

Smoothing Dynamic Positron Emission Tomography Time Courses Using Functional Principal Components

TECHNICAL NOTE

Ci-Ren Jiang* & John A D Aston^{†‡} & Jane-Ling Wang[§]

Summary: A functional smoothing approach to the analysis of PET time course data is presented. By borrowing information across space and accounting for this pooling through the use of a non-parametric covariate adjustment, it is possible to smooth the PET time course data thus reducing the noise. A new model for functional data analysis, the Multiplicative Nonparametric Random Effects Model, is introduced to more accurately account for the variation in the data. A locally adaptive bandwidth choice helps to determine the correct amount of smoothing at each time point. This preprocessing step to smooth the data then allows subsequent analysis by methods such as Spectral Analysis to be substantially improved in terms of their mean squared error.

Keywords: nonparametric smoothing; principal component analysis; functional data analysis; neuroreceptor ligands

1 Introduction

It is increasingly important in many types of neuroimaging studies to try to account for both the signal and noise efficiently while reducing the number of assumptions made about the nature of the underlying data. In addition, study designs are becoming increasingly complex and the

*crjiang@wald.ucdavis.edu; Department of Statistics, University of California, Davis, CA 95616, U.S.A.

†j.a.d.aston@warwick.ac.uk; CRiSM, Department of Statistics, University of Warwick, CV4 7AL, UK

‡jaston@stat.sinica.edu.tw; Institute of Statistical Science, Academia Sinica, Taipei, Taiwan

§wang@wald.ucdavis.edu; Department of Statistics, University of California, Davis, CA 95616, U.S.A.

signals under investigation closer to the limits of detection. In this note, we consider one specific example where techniques to account for noise can be difficult due to the nonlinear nature of the data, dynamic positron emission tomography (PET) time course analysis, and propose a non-parametric approach to noise reduction.

Dynamic PET data is collected in four dimensions, three spatial and one temporal, and it will be the time dimension that is of particular interest here, although the analysis will incorporate a spatial pooling of information. The time series data collected, while typically having independent errors across time, does have a nonlinear nature due to the chemical reactions that drive the system. Due to experimental constraints such as the need to limit the radiation dosage given to a subject or the count rate properties of the PET tomograph, the time course measurements are often also fairly noisy. In many fields of application, one method to deal with noisy data would be to smooth it.

However, this causes difficulty here due to the nonlinear nature of the time course which is measured on an unequal grid and contains distinct peaks. Some methods of incorporating spatial information have been investigated in the literature (O'Sullivan, 2006; Jiang and Ogden, 2008), but these tend to require techniques such as mixture modelling, which limit their potential to be combined with commonly used techniques in PET time course analysis. Methods such as wavelets in the time dimension (Millet et al., 2000) offer one possible solution using prespecified basis functions, and can be used with additional spatial wavelet bases to incorporate both spatial and temporal information (Alpert et al., 2006). In this paper, an alternative approach is taken, where the basis functions which best describe the data are estimated directly, incorporating both spatial and temporal information non-parametrically so that subsequent model assessment, whether it be parametric or otherwise, can be carried out as before. In practice, no smoothing or spatial pooling is usually performed and a particular model or model class is assumed (Gunn et al., 2001) and estimated, and the noise just seen as a measurement error to the model. However, fits to models can be adversely affected by noise

in the data (Slifstein and Laruelle, 2000), and a way of preprocessing the data to account for some of the noise could provide significant advantages in terms of reducing the Mean Squared Error (MSE) of the parameter estimation. However, any smoothing would have to be able to account for the nonlinear nature of the data.

In traditional analysis, either voxel or region data is analyzed separately, one voxel or region at a time. If region data is used, then an average is taken across voxels presumed to be in the region, making the implicit assumption that the only source of difference in the region is measurement error. However, in the proposed method, the data are treated as random functions coming from a population. Thus the data is pooled in a manner to both account for the mean of the population and also the way in which the population varies (through analysis of the covariance structure). This makes particular sense for PET data, where the data is very much related, but where each individual curve varies dependent on a large number of biological factors (such as receptor density, blood flow, specific and non-specific binding) which could be treated as random across the set of curves.

Here, an algorithm, based on functional principal component analysis (FPCA), is proposed to smooth PET time courses without making particular model assumptions in the preprocessing. The technique is related to Principal Component Analysis which has previously been used in PET blood flow studies (Friston et al., 1993) and in PET time course analysis (Kimura et al., 2002; Turkheimer et al., 2003; Joshi et al., 2008), but differs in that it includes a smoothing component due to the functional nature (ie the data can be consider as a continuous mathematical function) of the time course data. Functional data analysis (Ramsay and Silverman, 2005) has received only limited treatment in the neuroimaging literature with studies on the special cases of single subject pooled detection (Viviani et al., 2005) and on non-stationary noise estimation (Long et al., 2005). Here it is of interest to look at both the estimation procedure for the component in the presence of noise, but also how to take into account the large numbers of curves from different parts of the image. In addition, one particular advantage of the analysis is that

some of the noise and covariance characteristics well known to be present in PET can be easily incorporated using a new statistical model we have termed the multiplicative nonparametric random effects model. The algorithm also includes steps to choose the appropriate amount of smoothing at each temporal point and the number of representative functions needed to model the data. Another intrinsic advantage is that the algorithm takes advantage of the information available on multiple neighboring time courses when adjusting an individual time course. This also has the effect of measurement error being correctly accounted for without the need to choose a weight function as in most PET analysis, as the choice of weight function can influence the bias present in the estimation after analysis (Peng et al., 2008). It will be demonstrated through simulation and measured data that the method greatly reduces the MSE estimates of parameters subsequently estimated from models or model classes and stabilizes subsequent model fitting.

While PET data will be the focus of this note, it is indeed possible to extend this analysis to other types of data, particularly fMRI. However, the presumed smooth structure of the underlying system in PET makes the method particularly appealing in this case, and thus is a compelling motivation for the methodology.

The rest of the paper is as follows. A detailed explanation of the method is given in Section 2, including the new multiplicative nonparametric random effects model. The next section contains a set of simulations to evaluate the method, and a small application to measured PET data. Finally some discussion relating to the work is given.

2 Methods

We take the point of view that the observed PET time-course data at the i^{th} voxel are generated from an underlying smooth random function $C(t)$ that is measured with additional noise $e(t)$ on a time grid $\{t_{i1}, \dots, t_{ip_i}\}$, where $\{t_{ij}; j = 1, \dots, p_i\}$ belong to some interval \mathcal{I} and could be

different for different voxels but often measurements are taken according to the same schedule. This also applies to the data in section 3, where measurements are collected over 32 temporal positions, yielding $p = 32$. These smooth functions $C(\cdot)$, which in PET experiments represent the concentration levels of radio-tracer present in that location, are examples of “functional data” (Ramsay and Silverman, 2005). Basically, functional data are random functions (or stochastic processes) that are of interest to investigators but may not be observed directly, either because of noise in the data or discrete measurement schedules. This should not be confused with the notion of physiological function, but rather a smooth function in a mathematical sense.

Consequently, for simplicity if the time of measurements is assumed the same for each voxel, the observed data at the i^{th} voxel is:

$$Y_{ij} = C_i(t_j) + e_{ij}, \quad \text{where } t_j \in \mathcal{I}, \text{ for } j = 1, \dots, p;$$

$$\text{and } e_{ij} \text{ are independent mean zero random variables with } \text{var}(e_{ij}) = \sigma^2(t_j). \quad (1)$$

Here we assume that the noise e_{ij} is independent but possibly inhomogeneous across time, which is a reasonable assumption in PET (Aston et al., 2000). We also assume the measurement error is independent across i , which is equivalent to the assumptions used in the weighting schemes in Cunningham and Jones (1993) and Gunn et al. (2002) for example. In our method, this does not imply that there is no spatial correlation present, but rather that any spatial correlation present is also present in the $C(t)$ random function without measurement error. In addition, it is not necessary for Y_{ij} to be an instantaneous estimate of $C_i(t_j)$, but can rather be an averaged estimate over a range such as described in Gunn et al. (2002, Eq. 12). This would have the effect of presmoothing the data over the interval for the estimate. However, as the distinction does not change any subsequent analysis, it will not be explicitly considered.

Our main interest is to recover the latent PET concentration trajectory $C(\cdot) = \{C(t) : t \in \mathcal{I}\}$, from the noise-contaminated measurements taken at finitely many locations. An

ad hoc approach would be to smooth each observed individual voxel time-course data with a nonparametric smoothing method (Fan and Gijbels, 1996), but such an approach is not efficient as the number p of observations is typically moderate. Besides, the level of smoothing needs to be tuned for each individual trajectory, requiring different tuning parameters at different voxels. A much more effective way to restore the entire process $C(\cdot)$ is to perform the smoothing procedure on a uniform platform, which also allows information sharing between neighbouring voxels. We demonstrate in this paper that FPCA is well suited for such a purpose.

2.1 Multiplicative Nonparametric Random Effects Model

For multivariate data, principal component analysis amounts to conducting a spectral decomposition of the covariance matrix of the multivariate random vector. Such a spectral decomposition has been extended to functional data measured without error at a densely spaced temporal grid by, e.g. Castro et al. (1986), Dauxois et al. (1982), and Rice and Silverman (1991), among others. We adopt a new approach that accommodates unequally spaced temporal grid as well as noise, which is common for PET data. Specifically, let $\mu(t) = E(C(t))$ and $\Gamma(s, t) = \text{Cov}(C(s), C(t))$ denote the mean and covariance function of the process $C(\cdot)$. The functional principal component (Karhunen-Loève) representation for the i^{th} concentration temporal curve,

$$C_i(t) = \mu(t) + \sum_{k=1}^{\infty} A_{ik} \phi_k(t), \quad t \in I, \quad (2)$$

is based on the orthogonal eigenfunctions $\phi_k(t)$ of the covariance function $\Gamma(s, t)$. The A_{ik} are uncorrelated random effects variables (known as the principal component scores) with mean 0 and variances λ_k , which are the descending eigenvalues of $\Gamma(s, t)$ in descending order.

Many of the processes which are represented in the PET time course data have chemical rates associated with them. These rates are dependent on a large number of biological factors, too numerous and complex to be exhaustively represented or identified in the discretely and

noisily measured data. However, if an alternative viewpoint that the rates are random variables is taken, then a small additive random change in one rate will lead to a multiplicative change in the time course. This is a separate multiplicative random effect from the effect of the poisson decay process which is assumed to only affect the measurement error process and not the covariance function when the measurement error has been accounted for. In addition, due to the complexity of the underlying factors which influence the random nature of the curves, it is highly desirable not to make parametric assumptions about nature of the randomness, other than the likely presence of a multiplicative effect.

Moreover, from the random curve viewpoint, a source of variation may be in the scale difference of the curves suggesting that one of the eigenfunctions, e.g. $\phi_m(t)$, may be proportional to the mean function. Often, $\phi_m(t) = \phi_1(t) = \gamma\mu(t)$, for some constant γ , meaning that the largest variation in the temporal data arises from a scale deviation from the mean function. Since this is indeed the case for the PET data in section 3, without loss of generality we assume hereafter that $\phi_m(t) = \phi_1(t) = \gamma\mu(t)$ in (2), which leads to the following multiplicative random effects model:

$$C_i(t) = B_i\mu(t) + \sum_{k=2}^{\infty} A_{ik}\phi_k(t), \quad (3)$$

where $B_i = 1 + A_{i1}\gamma$ represents a multiplicative random effect with $E(B_i) = 1$. Since the first principal direction also often explains a large fraction (over 70%) of the variation in the data, we term model (3) a multiplicative random effects model, even though there are additional random effects reflected in $A_{ik}, k \geq 2$.

A form of multiplicative standard PCA model was used by Joshi et al. (2008) for PET time course data, as the mean was included in the regression used to reduce noise. However, it was not explicitly removed from the remaining covariance, and thus does not have a true multiplicative structure as defined above.

2.2 Algorithm

Since time course data are observed at each voxel, a question arises as to how to utilize the voxel positions when performing the multiplicative random effects model (3). We propose two ad hoc procedures in this subsection, and two additional extensions in the next section. All four procedures involve a functional PCA, which forms the bedrock of all nonparametric random effects models considered in this paper. Below we describe the estimation procedures involved in each of the building blocks. In addition, in Figure 1, a flowchart of the analysis procedure used in the measured data analysis is given.

Step 1 - *Estimation of the Mean Function:* Firstly, only in-brain voxels are used for the analysis, with voxels outside the brain not considered. The mean function can be estimated by pooling all the data points and their measurement times $\{(Y_{ij}; t_j) : i = 1, \dots, n; j = 1, \dots, p\}$ and then applying a one-dimensional smoother to smooth the responses Y_{ij} across the time points t_j . Such a smoothing procedure is called a “scatter plot smoother”, and any smoothing procedure could be used. We adopt local polynomial smoothing in this paper due to its conceptual simplicity, theoretical tractability and efficiency (Fan, 1993). In principle, a local polynomial smoother at a point t_0 finds the best polynomial fits for data in a selected neighborhood of t_0 using a weighted least squares estimate, with data further away from the point t_0 receiving less weight. Computer code (Lowess) for such a smoother is available in most standard statistical software, including the R-project and MATLAB. The resulting estimated mean function is denoted by $\hat{\mu}(t)$.

Step 2 - *Estimation of the Multiplicative Random Effect B_i :* For computational efficiency and model flexibility (no parametric assumptions are made on the distribution of B_i), we propose a least squares method to estimate the multiplicative random effect B_i for subject

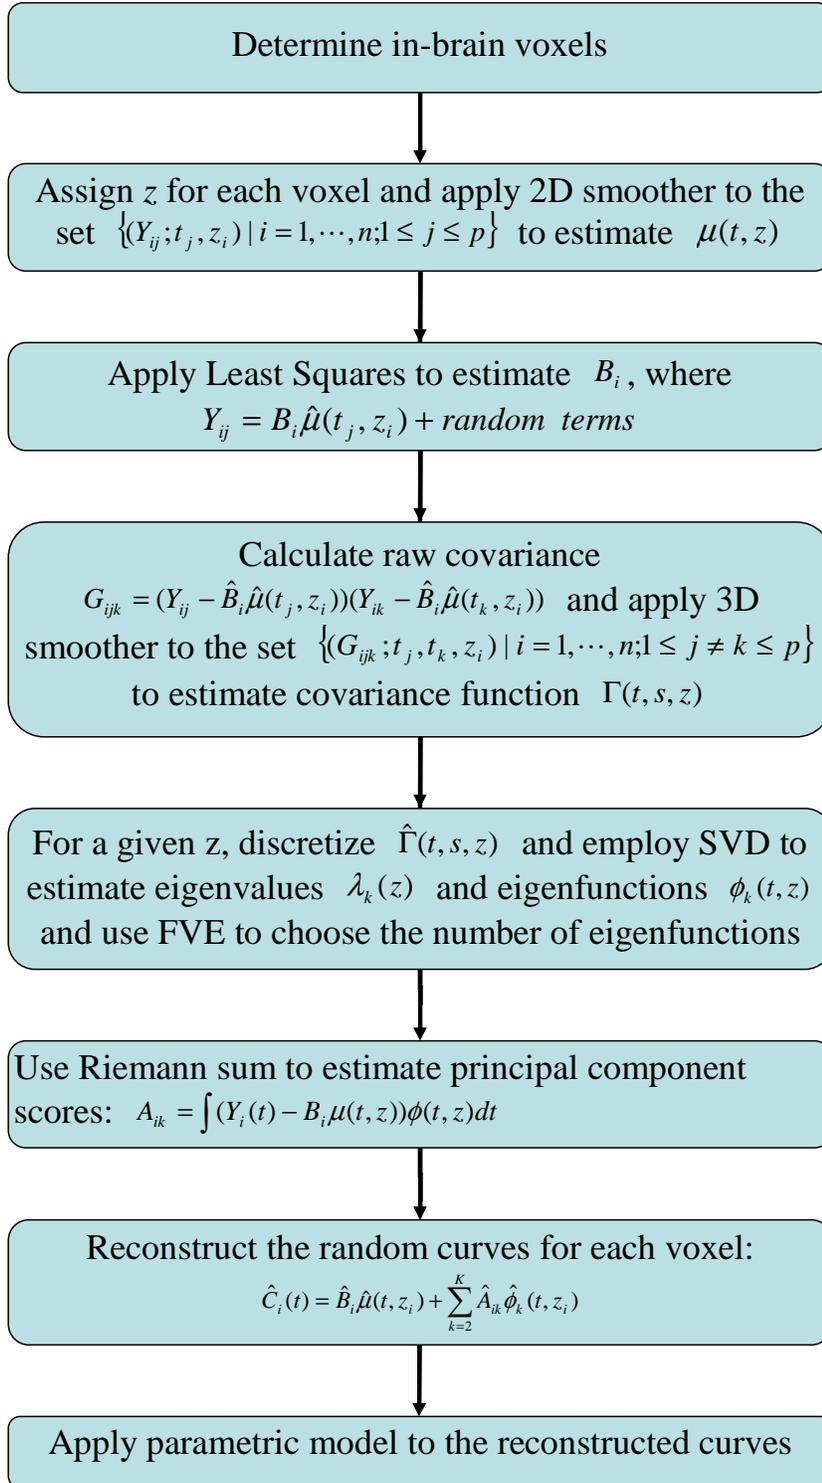


Figure 1: Flowchart of the proposed fFPCA analysis procedure from start to finish. This describes the procedure used in the ^{11}C -Diprenorphine analysis

i. Specifically, the estimate \hat{B}_i of B_i minimizes $\sum_{j=1}^p \{Y_{ij} - b\hat{\mu}(t_j)\}^2$ among all possible choices of b . Iterative weighted least squares could also be used, but empirically was found to offer no advantage and is computationally more intensive.

Step 3 - Estimation of the Covariance Function: After the mean function $\mu(t)$ and the multiplicative random effect B_i have been estimated for all subjects, the remaining covariance structure in $C(\cdot) - B\mu(\cdot)$ is $\Gamma_{(-B)} = \Gamma - \text{var}(B)\mu \otimes \mu$, where \otimes is the tensor product of two functions, and the notation $\Gamma_{(-B)}$ highlights that the variation due to the multiplicative random effects B has been removed from $C(\cdot)$. This covariance surface can be estimated from the residuals $R_{ij} = Y_{ij} - \hat{B}_i\hat{\mu}(t_j)$ by employing a two-dimensional smoother on the “raw covariances” $G_{ijl} = R_{ij}R_{il}$, excluding the diagonal elements because of the noise involved in model (1). To see this, combining (1) and (3), the covariance surface of $Y_i(\cdot) - B_i\mu(\cdot)$ is $\Gamma_{(-B)}(s, t) + \sigma^2(t)\delta_{st}$, where s and t are in the temporal interval I , and δ_{st} is 1 if $s = t$ and 0 otherwise. Therefore, the proper scatter plot to estimate the covariance surface $\Gamma_{(-B)}$ consists of $\{(G_{ijl}; t_j, t_l) : i = 1, \dots, n; 1 \leq j \neq l \leq p\}$. Applying a two-dimensional local polynomial smoother to the responses G_{ijl} over the two-dimensional time points (t_j, t_l) yields the desired estimate $\hat{\Gamma}_{(-B)}$. Details for a generic covariance surface estimate can be found in Yao et al. (2005).

Step 4 - Estimation of the Noise Levels: The variance $\sigma^2(t)$ of the noise function is also of interest and can be estimated by first estimating the variance of $R(t) = Y(t) - B\mu(t)$. To do so, apply a smoother to the scatter plot based on the diagonal elements $\{(G_{ikk}; t_k) : i = 1, \dots, n; 1 \leq k \leq p\}$, and call the resulting estimate $\hat{R}(t)$. This involves only one-dimensional smoothing, much like the smoother to estimate the mean function. The resulting estimate for $\sigma^2(t)$ is thus $\hat{\sigma}^2(t) = \hat{R}(t) - \hat{\Gamma}_{(-B)}(t, t)$.

Step 5 - Estimation of the Eigenvalues and Eigenfunctions: Once we have a covariance estimate for $\Gamma_{(-B)}$, the eigenfunctions and eigenvalues $\{(\phi_k, \lambda_k) : k \geq 2\}$ can be estimated

by solving, on a discrete temporal grid, the eigen-equations,

$$\int_{\mathcal{I}} \widehat{\Gamma}_{(-B)}(s, t) \hat{\phi}_k(s) ds = \hat{\lambda}_k \hat{\phi}_k(t), \quad (4)$$

where the $\hat{\phi}_k$ are subject to $\int_{\mathcal{I}} \hat{\phi}_k(t)^2 dt = 1$ and $\int_{\mathcal{I}} \hat{\phi}_k(t) \hat{\phi}_m(t) dt = 0$ for $m < k$. The solution of (4) is denoted as $\hat{\phi}_k$ and $\hat{\lambda}_k$. Details of the discretized solutions for the eigen-equation can be found in Rice and Silverman (1991) and Yao et al. (2005).

Two Ad Hoc Multiplicative FPCA Approaches: Since PET time-course data are available on each voxel, which belongs to a particular slice, there are several ways to perform FPCA. Here slice can be taken to mean any particular segmentation of the data in space. The natural one to use in the measured PET data will be the z-plane. Below we describe two ad hoc approaches.

- **pFPCA:** This approach combines all the PET time-course data across all voxels into a pooled sample of curves and then apply steps 1-5 of the algorithm. The prefix "p" stands for "pooling".
- **smFPCA:** This approach estimates the mean function for each slice separately by implementing the one-dimensional smoother in Step 1 using all the voxel curves on the same slice. Once the mean function on each slice has been estimated, Step 2 of the algorithm is employed for each voxel curve using the estimated mean function for the same slice in which this voxel curve lies. Afterwards, all the residual time-courses R_{ij} are pooled together from all voxels, and a common residual covariance is estimated through a two-dimensional smoothing approach as in Step 3 of the algorithm. Thereafter, Steps 4 and 5 are employed. The prefix "sm" of this approach suggests that the difference is in the sliced mean function, but not the covariance function, has been accounted for.

To summarize, these ad hoc approaches utilize information of the slice location of each voxel differently, with pFPCA ignoring this information completely. A similar FPCA approach that does not assume the multiplicative structure of our pFPCA was proposed in Viviani et al. (2005) for fMRI data. Our numerical findings in section 3 suggest efficiency loss when pooling all the voxel curves together to borrow information. An interesting question arises as how to best utilize the slice information. Below we describe two alternative approaches.

2.3 Multiplicative Nonparametric Random Effects Model with Covariates

The above question can be framed in a more general way:

How does one take into account available covariate information in PET time-course data under the platform of functional principal component analysis?

As an example, in the PET data in section 3, we have information on the thickness of each slice and its vertical location in the brain which can be used as information to assist the smoothing. The information is more informative than simply the voxel size as location gives relative position of voxels. Let Z reflect this covariate information; in this case Z is a scalar. Since we do not know for sure the form of the covariate effects, a nonparametric approach seems suitable as an exploratory tool. The only assumption that we will make is that the mean functions, and possibly the covariance surfaces as well, change smoothly across the slices, a very mild assumption given the spatial resolution of the data. Therefore sharing some information across slices likely will increase efficiency. There are two ways to model covariate effects, either through the mean function only, or through both mean and covariance functions. Both approaches assume that the mean function μ depends on a covariate value z , so that $\mu \equiv \mu(t, z)$. The difference is in the handling of the covariance structure. While the fully

adjusted approach has the covariance surfaces also depending on the values of the covariate, the mean adjusted approach does not account for the covariate effect on the covariances. Similar approaches to handle covariate effects are investigated in a technical report of Jiang and Wang (2008), but for a different model. Below we briefly describe how to adapt these approaches to the setting (3) of a multiplicative random effects model.

Mean Adjusted FPCA (mFPCA):

Here one assumes for the i^{th} voxel whose slice information is coded in z_i , that :

$$C_i(t) = B_i(z_i)\mu(t, z_i) + \sum_{k=2}^{\infty} A_{ik}(z_i)\phi_k(t), \quad (5)$$

Thus, only the mean function varies with z_i but not the residual covariance. This model is a counterpart of the smFPCA in section 2.2, where centering of the voxel time-course curves from the same slice z , as labeled by their common covariate value z , was done by estimating the mean curve $\mu(t, z)$, using the sample mean of all voxel curves on this slice. The difference here is that mFPCA estimates the mean function $\mu(t, z)$ differently by employing a two-dimensional smoother to smooth over both the t and z coordinates. Thus, a two-dimensional scatter plot smoother is employed to $\{(Y_{ij}; t_j, z_i) : \text{for } 1 \leq j \leq p, \text{ and } 1 \leq i \leq n\}$. Once we have an estimate for the mean function, the remaining steps 2-5 can be performed as in section 2.2 by replacing $\mu(t_j)$ and its estimate with $\mu(t_j, z_i)$ and the corresponding estimates.

Mean and Covariance Fully Adjusted FPCA (fFPCA):

This model assumes that both the mean function $\mu(t, z)$ and residual covariance function vary with the covariate information z . Thus, for the i^{th} voxel curve with covariate z_i ,

$$C_i(t) = B_i(z_i)\mu(t, z_i) + \sum_{k=2}^{\infty} A_{ik}(z_i)\phi_k(t, z_i), \quad (6)$$

where $\phi_k(t, z_i)$ are the eigenfunctions of the residual covariance function of this subject. Here

the mean function can be estimated exactly the same way as in mFPCA above, and so is step 2 of the algorithm, but a three dimensional smoother is needed in step 3 to estimate the residual covariance function by using the residuals $R_{ij} = Y_{ij} - \hat{B}_i(z_i)\hat{\mu}(t_j, z_i)$, whose values now depend on z_i . Following the same procedure as in step 3 of the algorithm in section 2.2, the “raw covariance” $G_{ijl} = R_{ij}R_{il}$ also depends on z_i . Thus, to estimate the covariance surface $\Gamma_{(-B)}$, we need to employ a three-dimensional smoother to the scatter plot $\{(G_{ijl}; t_j, t_l, z_i) : i = 1, \dots, n; 1 \leq j \neq l \leq p\}$. The algorithm in steps 4 and 5 can be adjusted accordingly. While a three-dimensional smoothing step involves higher computational cost and requires a much larger sample size, the rich data information contained in different voxels enables the practical implementation of the proposed fFPCA approach for PET time course data. We demonstrate its feasibility through simulations in section 3.

2.4 Bandwidth Selection

The aforementioned algorithms to perform the multiplicative FPCA require nonparametric smoothing methods involve the choice of a smoothing parameter to regulate the amount of desired smoothness. For the local polynomial smoothers used in the smoothing steps, the corresponding smoothing parameter is the bandwidth, which assigns less weights to data further away from the time point where the estimate is located. Usually, a single global bandwidth is used at all locations for simplicity and cross-validation methods are typically used to select the desired bandwidths. For PET time course data, a global bandwidth is appropriate along the covariate coordinate (z) but not desirable in the time coordinate due to the denser measurement schedule at the beginning of the time period and the sharp peak for each voxel curve near the left boundary (see Figure 2(a)). To retain these peaks and in order not to compromise the performance at other temporal locations, local bandwidths, which adapt to the temporal location, are recommended and used in our analysis.

Adaptive Local Bandwidths for the time dimension: In essence, a smaller bandwidth which involves less smoothing is preferred near the peak location, while larger bandwidths are used near the right boundary where the curve is relatively flat. The resulting local bandwidths used in our numerical analysis in section 3 are plotted in Figure 2(b). Specifically, we first selected 13 time locations where measurements of the time-course data were sampled and then determined the respective bandwidths for each of these 13 locations. To choose bandwidth $b(t)$ at location t we used a rule of thumb that the interval $[t - b(t), t + b(t)]$ included at least four observations. We also employed boundary correction to ensure a positive bandwidth choice. Overall, as the PET data were sampled more frequently in the beginning than near the end of the observation period, this choice led to increasing bandwidths for increasing time points. To ensure a smooth outcome, the bandwidth needs to vary smoothly across time locations, we therefore fitted a polynomial of order 4 to the 13 pairs $(t_j, b_j(t))$. The resulting curve $\hat{b}(t)$ was further multiplied by a constant α , determined by an additional cross-validation step. This extra constant α serves to facilitate calibration of the final local bandwidths, because the aforementioned choice of local bandwidths for $b(t)$ is subjective. Thus, selecting α by cross-validation, the final local bandwidths are $\hat{b}(t)$ as shown in Figure 2(b). The strategy of applying a local bandwidth choice for t and a global bandwidth for z seems to yield satisfactory results when estimating the mean and covariance functions.

The selection of α through cross validation follows standard practice that deletes one voxel curve while fitting the mean of covariance function using the remaining voxel curves. This is possible for two-dimensional smoothers such as in pFPCA, smFPCA, or mFPCA, but can be computing intensive, especially when a three-dimensional smoother is employed in fFPCA. We therefore employed a leave-one-slice-out cross-validation method for fFPCA by deleting one slice at a time to save computing time. For the remaining FPCA approaches, we also recommend to adhere to the leave-one-slice-out cross validation as numerical results in section 3 suggest that it performs well in general and the computational cost is manageable on a PC.

We further speed up computationally the bandwidth selection process by using a small approximation to the full multidimensional smoothing which was undertaken in the data analysis. Specifically, separable smoothing in the cross-validation step is used, in that the time dimensions are smoothed sequentially and then the covariate dimension to select the bandwidth. However, when the smoothing parameters have been selected, a full 2-D or 3-D smoothing incorporating all information is used to smooth the functions of interest (ie the final smoothing is not sequentially performed but performed in a full multidimensional manner at each point).

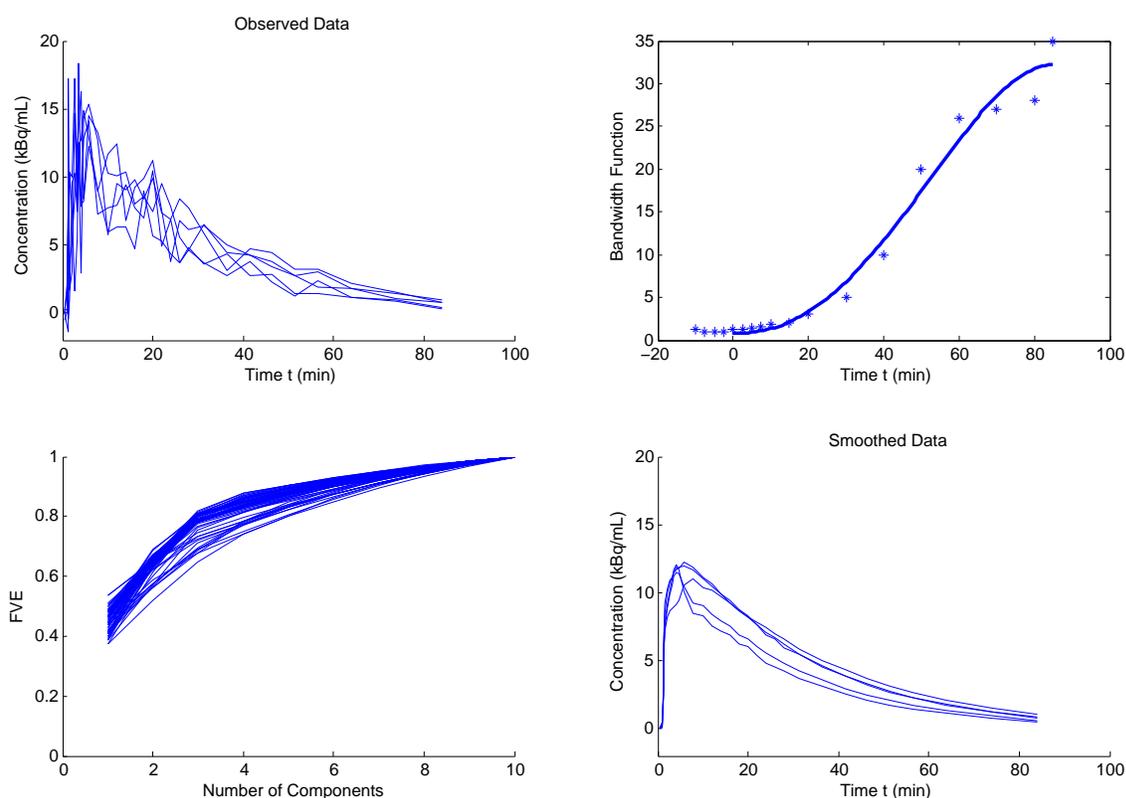


Figure 2: Example of analysis for five voxels from ^{11}C -Diprenorphine data. Top Left (a) Original Voxel Data; Top Right (b) Local Bandwidth Function; Bottom Left (c) Scree Plot; Bottom Right (d) Smoothed Voxel Data

2.5 Number of Components Selection

Several approaches could be applied to choose the number of components, such as AIC, BIC, and fraction of variance explained (FVE). However, AIC and BIC require model assumptions for all the random terms while FVE does not. Thus to avoid making further model assumptions, we employ FVE to select the number of components for the predicted trajectories. The FVE method is defined as the minimum number of components needed to explain at least a specified fraction, say 80%, of the total remaining variation not explained by the multiplicative random effects. That is, the number of components k selected by FVE can be represented as $\sum_{i=2}^k \hat{\lambda}_i / \sum_{i=2} \hat{\lambda}_i \geq 0.80$. In practice, the FVE method could be guided by a scree plot (see Figure 2(c)) and 80% provides satisfactory results in our simulation setting and for the measured [^{11}C]-Diprenorphine data in section 3. Examining the PET simulation data, two eigenfunctions explain about 80% of the variation when the noise level is close to that of measured data and thus was chosen most often as the number of components needed. In the measured data, three components seemed to most often represent at least 80% of the variation. Furthermore, under the multiplicative model assumptions as described in section 2.1, the eigenfunctions actually explain about 80% of the residual variation not explained by the multiplicative term, in essence making this value a strict lower bound of the amount of total variation explained, with it likely to be much large in practice. Therefore, even though only a very small number of eigenfunctions are used, the reconstruction of the concentration trajectories after smoothing would appear to be very accurate (see Figure 2(d)).

2.6 Reconstruction of the PET Time Course Data

Once all the unknown quantities in the multiplicative model (3) are estimated and the number of components K selected, the concentration temporal curve for the i^{th} voxel can be reconstructed as

$$\hat{C}_i(t) = \hat{B}_i \hat{\mu}(t) + \sum_{k=2}^K \hat{A}_{ik} \hat{\phi}_k(t). \quad (7)$$

2.7 Subsequent Parameter Estimation - Spectral Analysis

Having found a suitable decomposition of the time course data using the multiplicative models above, the data are reconstructed from its functional components and analyzed using standard techniques such as Spectral Analysis (Cunningham and Jones, 1993). While such methods often are affected by high noise levels, as the preprocessing reduces the noise by only selecting a few components to represent the data, the fitting becomes more robust. In addition, a particular feature of the preprocessing is that the measurement error does not need to be subsequently specified, as it is nonparametrically assessed in the procedure and then accounted for. This means that in the subsequent analysis no weight function needs to be specified. Weighting schemes in PET are usually based on a single global scheme for all voxels, but it has been observed that incorrect weights can lead to significant biases in analysis due to the non-linearities present (Peng et al., 2008). Therefore removing this step of requiring pre-specification of the weights is an interesting byproduct of the algorithm.

In the analysis below, the data was all analysed using Spectral Analysis (Cunningham and Jones, 1993) after smoothing. Spectral Analysis does not assume a known compartmental structure, but rather performs a model selection through a non-negativity constraint on the parameters. In particular, the curve $C(t)$ is parameterised by

$$C(t) = I(t) \otimes \sum_{j=1}^K \alpha_j e^{-\beta_j t} \quad (8)$$

where $I(t)$ is a known input function, and α_j and β_j are the parameters to be estimated subject to a non-negativity constraint. In practice K is set to a very large number and most α_j are

estimated to be zero. The parameter of interest V_T is then the integral of the impulse response function

$$V_T = \int_0^\infty \sum_{j=1}^K \alpha_j e^{-\beta_j t} dt = \sum_{j=1}^K \frac{\alpha_j}{\beta_j} \quad (9)$$

For more details see Cunningham and Jones (1993) or Gunn et al. (2002). This method is well known to be sensitive to noise with the bias being highly dependent on the level of noise present (Gunn et al., 2002; Peng et al., 2008). Thus if smoothing can reduce the noise level in the data, the analysis should be much less biased even at noise levels associated with voxel data.

3 Data Analysis

In order to determine the efficacy of the proposed methods in relation to standard analyses, simulations were performed along with the analysis of a measured Diprenorphine data set. The simulations were designed to determine whether combining information, through the different types of pooling proposed above, in different settings would be beneficial to the analysis, while the measured data set allowed an evaluation of the methods in practice. The results of the above models were compared to the results from Spectral Analysis performed on the original data, which neither has any intrinsic smoothing associated with it, nor makes use of any spatial pooling of the data, as the analysis is performed independently on each voxel.

3.1 Blurred Single Region Analysis

A set of simulations were generated for different measurement error noise levels based on the simulations in Gunn et al. (1997) for a blood/plasma input function model. Simulations were generated from a two-tissue compartment model with associated input function such that the simulated data was directly comparable to the measured data below. The rate constants used

in the simulations were taken from Gunn et al. (1997) along with a measured plasma input function (the input function associated with the measured data). Thus the rate constants [$K_1 = 6.7 \times 10^{-3}$ (mL plasma) \cdot sec $^{-1}$ \cdot (mL tissue) $^{-1}$, $k_2 = 3.3 \times 10^{-3}$ sec $^{-1}$, $k_3 = 6.7 \times 10^{-3}$ sec $^{-1}$ and $k_4 = 1.7 \times 10^{-3}$ sec $^{-1}$] for the target tissue were used, yielding a volume of distribution (V_T) value of 10. Noise added to the data was proportional to the true data divided by the frame duration, with the proportionality varying from 0.01 to 1.28, giving noise levels from lower than typical ROI analysis to slightly higher than voxel level analysis (voxel level is equivalent to approximately the value 1 in this scheme). We generated 50 images of 128×128 for each noise level. Here, for practicality of generating multiple simulations at different noise levels, we consider two dimensional simulations rather than three dimensional ones. In this case the covariate adjusted FPCA will be adjusted across columns rather than planes for each of the 50 realisations. For additional information concerning the definition of the model, of V_T and of the simulation setup for the individual time course curves, see Gunn et al. (1997).

An additional level of randomness was added to the simulations, by allowing the rate constants to vary in such a way that the associated true V_T in the simulations had a variability of approximately 6%, roughly equivalent to the voxelwise variability observed within regions in the measured data. Finally, the data was blurred using a standard Gaussian blurring kernel, with FWHM of 6mm, with the voxels in the image being presumed to be $2\text{mm} \times 2\text{mm}$ (as this is a 2 dimensional simulation).

The results of the single blurred region analysis can be seen in Table 1. Firstly, it does seem counterintuitive that the measurement noise level seems to have little effect on results. However, this can be explained by considering the fact the noise simulations are blurred across the planes, inducing a dependence structure. Due to the curves each being reconstructed from a set of random parameter values, this causes there to be multiple components present in the data even without the measurement error noise being added. These components can be difficult to detect, causing a large variance in the resulting estimates. The MSE is dominated by this

Noise	No Smoothing	pFPCA	mFPCA	fFPCA	smFPCA
0.01	12.869	5.7804	5.7384	5.5988	0.9564
	(0.15)	(1.5829)	(1.5655)	(0.4804)	(0.0785)
0.08	12.857	5.4421	5.4172	5.5926	0.9402
	(0.121)	(1.5)	(1.4848)	(0.3906)	(0.0825)
0.32	12.8697	5.7155	5.6923	5.5089	0.9563
	(0.1703)	(1.5095)	(1.4686)	(0.364)	(0.0911)
0.64	12.8781	5.9689	5.923	5.5522	0.9568
	(0.1375)	(1.3145)	(1.3067)	(0.4502)	(0.0849)
1.28	12.8835	6.0256	5.9867	5.5779	0.9535
	(0.151)	(1.4445)	(1.4667)	(0.5036)	(0.0863)

Table 1: MSE (standard error) values of V_T for the different methods for the blurred single region simulation study. It can be seen that all the smoothing methods lead to an improvement in the MSE in this situation.

variance in all cases. However, it is also very noticeable that the smoothed results all greatly improve on the MSE, in particular the sliced mean FPCA data analysis. This suggests that in principle it would be advantageous to smooth the data prior to analysis.

3.2 Blurred Phantom Simulations

The simulations of the previous section do however fail to account for the feature in PET data that multiple regions are present in the data, and that the smoothing takes advantage of a spatial pooling of the data. In order to assess the effect of different regions, a similar simulation was now performed using a brain phantom (Shepp-Vardi phantom, 128×128 pixels, repeated over 50 realisations) with five different regions of varying sizes. Different signals were placed in each of the regions based again on random parameters and then the phantom blurred as above. The values used in each region, along with its size is give in Table 2. The idea was to see how well different regions were recovered when they would be subjected to different partial volume effects based on their size and differing intrinsic V_T values. Again the covariate adjustment was across columns in the 2-D simulation.

The results of the blurred phantom simulation study are somewhat consistent with the

Region	Size	K_1 $((\text{mL}_{\text{plasma}}) \cdot \text{sec}^{-1} \cdot (\text{mL}_{\text{tissue}})^{-1})$	k_2 (sec^{-1})	k_3 (sec^{-1})	k_4 (sec^{-1})	V_T
1	14	6.70×10^{-3}	3.30×10^{-3}	6.70×10^{-3}	1.70×10^{-3}	10
2	701	6.70×10^{-3}	6.70×10^{-3}	1.70×10^{-2}	1.70×10^{-2}	2
3	5351	6.70×10^{-3}	6.70×10^{-3}	-	-	1
4	704	1.70×10^{-2}	1.70×10^{-2}	-	-	1
5	9614	-	-	-	-	0

Table 2: Parameter values and region sizes for the phantom simulation

results of the previous simulations, although one marked difference with the performance of smFPCA occurs. In this case there is no covariate information in the first simulation set, so pooling over all the curves and only adjusting the means in each column actually helped capture more information. However, in the phantom case, the performance is enhanced by assuming that there is a covariate function across the columns, as the differing regions have different variance structures (due to the random parameter values), and thus smFPCA performance is affected as it does not use this information. pFPCA is somewhat effective at low noise levels, but becomes inefficient at higher noise levels, while mFPCA is also beneficial at all noise levels except the largest. In contrast, in all cases using fFPCA, where both the mean and covariance functions are adjusted, greatly improves the MSE for the data. The highest noise level is approximately of the same magnitude as the noise level observed in voxelwise measured data (in addition to the within region variability added in the simulations) and fFPCA still yields considerable gains over the method without smoothing.

3.3 Measured ^{11}C -Diprenorphine Data

A dynamic scan from a measured ^{11}C -diprenorphine study of normal subjects, for which an arterial input function was available, were analysed. The subject underwent a 95-min dynamic ^{11}C -diprenorphine PET baseline scan. The subject was injected with 185 MBq of ^{11}C -diprenorphine. PET scans were acquired in 3D mode on a Siemens/CTI ECAT EXACT3D

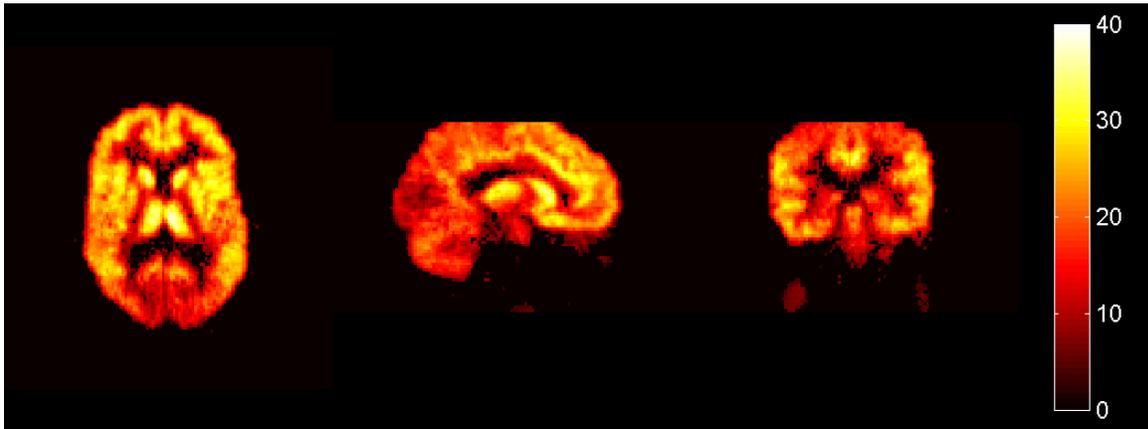
Noise	No Smoothing	pFPCA	mFPCA	fFPCA	smFPCA
0.01	1.2271	0.9733	0.9103	0.6533	1.5724
	(0.0454)	(0.3815)	(0.2236)	(0.037)	(0.9008)
0.08	1.2424	0.8319	0.6859	0.636	2.1873
	(0.0435)	(0.4554)	(0.2359)	(0.0862)	(0.5765)
0.32	1.2963	0.862	0.7064	0.5837	2.2461
	(0.0545)	(0.685)	(0.3171)	(0.1268)	(0.4022)
0.64	1.394	1.4905	1.1542	0.759	2.4844
	(0.0762)	(1.2525)	(0.6327)	(0.1607)	(0.583)
1.28	1.5834	2.1756	1.6746	1.1869	2.4227
	(0.1299)	(1.0299)	(0.8057)	(0.2136)	(0.5206)

Table 3: MSE (standard error) values of V_T for the different methods for the blurred phantom simulation study. The smoothing results are more variable in this case. However, fFPCA does outperform the analysis with no smoothing in all cases.)

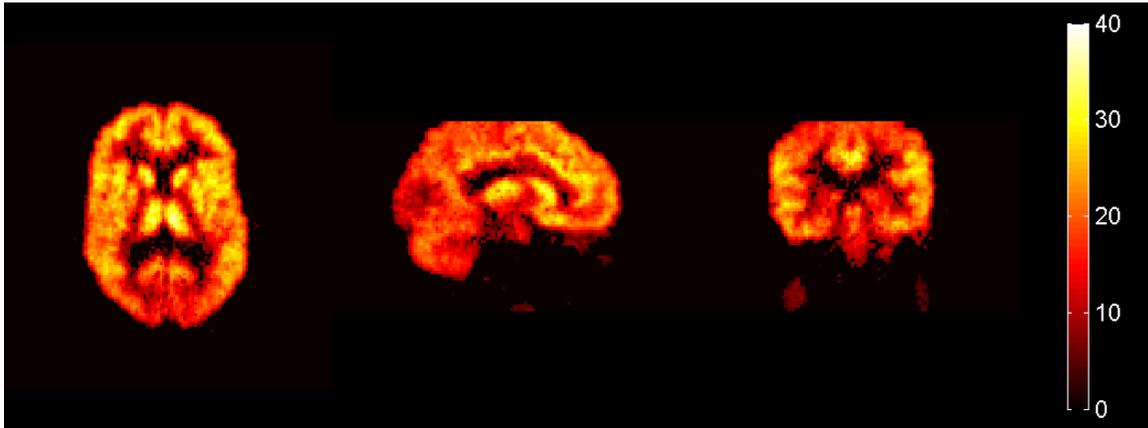
PET camera, with a spatial resolution after image reconstruction of approximately 5 mm. Data were reconstructed using the reprojection algorithm (Kinahan and Rogers, 1989) with ramp and Colsher filters cutoff at Nyquist frequency. Reconstructed voxel sizes were 2.096 mm \times 2.096 mm \times 2.43 mm. Acquisition was performed in listmode (event-by-event) and scans were rebinned into 32 time frames of increasing duration. Frame-by-frame movement correction was performed on the dynamic [^{11}C]-diprenorphine PET images.

From the results of the simulation it was deemed that fFPCA was most applicable to the analysis of data where there are multiple regions. The analysis was carried out in three dimensions and the covariate adjustment carried out using z-plane information (rather than column information as was done in the simulations). A scree plot was used to determine the number of eigenfunctions needed for the smoothing, and on average three eigenfunctions ($K = 3$) were needed to explain 80% of the variation of the remaining eigenfunctions after accounting for the multiplicative random effect component. The flowchart of the analysis procedure is given in Figure 1.

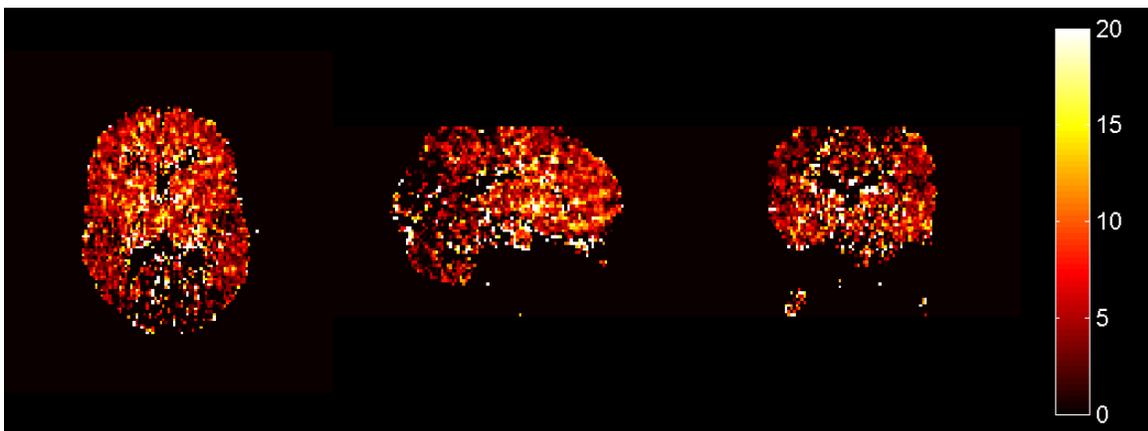
Figure 2(a) shows five typical voxels from the study prior to smoothing while Figure 2(d) shows the same voxels after smoothing. The results show that the smoothing retains the



(a) Spectral Analysis (no smoothing)



(b) Functional Smoothing Spectral Analysis using fFPCA



(c) Percentage Difference image (Spectral Analysis-Smoothed Spectral Analysis)

Figure 3: Comparison of Spectral Analysis results of V_T parameter for ^{11}C -Diprenorphine with and without smoothing using the multiplicative model. As can be seen the second has lower values than the first (about 4.5% on average for in brain voxels), which is consistent with the fact that Spectral Analysis without smoothing has high positive bias at voxel noise levels.

functional form of the data while removing almost all the noise from the signals. Indeed in terms of mean squared residuals, the noise is reduced on average by 72% (range 55%-90%). As can be seen in Figure 3, the data without smoothing has higher V_T on average than that after smoothing. The average difference was about 4.5% and this value was fairly uniform over the range of V_T values, although there is an increase in difference as V_T becomes higher. This is consistent with the knowledge that Spectral Analysis has high positive bias at voxel noise levels of around 5% (Gunn et al., 2002; Peng et al., 2008). By smoothing the data, the noise level is reduced, and thus the level of bias is also reduced. In addition, overall, the average reduction in mean squared residuals over all image voxels was 71.82%.

4 Discussions

In this paper, a methodology has been presented for smoothing PET time course data, using a new statistical model based on multiplicative functional principal components. It has been demonstrated that using this method reduces the MSE for subsequent parameter estimation as many techniques such as Spectral Analysis are known to have noise dependent bias, and thus the removal of noise prior to the analysis will improve their accuracy. In terms of analysis the model improves estimation in two distinct ways. Firstly, the error distribution is accounted for automatically yielding reduced noise after smoothing with the remaining noise having homogeneous variance across time removing the need to specify a variance weight function for the noise in the data. Secondly and most importantly, the analysis borrows information across space and when adjusted using a non-parametric covariate for the spatial information, this yields dramatic improvements in MSE. Although we have concentrated on using Spectral Analysis in subsequent analysis, due to both its popularity in clinical applications and also its well known noise dependent bias, the multiplicative model developed here is completely independent of any subsequent parameter estimation method used. In addition, it is not dependent on whether a

plasma input function is available (as was the case in the measured data here) or not, in that the output of the analysis is image data that has been temporally smoothed.

One could also perform a standard FPCA without assuming a multiplicative random effects model. The algorithm is actually simpler by skipping the estimation of the multiplicative random effects, or equivalently, setting $B_i = 1$ for all voxels. However, given the nature of the data, and the likely way the random nature of the curves is expressed through the unknown rate constants, a multiplicative model would seem more appropriate.

Several different methods have been presented to carry out the smoothing, all attempting to borrow information spatially to inform the temporal smoothing. The more naive methods of pooling such as pFPCA have been shown to be less effective than methods which try to account for the spatial structure to some degree through the use of non-parametric covariate adjustment, particularly using a full covariate adjustment on both the mean and the covariance structure ffPCA. However, the more complex methods do carry a computational penalty associated with them, in that carrying out the covariate adjustment can be time consuming. In addition, it would be possible to perform a full FPCA analysis on each slice separately, sfFPCA as a natural extension of smFPCA. However, this would be equivalent to doing a slicewise pFPCA which lacks the necessary covariate adjustment to make full use of the data. In addition, due to the need to perform a new bandwidth choice for each slice (the most computationally intensive part of the procedure due to the cross-validation), this method is much more computationally expensive than the other methods. We did however verify that no advantage would be gained by attempting this method in the simulations (data not shown) and it was found to be worse than the other smoothing methods for the blurred phantom multiple region data. In addition, we do not claim to find unbiased estimates of the signal after smoothing as smoothing naturally introduces a bias into the estimation procedure. One source of bias is the possible autocorrelation that enters the system due to the smoothing. However, given the non-parametric approach undertaken, it is difficult to assess the exact nature of any autocorrelation present with respect

to any other particular parametric model. However, we do believe that any bias is small in comparison to the bias which arises from applying a procedure such as Spectral Analysis to highly noisy data as is evidenced by the simulations.

In all non-parametric smoothing methods there is the problem of choosing the bandwidth parameter. It has been shown that a locally adaptive smoothing structure is more applicable here, given the nature of the curve, and cross-validation is used to select the final bandwidth choice. In addition, the scree plot of percentage of variance explained is also used to select the number of components. This could be modified to an AIC or BIC approach to choose the number of components if that were preferred, but the approach used here seemed robust both in the simulations and real data analysis. The fact of using the multiplicative model means that if the FVE criteria is used, the amount of total variation explained will be considerably more than the level chosen when the mean is proportional to the first eigenfunction. By examining a non-multiplicative FPCA model for the simulation data, it was found that the first functional principal component not only resembled the mean to a great extent, but that it accounted for 70% of the total variation. This would imply that an FVE criteria for the multiplicative model of 80% would actually account for closer to 95% of the total variation in the data. Alternatively, more localised measures could be used but these are more difficult to assess objectively than the global measures such as FVE and as such FVE was chosen.

While the current choice of the non-parametric covariate function could justifiably be considered slightly naive, just being the z -location in space, it would appear that this function captures enough of the spatial structure to yield the pooling of the information over this covariate very useful in reducing the MSE of subsequent fitting. However, this is just a first step and it is intended to make the covariate information more complex in future. Several possibilities are available. The use of surrogate information such as structural MRI could be combined to give some covariate function that would contain more information, or a full three dimensional spatial covariate used. This second approach would be particularly computationally intensive

in that the scatter plot smoother would now need to be over four dimensions in mFPCA and five dimensions in fFPCA. However, by using a simpler smoothing algorithm it may be possible to achieve this, and this is one of the lines of enquiry in our current research.

The methodology presented here is still very much in its infancy especially when it is applied to neuroimaging data. The aim of this note was to show that the methods of functional principal component analysis can be applied in practice to measured PET data using spatially pooled information. While there is still much scope to improve how the spatial information is incorporated in the pooling, it is interesting to see that even a simple and somewhat naive approach of pooling in one dimension can dramatically reduce the MSE error estimates. We believe that further research into the use of functional data analytic methods in neuroimaging will lead to many fruitful advances.

References

- Alpert, N. M., A. Reilhac, T. C. Chio, and I. Selesnick (2006). Optimization of dynamic measurement of receptor kinetics by wavelet denoising. *Neuroimage* 30, 444–451.
- Aston, J. A. D., R. N. Gunn, K. J. Worsley, Y. Ma, A. C. Evans, and A. Dagher (2000). A statistical method for the analysis of positron emission tomography neuroreceptor ligand data. *Neuroimage* 12, 245–256.
- Castro, P. E., W. H. Lawton, and E. A. Sylvestre (1986). Principal modes of variation for processes with continuous sample curves. *Technometrics* 28, 329–337.
- Cunningham, V. J. and T. Jones (1993). Spectral analysis of dynamic PET studies. *Journal of Cerebral Blood Flow and Metabolism* 13, 15–23.
- Dauxois, J., A. Pousse, and Y. Romain (1982). Asymptotic theory for the principal components

- analysis of a vector random function: Some applications to statistical inference. *Journal of Multivariate Analysis* 12, 136–154.
- Fan, J. and I. Gijbels (1996). *Local Polynomial Modelling and Its Applications*. London: Chapman and Hall.
- Friston, K., C. Frith, P. Liddle, and R. Frackowiak (1993). Functional connectivity: The principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism* 13, 5–14.
- Gunn, R. N., S. R. Gunn, and V. J. Cunningham (2001). Positron emission tomography compartmental models. *Journal of Cerebral Blood Flow and Metabolism* 21, 635–652.
- Gunn, R. N., S. R. Gunn, F. E. Turkheimer, J. A. D. Aston, and V. J. Cunningham (2002). Positron emission tomography compartmental models: A basis pursuit strategy for kinetic modeling. *Journal Of Cerebral Blood Flow And Metabolism* 22, 1425–1439.
- Gunn, R. N., A. A. Lammertsma, S. P. Hume, and V. J. Cunningham (1997). Parametric imaging of ligand-receptor binding in PET using a simplified reference region model. *Neuroimage* 6(4), 279–287.
- Jiang, C. R. and J. L. Wang (2008). Covariate adjusted functional principal components analysis for longitudinal data. *submitted*.
- Jiang, H. and R. T. Ogden (2008). Mixture modelling for dynamic PET data. *Statistica Sinica* 18, 1341–1356.
- Joshi, A. D., J. A. Fessler, and R. A. Koeppe (2008). Improving PET receptor binding estimates from logan plots using principal component analysis. *Journal of Cerebral Blood Flow and Metabolism* 28, 852–865.

- Kimura, Y., M. Senda, and N. M. Alpert (2002). Fast formation of statistically reliable FDG parametric images based on clustering and principal components. *Physics in Medicine and Biology* 47, 455–468.
- Kinahan, P. and J. Rogers (1989). Analytic 3D image reconstruction using all detected events. *IEEE Transactions on Nuclear Science* 36, 964–968.
- Long, C. J., E. N. Brown, C. Triantafyllou, I. Aharon, L. L. Wald, and V. Solo (2005). Non-stationary noise estimation in functional mri. *NeuroImage* 28, 890 – 903.
- Millet, P., V. Ibáñez, J. Delforge, S. Pappata, and J. Guimón (2000). Wavelet analysis of dynamic PET data: application to the parametric imaging of benzodiazepine receptor concentration. *Neuroimage* 11, 458–472.
- O’Sullivan, F. (2006). Locally constrained mixture representation of dynamic imaging data from pet and mr studies. *Biostatistics* 7, 318–338.
- Peng, J. Y., J. A. D. Aston, R. N. Gunn, C. Y. Liou, and J. Ashburner (2008). Dynamic positron emission tomography data-driven analysis using sparse bayesian learning. *IEEE Transactions on Medical Imaging* 27, 1356–1369.
- Ramsay, J. O. and B. W. Silverman (2005). *Functional Data Analysis* (2nd ed.). Berlin: Springer.
- Rice, J. A. and B. W. Silverman (1991). Estimating the mean and covariance structure non-parametrically when the data are curves. *Journal of the Royal Statistical Society: Series B* 53, 233–243.
- Slifstein, M. and M. Laruelle (2000). Effects of statistical noise on graphic analysis of PET neuroreceptor studies. *Journal of Nuclear Medicine* 41, 2083–2088.

- Turkheimer, F. E., R. Hinz, R. N. Gunn, J. A. D. Aston, S. R. Gunn, and V. J. Cunningham (2003). Rank-shaping regularization of exponential spectral analysis for application to functional parametric mapping. *Physics in Medicine and Biology* 48, 3819–3841.
- Viviani, R., G. Grön, and M. Spitzer (2005). Functional principal component analysis of fMRI data. *Human Brain Mapping* 24, 109–129.
- Yao, F., H. G. Müller, and J. L. Wang (2005). Functional data analysis for sparse longitudinal data. *Journal of American Statistical Association* 100, 577–590.