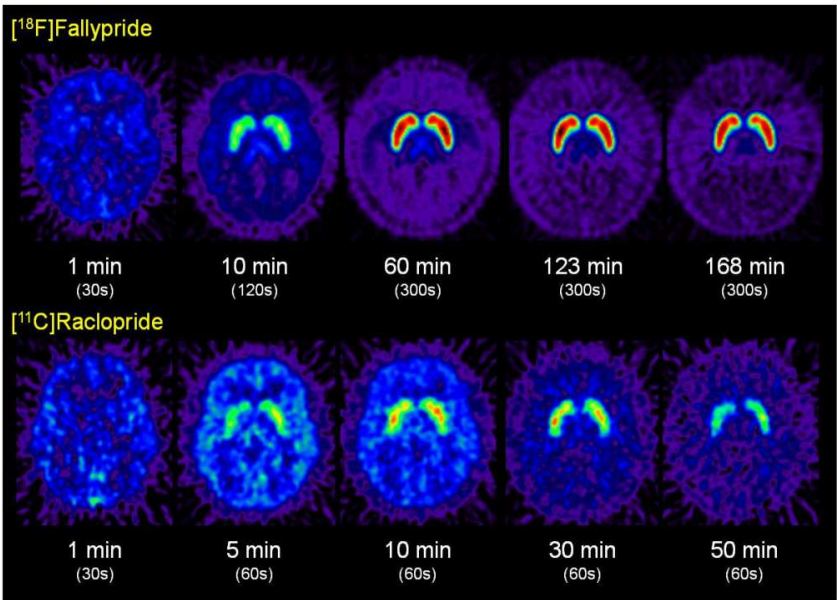
(a)
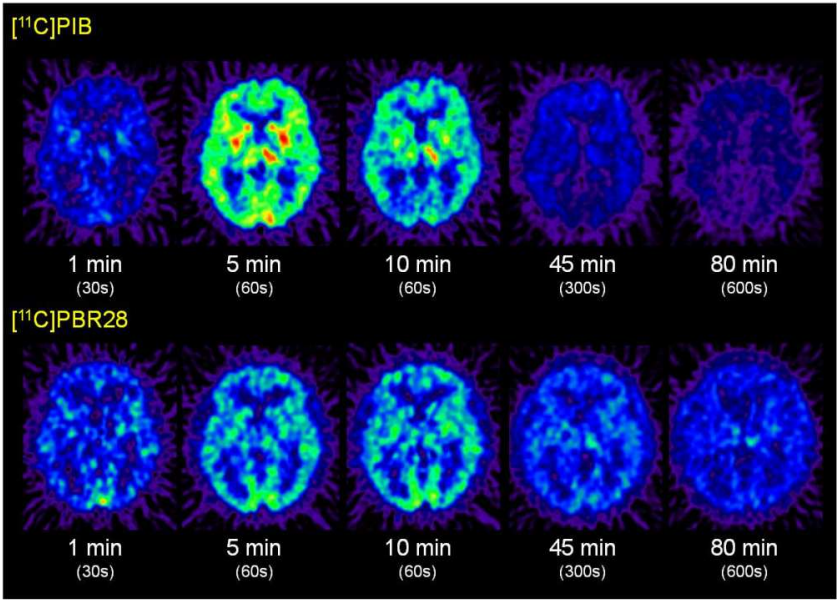
(b)

Figure 2. (a) Differential behavior of tracers over time. Multi-panel figure of early, mid, and late time frames from [^{18}F]fallypride (top panels) and [^{11}C]raclopride (bottom panels). Both tracers are dopamine D_2/D_3 antagonists; each has different kinetic and signal-to-noise properties. Underneath each panel: start time of each frame relative to tracer injection (min), and duration of each frame (s). **(b)** Differential behavior of tracers over time. Multi-panel figure of early, mid, and late time frames from [^{11}C]PiB in an Alzheimer's Disease subject (top panels) and [^{11}C]PBR28 in a healthy elderly control (bottom panels). [^{11}C]PiB binds to β-amyloid plaques, one of the primary pathological hallmarks of Alzheimer's Disease. The majority of healthy elderly subjects have no discernable [^{11}C]PiB uptake. [^{11}C]PBR28 binds to the Translocator Protein 18kDa, a mitochondrial marker associated with inflammation. There is some degree of consistent [^{11}C]PBR28 brain uptake in healthy subjects; the pathological patterns of [^{11}C]PBR28 in neurological and psychiatric disease are not yet well-understood. Underneath each panel: start time of each frame relative to tracer injection (min), and duration of each frame (s).
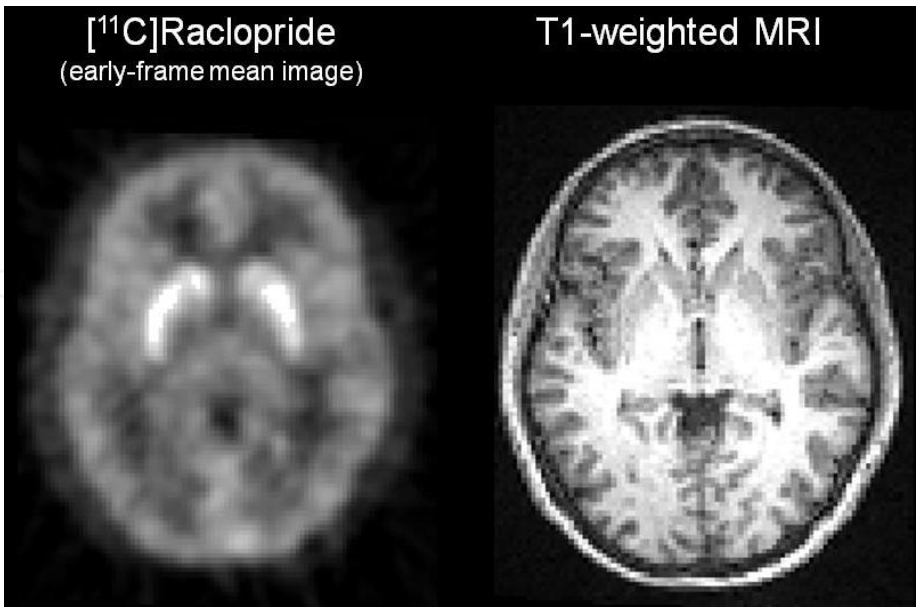
Figure 3. Example of how a mean dynamic image can be used to facilitate successful co-registration with an anatomic MRI. Left, spatially normalized "early" mean PET, consisting of the first ~10 minutes of dynamic [^{11}C]raclopride data. The number of frames required to achieve a balance of flow/binding for co-registration with the MRI depends heavily on the individual tracer kinetics, and must be determined empirically by each investigator for the particular tracer and acquisition sequence. This particular combination happens to work well for [^{11}C]raclopride. Note the general similarity to the FDG scan in Figure 2. Right, corresponding spatially normalized MRI from the same subject.
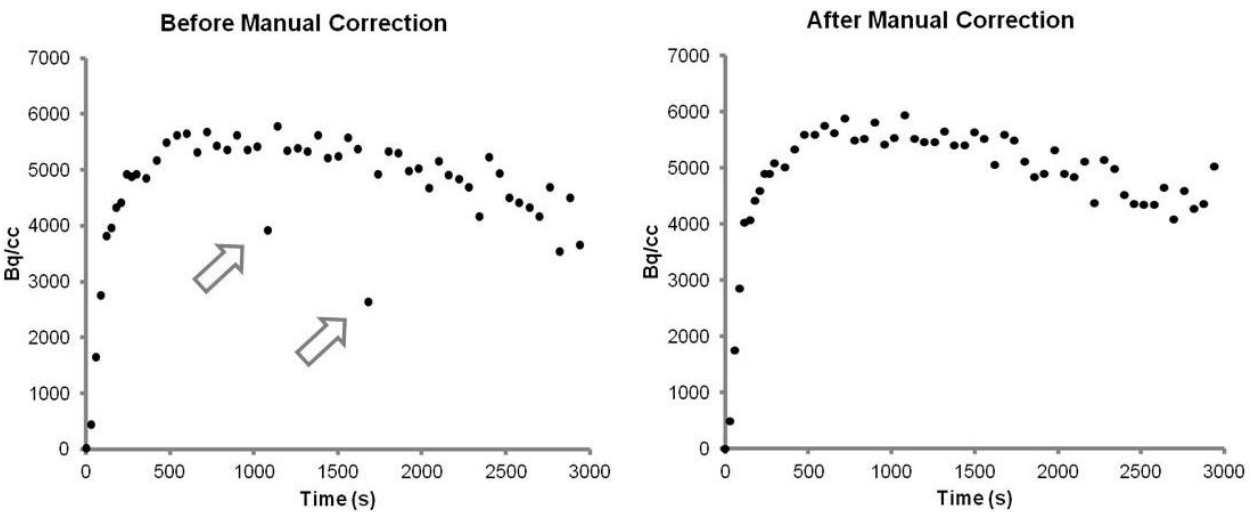

Figure 4. Time-activity curves (TACs) from a [^{11}C]raclopride scan with and without manual motion correction. Left: TAC from the right putamen of a subject after initial automated motion correction was conducted. Subject motion was severe enough that at least two frames could not be corrected by the algorithm (arrows). Right: resultant TACs after several frames were re-oriented manually and the co-registration algorithm was re-run. The improved initial guesses given by the manual manipulation resulted in better convergence for the algorithm and much smoother curves. This illustrates the need for use of TACs to check for motion in addition to the use of cine loops. It also illustrates the advantage of shorter time frames for capturing motion artifacts.

Nonlinear transformations. Nonlinear transforms are most commonly used to "warp" anatomic MRI brain images into a common stereotaxic coordinate space. This is necessary when a "voxel-wise" approach for data analysis is desired (see below). Typically, each individual subject brain is "warped" to a canonical template brain, typically supplied by the program that hosts the spatial transformation algorithm. However, canonical templates may not be the best representation of a given population sample, especially in patient populations that have unique structural disorders. Many investigators prefer to generate study-specific unique templates; this is definitely desirable in animal studies. Approaches to creating templates range from simple averaging of MRI (or PET) data across a sample to more sophisticated approaches that carefully map subject brains onto an existing defined coordinate system (e.g., Schweinhardt et al., 2003). It is up to the investigator to determine the optimal approach for their respective study.

If the investigator intends to rely on region-of-interest (ROI) analysis based on subject-specific ROIs from native-space MRIs, then this step may not be needed. Our laboratory uses a combined approach for ROI analysis (see below).

It should be noted that while the working assumption is that the deformations applied to subject data will render the brain totally "warped" to the template, not all individual variation in anatomy is lost. This should be taken into consideration when interpreting voxel-wise analyses (see below), or using template or group-averaged normalized MRs as starting points for ROIs.

4. Partial volume effects and partial volume correction

The term "partial volume effect" (PVE) and "partial volume correction" have become general terms in neuroimaging. However, it can mean very different things to MRI and PET experts. Even in PET, there can be confusion about between PVE and "spill-out/spill-in" effects. Therefore, definitions are warranted to prevent confusion.

Spatial Resolution Effect: This is often what is referred to as PVE. However, the term "spatial resolution effect" is more accurate. PET does not provide a spatially pristine representation of the radioactivity in the tissue – it is a "fuzzy" picture of the true concentration of radioactivity. This can be especially problematic when attempting to measure radioactivity in very small or very thin structures that are smaller than the inherent resolution of the scanner. For example, imagine a very small, very "hot" object (like a 1 mm sphere) that is surrounded by tissue that contains no radioactivity. If the intrinsic resolution of the scanner is e.g., 5 mm (please see Phelps (2006) and Bailey (2005) for details on how resolution is defined and specified), the geometry and properties of the scanner will inevitably blur the apparent concentration of radioactivity, effectively assigning spatial location of the radioactivity that originated from the object to the surrounding tissue. The amount of radioactivity measured in the object will be underestimated (from "spill-out"), and surrounding tissue will appear to have radioactivity ("spill-in"; please refer to Morris et al., 2004, for an excellent illustration and mathematical explanation of this phenomenon). Several strategies exist to correct for spatial resolution

problems. They typically are both computationally and labor-intensive, and require a very detailed anatomic MRI image and robust a priori knowledge of the tracer distribution (Mawlawi et al., 2001; Morris et al., 2004). In deciding on whether or not to apply correction for spatial resolution, the key question is: How important is absolute quantitation? Does having absolute certainty of radioactivity concentration increase the ability to detect differences between groups or conditions? In many cases, it may be reasonable to assume that variability contributed by spatial resolution effects is homogenous across/within subjects, and therefore spatial resolution correction is not warranted. However, in cases where "spill-out" drastically reduces the dynamic range of signal, or interferes with the ability to detect signal above background (such as in small structures, especially in rodent PET studies), spatial resolution correction may be an important option to consider.

Partial Volume Effect: True PVE actually is a problem of tissue heterogeneity within a "volume." A volume could be a voxel, or a large region of interest that spans many different tissue types (e.g., a brain lesion). Even with MRI's superior spatial resolution, MRI voxels that sit on the borders between gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF) may actually contain components of more than one type of signal. In MRI, this problem is addressed (in part) with probabilistic "segmentation" algorithms that assign voxels to either GM, WM, or CSF. These algorithms create tissue-specific maps, which are useful for many purposes, including creation of anatomic masks which can be used to restrict spatial extent of voxel-wise statistical analyses. Gray matter maps are also a good starting point for generation of subject-specific ROIs.

The source of concern of PVE in PET focuses mainly on quantitative analyses. Regardless of tracer, GM, WM, and CSF will have inherently different kinetics (although CSF does not have "kinetics" per se). This heterogeneity would necessitate accounting for multiple sets of tracer behaviors, complicating and potentially confounding quantitation via mathematical modeling. However, tissue heterogeneity in neuroligand PET data is typically not addressed. This is in part because scanner resolution has improved significantly, and in part because PET processing and analyses rely heavily on structural information from the MRI, which helps restrict the analyses to specific structures/tissue types.

5. Data analyses

So, the PET studies have been designed and data have been collected. Now what?

Quantitative Analyses. If dynamic data were acquired, it is possible to get quantitative information about the ligand in the brain- provided that an "input function" and software implementation of the proper tracer kinetic models are available. The two most common parameters of interest are "Volume of Distribution" (V_T) and "Binding Potential" (BP), both of which provide physiologically relevant data regarding tracer retention. Input functions are necessary to drive the tracer kinetic modeling procedures – they provide key information about free tracer concentration in the plasma (or parameters related to it). There are three main types of input functions: arterial plasma, reference region, and image-derived (for example, informa-

tion from the carotid artery, left ventricle of the heart, or even lungs). Obtaining quantitative endpoints with tracer kinetic modeling and arterial plasma input functions is the "gold standard." The choice of either reference region and image-derived inputs must be substantiated by the literature and validated by extensive testing by kinetic modeling experts. Once time-activity curves (TACs) from both your input function and tissue of interest are at hand, the parameter estimation can begin. Explanation of the types of kinetic models, the assumptions of each, the parameters they yield, and advantages/disadvantages of each are beyond the intent of the present text. Regardless, the investigator should be aware that different model implementations may behave differently with different tracers. Consult with your local PET modeling expert to determine which methods are most appropriate and most convenient.

Semi-quantitative. This discussion refers to "static" images (see above). The voxel values within the PET image are in values of tissue radioactivity concentration, such as Bq/cc or kBq/mL. However, taken in isolation, these values are not meaningful and cannot be used as the final endpoint for analysis. Too many factors affect the radioactivity concentration, including the total dose injected, and the body weight of the individual. At a minimum, the data must be normalized to account for injected tracer dose.

The most common method used for data normalization is the index of "Standardized Uptake Value" (SUV), in which the radioactivity concentration is divided by injected dose per body weight (e.g., MBq/kg). This index comes with one major assumption, which is that the tracer has been distributed equally across the entire body- that is, all tissues have had an equal opportunity to be exposed to the tracer. If a "sink" for the tracer exists outside the target of interest, such that a great amount of tracer is sequestered during first pass circulation, then the tracer is not being distributed equally across the body. The whole-body distribution assumption has then been violated, and body weight is no longer the proper denominator. In this case, SUV measurements are rendered incorrect, and become unreliable as a dependent variable. Another type of normalization with SUV is "SUVR", which is the ratio of SUV from tissue that has specific binding of the tracer to tissue that does not (a reference region). This method was proposed for [^{11}C]PiB, and was evaluated thoroughly for this tracer against arterial and reference region kinetic approaches (Lopresti et al., 2005). If investigators are using a relatively new neuroligand and seek to use SUV or SUVR as endpoints, it is highly recommended the stability of the semi-quantitative index be assessed against either V_T or BP, either with real data or via simulation studies.

What kind of analysis do you want? Region-of-Interest versus Voxel-Wise.

Anatomic ROI Analyses

Broadly speaking, a region-of-interest (ROI) refers to a user-defined set of voxels (or voxel) on an image, from which PET data are extracted. First, we will address anatomically-defined ROIs.

Using anatomically-defined ROIs remains a popular approach for analyzing neuroligand PET data. Typically, anatomic ROIs are defined on a subject's MRI, and then transferred to the dynamic PET data (which is in register with the MRI). An average time-activity curve for the ROI is generated (that is, the time-activity curves of all voxels within the ROI are averaged), and the TAC is fed into a model for estimation of V_T or BP.

Pros: The anatomic ROI approach is a good choice when the study hypothesis anticipates that the effect of group (or condition) with be consistent across the entire anatomic extent of the structure (that is, the ROI is chemically and/or functionally homogeneous). Additionally, if the ROIs are reasonably sized, the TACs are usually smooth. This results in more robust parameter estimates (noisy time-activity curves typically induce a negative bias, or lower values) (see below, and Figure 5).
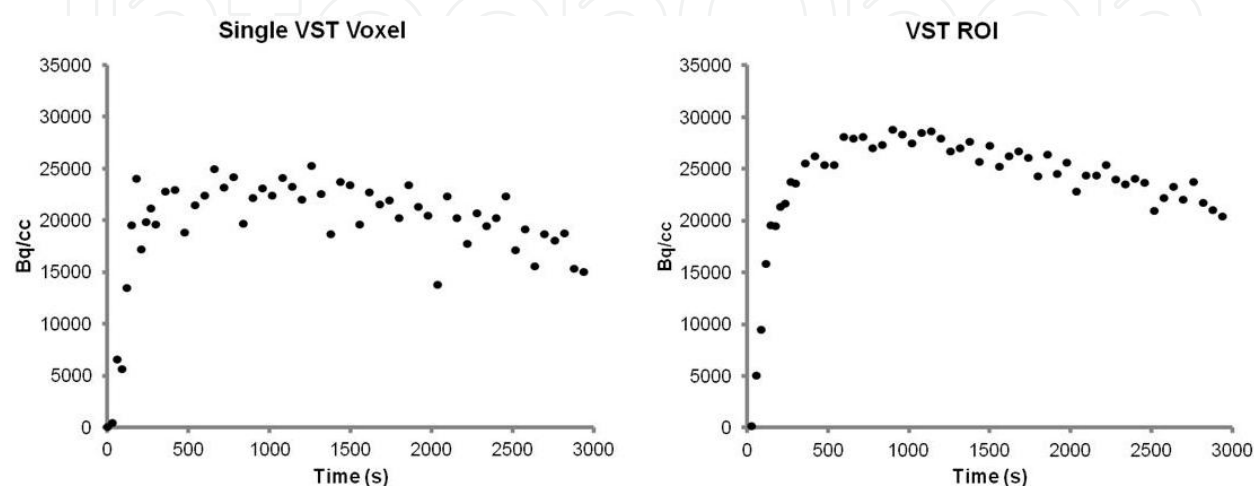

Figure 5. Time-activity curves (TACs) of [^{11}C]raclopride from, Left: a single voxel in the left ventral striatum (BP$_{ND}$ = 2.09) and Right: from the whole left ventral striatum region-of-interest (ROI) (BP$_{ND}$ = 2.74). BP$_{ND}$ values were estimated with MRTM (Ichise et al., 2003), using the same cerebellar input function. Note that the average TAC from the whole ROI is much smoother than the single voxel TAC. The slight difference in intensity scale of the single voxel could be attributed to its location in the more ventral aspect of the striatum, close to the base of the brain, which makes it more susceptible to "spillout" artifact.

Cons: ROI analyses may miss subtle effects that are spatially constrained to a small area within the larger ROI. If an effect is only present in a subset of voxels, then this may get lost (smoothed out) when all the TACs from all voxels are averaged together.

There are many ways to generate anatomic ROIs. Often, ROIs are painstakingly drawn by hand, which is labor- and time-intensive (and fairly boring for the individual charged with this task). This also has the risk of inducing subjective bias to the ROI definition, with the ensuing possibility data may not be completely comparable across institutions. However, adherence to strict and consistent anatomic definitions based on accepted atlas(es) (e.g., Martinez et al., 2003; Mawlawi et al., 2001) helps mitigate any investigator-induced bias. There are many software programs that offer sets of pre-defined ROIs, which are often defined from a single-subject MR. In our experience, these ROIs are not very representative and do not match well to our subject samples. We have also found that ROIs drawn by our lab on "canonical" average multi-subject T1 templates (again, available in many software packages) do not conform well to our subject samples. Yet another option is to utilize sophisticated software that automatically extract hundreds of ROIs by parcellation of a subjects' MRI. Our laboratory uses a combined approach, in which we start with an individual subjects' spatially normalized gray matter map and a template ROI (e.g., ventral striatum) generated from an average MR

from our subject sample. These two sources are combined to generate a "starting point" MRI, which is then edited to explicitly conform to an individual's subject anatomy. Regardless of chosen method, the investigator should take care to ensure that the anatomic ROIs are spatially appropriate for each individual subject.

Choice of statistical analysis of ROIs depends on the study design: independent t-tests, paired t-tests, one-way ANOVA, mixed effects models, ANCOVA, correlations, etc. Regardless of the test, when multiple ROIs are being tested for between-group or between-condition effects, or for correlations with e.g., a particular subject characteristic, there is always the question of whether the results need to be corrected for multiple comparisons. This is a relevant but somewhat controversial issue. It is indeed the case that multiple comparisons can lead to false positive results (Type I error), and results that survive statistical adjustments for the multiple tests (e.g., a Bonferroni correction) can help assure the investigator that the effects are real. However, arguments have been made that in pilot studies and/or with exploratory data, such corrections are overtly stringent and unwarranted (Perneger, 1998). Investigators should be prepared to justify omission of correction for multiple comparisons based on the exploratory nature of the study and/or sample size.

Voxel-Wise Analyses (voxels are ROIs, too)

Technically speaking, a voxel is the smallest ROI that is possible within an image. Voxel-wise analyses assume that all subject brain data are in the same coordinate space (see above). Voxel-wise studies demand "parametric images", that is, voxels cannot be in units of radioactivity concentration, but must be converted to either a quantitative (e.g., V_T or BP) or semi-quantitative (e.g., SUV) value. (In this context, "parametric" simply carries the general meaning of a uniformly normalized or an explicit physiologically descriptive value, and should not be confused with "parameter estimation" used to describe the process of kinetic modeling). In the case of quantitative values, the parameter of interest is generated based on the time-activity curve for each voxel (the input function is the same for all voxels). Taking a page from MRI processing procedures, many investigators will spatially smooth the parametric images to remove any spuriously high or low voxel values. The smoothing kernel should be roughly the size of the practical resolution of the PET scanner (not the ideal, intrinsic resolution). Statistical models are specified based on study design, and statistical testing is performed at each voxel. Most image analysis packages include the flexibility to specify different statistical thresholds, which allows investigators to interrogate the data for subthreshold effects. They also have the capacity to apply stringent corrections for a true multiple comparisons problem: performing statistical tests at tens of thousands of voxels across the brain simultaneously. Areas of significant results are shown as "clusters" (groups of contiguous voxels).

Although first-pass voxel-wise analyses does not necessarily have to correct for multiple comparisons, there may be logical reasons to spatially restrict the initial voxel-wise analyses. If the tracer is only anticipated to have specific binding in gray matter, use of an average gray matter mask (derived from the sample) would be appropriate to exclude WM and CSF voxels. [^{11}C]raclopride (a dopamine D_2/D_3 antagonist) is another good example- the signal-to-noise properties of this tracer are such that it cannot be used to quantitate D_2/D_3 receptor binding in areas outside the striatum (which has the highest concentration of D_2/D_3 in the brain). In our

laboratory, we use a striatal mask to restrict the search area to the striatum. However, with tracers that can bind to processes that are not restricted to gray or white matter, whole-brain sampling would be more appropriate (unless the investigator has an a priori hypothesis that targets a specific region). An example of this would be [^{11}C]PBR28, which is a marker of neuroinflammatory processes. Additionally, if the investigator has specific a priori hypotheses about a structure of interest, it is reasonable to use an anatomic ROI to restrict analysis to a particular nucleus or cortical area. By now, the reader should appreciate that the distinctions between voxel-wise and ROI analyses begin to blur a bit.

Pros: Voxel-based analyses have two main advantages over ROI analyses. First, they sample the entire brain (or a spatially restricted region) objectively. Second, a voxel-wise approach they can pick up spatially discrete areas of effect that may be "washed away" by an ROI approach.

Cons: TACs at the voxel level can be extremely noisy (Figure 5). In many kinetic models, noise can cause underestimation of the quantitative parameters. However, if subjects have received approximately the same dose of tracer, then the noise at the voxel level should be uniform across subjects. Gray-matter voxels that share boundaries with CSF, air (sinuses), or white matter are especially subject to spatial resolution problems (see above); results that seem to outline "edges" of ventricles or other structures near white matter should be considered with caution. Also, even including one or two subjects with severe atrophy not corrected for by spatial normalization can skew voxel-wise results. Again, this brings to light the care an investigator must take with understanding the data (see QC section, below).

Congratulations, you've got statistically significant results with a voxel-wise analysis. Now what?

The output of most voxel-wise analyses is a series of parametric maps with *t*-statistic values at each voxel where there was a statistically significantly effect at a given *p*-threshold. However, the programs typically don't give you direct information about the actual effect size you are detecting. In addition to the statistical results, investigators should report the quantitative description of the data such as what the percent change was between groups/conditions, what the nature of the correlation was, etc.

Many analysis programs will allow you to save out a cluster of significant voxels as an ROI; single voxels may even be used as ROIs (this would be useful for characterizing peak effects). You may also choose to use a predefined anatomic ROI, especially when the effect of interest spans areas of interest, or the region was part of an *a priori* hypothesis. Once the ROI has been designated, there are two main methods for extracting the data. The ROIs can be applied to all subject's parametric images, and then the mean values (e.g., V_T, BP, SUV) for the ROI can be compiled across groups, conditions, etc. Alternatively, the ROIs can be applied directly to the dynamic data to extract the average TAC from the ROI, which is then fed into a modeling program to estimate the parameter of interest (as described above).

6. Quality control and automation

Many of the processes discussed here involve computer-based procedures. However, it is unwise to assume that the algorithms will work perfectly and that the data will always be robust. In order to assure the quality of the study, quality control by real humans is required- at every point along the processing and analysis stream. Simple visual checks can be made to determine the success of the co-registration, motion correction, and normalization steps. Here is a sample checklist:

1. MR-PET co-registration: Is the mean/summed PET in the same space as the native MR? Multiple anatomic landmarks should be assessed- from outer cortical layers to subcortical landmarks such as striatal boundaries, ventricles, and corpus callosum. Cortical land-marks may be the best visual assessment for PET studies that do not contain much subcortical binding (for example, [11C]PiB in healthy controls)

2. MR normalization. The "warping" of the native space MR to the target coordinate space should be checked against the coordinate template. This step does not always converge appropriately, and some very strange brains can result from incorrect convergence of the spatial normalization algorithm. A helpful step to ensure successful normalization is to first perform a rigid-body co-registration between the subject native space MR and the canonical template, then perform the spatial normalization step.

3. PET normalization. To perform this QC step, create a mean or summed PET from all the spatially normalized dynamic PET frames (or spatially normalized single static frame), and compare it to both the subject's normalized MR and the template MR. If the MR normalization step was successful, then in most cases the PET normalization will be fine – but this should never be taken for granted.

4. Motion-correction. Programs that read in multiple 3D volumes in a cine loop are extremely useful for visually detecting motion. Make sure that the program can read in the numerical convention of sequential images (i.e., decimal or hexadecimal). Investigators must learn to distinguish between true anatomic subject motion and the random noise inherent in PET images- which can create an optical illusion of motion.

5. Anatomic regions of interest. This QC should be done during the generation of the regions of interest, but is still a crucial trouble-shooting step, especially if a subject's data appears to be a high or low outlier relative to the sample. It is extremely important to check the overlay of the ROIs on the PET data. Reslicing of the PET data during the motion correction and spatial normalization can result in "chopping off" of brain regions, especially cerebellum and frontal cortex. It is very important to make sure the anatomic ROIs are not sampling image "air". This is especially important for quantitative analyses that utilize BP with a "reference regions" (see above). If the reference region is corrupted by white matter, CSF, or "air", then the BP estimate of the target region (or voxel) will be corrupted.

6. Time-activity curves. Even if the investigator intends to only run a voxel-wise analysis, it is always a good idea to visually check the TACs from at least one, if not several, anatomic

regions of interest. This may help identify motion that needs to be corrected, or may shed light on other data quality problems that should be addressed.

7. Parametric images. Regardless of outcome variable (e.g., SUV V_T, BP), the parametric images should be examined to make sure the values are reasonable. If outlying values are observed, then more intensive investigation is warranted to identify the source of the (apparently) aberrant data.

Executing all the image processing steps individually can be time-consuming and labor-intensive. With many programs, it is possible to automate, or "batch", many steps together through the use of scripts. In fact, automation is often implemented at the level of multiple subjects at once. The degree of automation implemented must be up to the discretion of the laboratory and what works best with the current laboratory culture, which encompasses both available manpower and study completion rate. Some labs may run hundreds of subjects, then perform QC steps for each step in batches for each QC point. Other labs may choose to implement QC for each individual subject as those data come through. Workload distribution of QC is ultimately up to each investigator, based on their needs.

7. Small animal PET considerations (especially rodent)

In general, the same principles described above regarding types of data (semi-quantitative, quantitative) and analyses (ROI, voxel-wise) apply to PET imaging in small animals. Having said that, some special concerns need to be addressed (or alleviated). In most studies, animals will be anesthetized during imaging. If the animal is restrained by a device that prohibits head motion (e.g., for neuroimaging, a stereotaxic head-holder), then motion correction for the dynamic PET data may not be needed. (Again, gating acquisition methods for thoracic and abdominal imaging are beyond the scope of this discussion). If the animal's skull is not explicitly restrained, then head motion may occur from breathing, and motion-correction algorithms may be warranted.

If one acquires parallel data in other modalities for the purposes of co-registration, the nature of the PET data must be considered within the context of aligning one modality to another. Tracer kinetics in rodents can be vastly different than what is observed in humans. Some tracers may have very little apparent brain uptake, and therefore the outline of the brain may not be obvious. If there is little information about brain shape in the PET data, a co-registration algorithm may not work accurately, or may even crash. If alternate modality images are acquired and needed for co-registration (e.g., CT for attenuation correction; MR for anatomic localization of anatomic structures), then image editing may be required for successful co-registration between the PET data and MRI and/or CT. For example, if the lack of a coherent brain outline in the PET image data is problematic, it may be useful to edit out the skull in CT or MR images. On the other hand, in rodents, tracers with high uptake areas (e.g., dopaminergic ligands in the striatum) may not register well to a rodent MR, which will not have clear delineation between subcortical nuclei in rodents. In this case, an early-time PET image may be useful for registration. Finally, if the animal is restrained with a headholder, the extra image

information from the headholder material may need to be edited out to render a more purely anatomic image and facilitate co-registration with the PET (which will not show the presence of the holder). Again, there is no "set" approach, and it is up to the investigator to empirically determine what is most appropriate for their particular imaging system.

Regardless if an ROI or voxel-wise approach is used for small animal PET data analysis, it is important that the investigator appreciate the practical resolution limitations of the imaging modality. Even with the most advanced small animal PET scanners, it is difficult to resolve structures below a ~1mm³ volume. This is true even if high-resolution anatomic MR data are available. Two key concepts are worth emphasizing:

1. The presence of micron resolution in an MR image co-registered to a PET image does not guarantee micron resolution in the PET dataset (see above section for discussion on partial voluming artifacts).

2. The voxel size in a PET image does not necessarily correspond to the practical resolution of the scanner. Images can be resliced almost infinitely to very small voxels by relatively simple interpolation of the original PET data. However, this does not change the actual resolution of the scanner. If the final PET voxel size is one 0.25mm x 0.25mm x 0.25mm, and the practical scanner resolution is 1mm x 1mm x 1mm, then the resolution is still 1mm x 1mm x 1mm. Information is not gained by voxel sizes that are smaller than the intrinsic scanner resolution. Check with your local PET expert to determine what the resolution of the small animal scanner is.

Voxel-wise studies in small animal PET data require careful consideration with respect to how spatial normalization will be achieved, especially with rodent data. Brain structure between rodent strains is likely to be quite different; it should not be assumed that one-rat (or mouse) brain-fits-all. Additional factors like gender, age, and weight of the animal also influence brain shape and structure. Of particular note is that brain development and growth in rodents is nonlinear across structures (Sullivan et al., 2006), and therefore a younger brain should not simply be scaled up to the size of an older rodent brain. Three-dimensional templates and atlases of mice and rat brains are becoming more common. However, when possible, the investigator should consider generating an in-house brain template specific to the strain, age, gender, and weight of the sample being studied.

8. Summary

Image processing and data analysis of neuroligand PET data requires multiple steps. There are an almost infinite number of possible iterations and refinements to data processing streams. Hopefully, this chapter has provided a useful overview of the key concepts investigators need to consider when working with these expensive and often complex datasets. Regardless of the exact sequence of processing procedures selected, a thorough working knowledge of the rationale behind each step will help ensure the fidelity and quality of the laboratory's datasets.

Acknowledgements

Author details

Karmen K. Yoder

Indiana University School of Medicine, Department of Radiology and Imaging Sciences, Indiana Institute for Biomedical Imaging Sciences, Center for Neuroimaging, Indianapolis, Indiana, USA

References

[1] Bailey DL. 2005. Data acquisition and performance characterization in PET. In: Bailey DL, Townsend DW, Valk PE, Maisey MN, editors. Positron emission tomography: Basic sciences. London: Springer-Verlag. p. 41-62.

[2] Ichise M, Liow JS, Lu JQ, Takano A, Model K, Toyama H, Suhara T, Suzuki K, Innis RB, Carson RE. 2003. Linearized reference tissue parametric imaging methods: application to [11C]DASB positron emission tomography studies of the serotonin transporter in human brain. J Cereb Blood Flow Metab 23(9):1096-1112.

[3] Innis RB, Cunningham VJ, Delforge J, Fujita M, Gjedde A, Gunn RN, Holden J, Houle S, Huang SC, Ichise M, Iida H, Ito H, Kimura Y, Koeppe RA, Knudsen GM, Knuuti J, Lammertsma AA, Laruelle M, Logan J, Maguire RP, Mintun MA, Morris ED, Parsey R, Price JC, Slifstein M, Sossi V, Suhara T, Votaw JR, Wong DF, Carson RE. 2007. Consensus nomenclature for in vivo imaging of reversibly binding radioligands. J Cereb Blood Flow Metab 27(9):1533-1539.

[4] Lopresti BJ, Klunk WE, Mathis CA, Hoge JA, Ziolko SK, Lu X, Meltzer CC, Schimmel K, Tsopelas ND, DeKosky ST, Price JC. 2005. Simplified quantification of Pittsburgh Compound B amyloid imaging PET studies: a comparative analysis. J Nucl Med 46(12):1959-1972.

[5] Martinez D, Slifstein M, Broft A, Mawlawi O, Hwang DR, Huang Y, Cooper T, Kegeles L, Zarahn E, Abi-Dargham A, Haber SN, Laruelle M. 2003. Imaging human mesolimbic dopamine transmission with positron emission tomography. Part II: amphetamine-induced dopamine release in the functional subdivisions of the striatum. J Cereb Blood Flow Metab 23(3):285-300.

[6] Mawlawi O, Martinez D, Slifstein M, Broft A, Chatterjee R, Hwang DR, Huang Y, Simpson N, Ngo K, Van Heertum R, Laruelle M. 2001. Imaging human mesolimbic dopamine transmission with positron emission tomography: I. Accuracy and precision of D(2) receptor parameter measurements in ventral striatum. J Cereb Blood Flow Metab 21(9):1034-1057.

[7] Morris ED, Endres CJ, Schmidt KC, Christian BT, Muzik Jr. RF, Fisher RE. 2004. Kinetic modeling in positron emission tomography. In: Wernick MN, Aarsvold JN, editors. Emission tomography: the fundamentals of PET and SPECT: Elsevier. p. 499-540.

[8] Perneger TV. 1998. What's wrong with Bonferroni adjustments. BMJ 316(7139): 1236-1238.

[9] Phelps ME, editor. 2006. PET: Physics, instrumentation, and scanners. New York: Springer.

[10] Schweinhardt P, Fransson P, Olson L, Spenger C, Andersson JL. 2003. A template for spatial normalisation of MR images of the rat brain. J Neurosci Methods 129(2): 105-113.

[11] Sullivan EV, Adalsteinsson E, Sood R, Mayer D, Bell R, McBride W, Li TK, Pfefferbaum A. 2006. Longitudinal brain magnetic resonance imaging study of the alcohol-preferring rat. Part I: adult brain growth. Alcohol Clin Exp Res 30(7):1234-1247.