

Dynamic Positron Emission Tomography Data-Driven Analysis Using Sparse Bayesian Learning

Jyh-Ying Peng, John A. D. Aston*, Roger N. Gunn, Cheng-Yuan Liou, and John Ashburner

Abstract—A method is presented for the analysis of dynamic positron emission tomography (PET) data using sparse Bayesian learning. Parameters are estimated in a compartmental framework using an over-complete exponential basis set and sparse Bayesian learning. The technique is applicable to analyses requiring either a plasma or reference tissue input function and produces estimates of the system's macro-parameters and model order. In addition, the Bayesian approach returns the posterior distribution which allows for some characterisation of the error component. The method is applied to the estimation of parametric images of neuroreceptor radioligand studies.

Index Terms—Basis pursuit, compartmental models, DEPICT, nonnegative least squares, time course analysis.

I. INTRODUCTION

DYNAMIC positron emission tomography (PET) is a technique that can be used to study the function of the brain *in vivo*. Reconstructed images contain measurements of the concentration of the injected tracer over time. These data can then be analyzed quantitatively using compartmental models [1] where the compartments represent distinct pools of tracer substance, for example tracer in the blood, free in tissue, or bound to specific receptors. Compartmental analysis assumes that the injected tracer causes negligible perturbation in the process under study and that each compartment contains a homogeneous tracer concentration (instantaneous mixing within compartments). These models can have different inputs: the blood/plasma input model can be used when blood/plasma tracer input measurements are available in addition to the dynamic PET image data. If a suitable reference tissue exists within the image then a reference tissue input model may be used.

Manuscript received October 25, 2007; revised February 21, 2008. First published July 22, 2008; current version published August 27, 2008. This work was supported by the National Science Council (Taiwan) under Grant NSC-94-2118-M-001-014 and Grant NSC-95-2118-M-001-003. *Asterisk indicates corresponding author.*

J.-Y. Peng is with the Institute of Statistical Science, Academia Sinica, Taipei 11529, Taiwan and also with the Department of Computer Science and Information Engineering, National Taiwan University, Taipei 10617, Taiwan (e-mail: jaston@stat.sinica.edu.tw).

*J. A. D. Aston is with the Institute of Statistical Science, Academia Sinica, Taipei 11529, Taiwan and also with the Centre for Research in Statistical Methodology, Warwick University, CV4 7AL Coventry, U.K. (e-mail: jaston@stat.sinica.edu.tw).

C.-Y. Liou is with the Department of Computer Science and Information Engineering, National Taiwan University, Taipei 10617, Taiwan.

R. N. Gunn is with the Clinical Imaging Centre, GlaxoSmithKline, W12 0NN London, U.K. and also with the Department of Engineering Science, University of Oxford, OX1 3PJ Oxford, U.K.

J. Ashburner is with the Wellcome Centre for Neuroimaging, Institute of Neurology, UCL, WC1N 3BG London, U.K.

Digital Object Identifier 10.1109/TMI.2008.922185

PET compartmental analysis can be divided into model-driven methods and data-driven methods as defined in [2]. When using model-driven methods, a specific compartmental model is selected *a priori* and the model parameters estimated from the data. For data-driven methods, both the model order and model parameters are obtained from the data, hence *a priori* assumptions about the number of compartments and the topology of particular compartmental structure are not needed. While data driven methods may not have the power of model-driven methods when the correct model is used, data-driven methods do allow a more general approach to be taken to data analysis. Bias that might be introduced by using an incorrect model choice will not be present, and in addition biological information may be obtained about the system by examining the implied model determined from the data driven methods. Thus, data-driven approaches allow for the extension of biological questions from those about parameters conditioned on the model, to those concerning the parameters within a general class of models, without being restricted to one particular model. A class of data-driven methods constructs the modelling approach in terms of a set of basis functions. The new algorithm presented in this paper falls into this category.

The differential equations that are derived from compartmental models can be solved using exponential functions. Data-driven analysis proceeds using a predefined range of possible exponential functions, together with the input function (from blood/plasma or reference tissue), to obtain the weight for each basis function. Since the number of basis functions (in the hundreds) is often much larger than the number of identifiable physiological compartments (usually one to three), the correct weight vector should be sparse. Such analysis can be applied voxel-by-voxel to obtain parametric images, or applied to region of interest (ROI) data to obtain regional parameter estimates. From the coefficients of these exponential basis functions, it is straightforward to derive the relevant macro-parameters such as volume of distribution or binding potential [1].

Two data-driven methods, which have previously approached kinetic modelling with exponential bases, are spectral analysis [3] and basis pursuit denoising via DEPICT [2]. Spectral analysis is a data-driven method for blood/plasma input models. At its core is the nonnegative least squares (NNLS) algorithm [4]. Spectral analysis finds a solution by noting that in the blood/plasma input model weights are nonnegative for nearly all tracer compartmental structures [5], and the nonnegativity constraint in the NNLS algorithm is sufficient to ensure a sparse and unique solution.

Data-driven estimation of parametric images based on compartmental theory (DEPICT) [2] solves the over-complete ex-

ponential basis problem using a mixed norm objective function containing the residual sum of squares (2-norm) and a penalty term of the sum of the weights (1-norm) using basis pursuit denoizing [6]. The 1-norm is used as a surrogate for minimization of the 0-norm for sparsity. Basis pursuit obtains the weight vector by minimizing the mixed norm objective function, the relative contribution of the two terms also being obtained automatically from the data (via cross validation). Since basis pursuit does not need the nonnegativity constraint to ensure sparsity, it can be applied to both blood/plasma input models and reference tissue models (the latter can have negative weights).

The purpose of this work is to investigate whether the application of sparse Bayesian learning (SBL) [7], [8] techniques can improve the estimation of parameters and facilitate the straightforward estimation of model errors. SBL assumes an error distribution for the linear model, and also prior distributions for each weight, and defines hyper-parameters that govern these distributions. The hyper-parameters themselves also have prior distributions, and the analysis proceeds using Bayesian probability theory. Sparsity in this case is achieved by using the concept of automatic relevance determination (ARD) [9], where each weight has its own governing hyper-parameter, leading to individual distributions for each weight. Hence relevance of each weight (basis) is then determined by its distribution, with irrelevant weights having distributions that peak at zero.

II. COMPARTMENTAL MODELS

The two main classes of compartmental model input used in PET are reviewed: the blood/plasma input model and the reference tissue input model. Since only data-driven methods are considered here, the prime emphasis will be on the corresponding linear regression problems implied by the model differential equations.

A. Blood/Plasma Input Models

In this model, it is assumed that in addition to the dynamic PET data, measurements of the blood/plasma input function are also available, which is the concentration of tracer in blood/plasma over time. From the dynamic PET data, the tissue time-activity function can be obtained, which is the total concentration of tracer in tissue. For parametric imaging it is assumed that the blood/plasma input is the same at all voxels, but that each voxel can have a different number of compartments. Since each voxel is analyzed independently, the model for one voxel is presented, and extension to the whole image can be readily inferred.

As given in [1], in the blood/plasma input compartmental model, the tissue time-activity function $C_T(t)$ is modelled as the convolution of the blood/plasma input function $C_P(t)$ with the tissue response function $h(t)$ (a weighted sum of exponential functions)

$$C_T(t) = C_P(t) * h(t) \quad (1)$$

$$h(t) = \sum_{i=1}^{n_c} w_i \exp(-\nu_i t) \quad (2)$$

where n_c is the number of compartments, w_i and ν_i are the weights and exponents respectively, and $*$ represents the convolution operator.

The tissue time-activity function $C_T(t)$ is measured for m discrete (not necessarily equally spaced) time frames ending at times t_1, \dots, t_m . In this paper, the data for $C_P(t)$ is measured nearly continuously in time, although it is possible to extend the methods to discretely measured data for $C_P(t)$. Reformulating (1)

$$C_T(t) = \sum_{i=1}^{n_c} w_i \int_0^t C_P(\tau) \exp(-\nu_i(t - \tau)) d\tau. \quad (3)$$

To solve this linear equation, n_c is replaced by a large number n and the ν_i 's take values in the range of all possible kinetics (basis functions that are not present will have weight $w_i = 0$ in the solution), and define the discrete target (the measured data) as

$$y_j \equiv \frac{1}{t_j - t_{j-1}} \int_{t_{j-1}}^{t_j} C_T(t) dt \quad (4)$$

where $t_0 = 0$. If the basis/design matrix is defined as

$$\Phi_{ji} \equiv \frac{1}{t_j - t_{j-1}} \int_{t_{j-1}}^{t_j} \int_0^t C_P(\tau) \exp(-\nu_i(t - \tau)) d\tau dt \quad (5)$$

then the problem can be expressed as the linear equation

$$\mathbf{y} = \Phi \mathbf{w} \quad (6)$$

where each column of Φ is a basis vector. Solving this equation yields the number of compartments n_c (as the number of nonzero weights) and the parameters w_i and ν_i from which the macro-parameters can be derived.

In blood/plasma input model PET analysis, the macro-parameter of interest is the total tissue volume of distribution

$$V_T \equiv \int_0^\infty \sum_{i=1}^n w_i \exp(-\nu_i t) dt = \sum_{i=1}^n \frac{w_i}{\nu_i}. \quad (7)$$

It is defined as the integral of the tissue response function, which can be calculated from w_i and ν_i .

B. Reference Tissue Input Models

A reference tissue input model [10], [1] may be used when a suitable reference region exists. The input function is obtained as the average time activity curve for the reference region. If a large region is taken, the average time-activity curve can be assumed to be essentially noise free.

$C_T(t)$ is modelled as the convolution of the reference tissue input function $C_R(t)$ with the response function $h(t)$ (in this model a sum of exponentials plus an impulse at $t = 0$)

$$\begin{aligned} C_T(t) &= C_R(t) * h(t) \\ &= w_0 C_R(t) + C_R(t) * \sum_{i=1}^{n'_c} w_i \exp(-\nu_i t) \end{aligned} \quad (8)$$

$$h(t) = w_0\delta(t) + \sum_{i=1}^{n'_c} w_i \exp(-\nu_i t) \quad (9)$$

where $\delta(t)$ is the dirac delta function and n'_c is the equivalent number of compartments with the addition of reference tissue compartments. The same (6) can be obtained by discretizing over the time frames and replacing $n'_c = n$ as before, with one extra basis function, $C_R(t)$, in Φ . The linear problem for this model differs from the blood/plasma input model in that it is possible for the basis function weights w_i to be negative [2].

It should be noted that (1) and (8) are very much related. The reference region is used as a surrogate for the input function, with the reference region being available directly from the data rather than from a separate measurement. However, it is often difficult to obtain a reference region, as it must be devoid of specific binding of the tracer [1], and thus plasma input methods are preferable when plasma input function data is available. Also if very complex structures are estimated from the data, while macro parameter estimates will still be available, the exact compartmental configuration may not be identifiable.

The macro-parameter of interest for the reference tissue input model is the binding potential for the tracer, defined as

$$\text{BP}_{\text{ND}} \equiv w_0 + \sum_{i=1}^n \frac{w_i}{\nu_i} - 1 \quad (10)$$

the integral of the response function minus one. For a radioligand, this is a function of the tissue free fraction, the receptor availability and the equilibrium dissociation rate constant of the radioligand. For the definition of the radioligand nomenclature see [11].

III. DATA-DRIVEN METHODS

All the methods presented here try to solve the general linear regression problem, (6), to obtain model order and basis function weights. The spectral analysis [3] method uses NNLS [4], and the DEPICT [2] method uses basis pursuit denoising [6]. A new algorithm is proposed that uses SBL [7], [8], a statistical learning algorithm based on Bayesian probability theory.

It can be noted that while all the methods find different ways to constrain the solution to the overcomplete linear fit problem, the same underlying basis set is used for each method. The basis is selected over a suitable kinetic range, selected between values from half-life decay (slowest exponential rate) to arterial transfer (fastest exponential rate). In practice, the basis is usually brought in a little from the half-life decay to reduce bias, thus the basis used for all methods in this paper was generated using the range from 8×10^{-4} to 0.01 sec^{-1} , which is the default range in the DEPICT software. However, it should be noted that choice of the range can be problematic, and is one of the known sources of bias in the output of NNLS. It has also recently been shown that basis pursuit also is sensitive to the choice of basis range [12], and thus a method with little bias would be of considerable interest.

The following methods are all applied voxel-by-voxel to generate parametric images, thus each individual voxel is assumed to be analyzed separately and their results independent. There is no loss of generality since voxel based analysis is just ROI

analysis with a small ROI, namely a voxel. However, traditional larger ROI analysis would proceed identically.

A. Spectral Analysis

Spectral analysis was the first data-driven method developed for PET and can be used for plasma input function studies. Spectral analysis formulates the linear problem as above and solves the problem using the NNLS algorithm

$$\min_{w_i \geq 0, 1 \leq i \leq n} \|\mathbf{y} - \Phi \mathbf{w}\|_2^2. \quad (11)$$

Least squares with nonnegative constraints is used here because, by definition, the weights of basis functions in this model must be positive, or zero if that basis function is not present. The constraints have the effect of ensuring a sparse solution. Since the time-activity curve and the basis vectors in the design matrix are, in general, all positive (since they are all concentrations), the nonnegative constraint prevent the case where one “negative” weighted basis almost cancels out another (spurious) positive weighted basis.

It is not possible to study the accuracy of the weight micro-parameter estimates obtained with NNLS directly, since the asymptotic theoretical statistical properties of NNLS are not well known. However, general bootstrap resampling procedures can be applied to get an estimate of the possible variance of weight solutions found conditional on the current PET data [13], with no assumptions about the error distribution. However, bootstrap procedures can require many repetitions. For example, in the case of NNLS about a thousand bootstrap samples are needed, which means the same algorithm has to run on a thousand times more data to obtain an error estimate.

B. DEPICT

The DEPICT method can be used for both blood/plasma input models and reference tissue input models. DEPICT formulates the solution to the linear regression problem in terms of a mixed norm objective function

$$\min_w (\|\mathbf{y} - \Phi \mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_1). \quad (12)$$

The DEPICT algorithm includes a cross-validation procedure to obtain the regularization parameter λ needed in the basis pursuit objective function.

To enable estimation and to avoid dense weight vectors resulting from fitting of over-complete basis/design matrices, basis pursuit uses a weight vector 1-norm term in its objective function, which is an approximation to model order (0-norm). Both the error of fit and 1-norm are minimized together. A value for the regularization parameter, which determines the relative importance of the error of fit term and 1-norm term, is estimated from the data by leave one out cross validation.

Like NNLS in spectral analysis, the error properties of estimates obtained from basis pursuit are not well known, mainly due to the addition of the 1-norm term in the objective function. Bootstrap resampling procedures can be applied to DEPICT to obtain the probability distribution [14], where about two hundred samples are needed to obtain an accurate estimate.

C. SBL

SBL, originally developed in the machine learning literature [7], solves the supervised learning problem

$$y_j = f_j(\mathbf{x}_i; w_i) + \varepsilon_j \quad (13)$$

where the basis function $f_j(\mathbf{x}_i; w_i)$ can be of any general form parameterized by model parameters \mathbf{w} , and ε_j are assumed to be independent samples from a zero-mean Gaussian noise process. SBL outputs the posterior distribution of the model parameters $p(\mathbf{w} | \mathbf{y})$, from which an estimation of the model parameters \mathbf{w} and the error component can be determined.

For application to PET data analysis, the basis functions are chosen to be of the form

$$f_j(\mathbf{x}_i; w_i) = \Phi_{ji} w_i \quad (14)$$

where Φ is the basis/design matrix for the PET data. Treating the targets y_j and disturbances ε_j as vectors, we have

$$\mathbf{y} = \Phi \mathbf{w} + \varepsilon \quad (15)$$

which is the linear equation corresponding to (6) with an additive noise assumption. Using the tissue time-activity function $C_T(t)$ to obtain the target \mathbf{y} , and input function $C_P(t)$ or $C_R(t)$ with exponential coefficients ν_i to obtain the basis/design matrix Φ , SBL is used to solve for the basis function weights \mathbf{w} .

SBL solves the linear (15) by assuming that both the disturbance ε and weights \mathbf{w} are Gaussian distributed with zero mean, and incorporates two sets of hyper-parameters: α_i as the inverse variance of each w_i and β as the inverse variance of the noise. Setting a separate variance for each weight parameter is known as ARD [9], as irrelevant parameters will have $\alpha_i = \infty$, yielding a zero-mean Gaussian distribution with zero variance, and hence $w_i = 0$ with probability one. So the weight priors are given by

$$p(w_i | \alpha_i) \sim \mathcal{N}(0, \alpha_i^{-1}), \text{ for each basis function } i, \quad (16)$$

and noise distribution

$$p(\varepsilon_j | \beta) \sim \mathcal{N}(0, \beta^{-1}), \text{ for each time point } j. \quad (17)$$

The data distribution (given model parameters) is then Gaussian

$$p(\mathbf{y} | \mathbf{w}, \beta) \sim \mathcal{N}(\Phi \mathbf{w}, \mathbf{B}^{-1}) \quad (18)$$

where \mathbf{B} is just a matrix with β on the diagonals and zeros elsewhere. The prior distributions of the hyper-parameters α_i and β (hyper-priors) are assumed to be log-uniform, which is the standard uninformative prior for scale parameters [15].

Since the noise in the tissue time-activity curve is actually time varying and signal dependent, appropriate weighting is needed for the target \mathbf{y} and corresponding entries in the basis matrix Φ before running SBL, so the noise at each time point can be treated as coming from the same distribution. This is implicitly assumed in both previous methods too, as they also use signal dependent weighting. The true time course weighting

(noise structure) depends on many functions of both the data itself, and the reconstruction algorithm used, but many algorithms have used frame dependent counts as weights successfully [3], [16], [2] and this approximation will also be used here. This weighting correction is implicitly assumed in subsequent discussion.

SBL uses a type-II maximum likelihood procedure [15]. The posterior distribution of the weights (given data and hyper-parameters) can be calculated analytically given that both the data and weight prior distributions are Gaussian

$$p(\mathbf{w} | \alpha, \beta, \mathbf{y}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (19)$$

where $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ can be computed from $\mathbf{A} = \text{diag}(\alpha), \beta$

$$\boldsymbol{\Sigma} = (\beta \Phi^T \Phi + \mathbf{A})^{-1}, \quad (20)$$

$$\boldsymbol{\mu} = \beta \boldsymbol{\Sigma} \Phi^T \mathbf{y}. \quad (21)$$

The hyper-parameter posterior can usually be approximated by a delta function [7]

$$p(\alpha, \beta | \mathbf{y}) \approx \delta(\alpha_{\text{MP}}, \beta_{\text{MP}}) \quad (22)$$

where $\alpha_{\text{MP}}, \beta_{\text{MP}}$ is the most probable value obtained after the maximization procedure converges. The posterior $p(\mathbf{w} | \mathbf{y})$ can be approximated by $p(\mathbf{w} | \alpha_{\text{MP}}, \beta_{\text{MP}}, \mathbf{y})$. The reason for choosing type-II maximum likelihood instead of directly maximizing the parameter posterior is because the shape of the parameter posterior distribution with the assumed priors is extremely multimodal, for details see [9].

In the type-II maximum likelihood procedure, firstly an appropriate initial value for the hyper-parameters α and β is chosen (e.g., $\alpha_i = n \text{Var}(\mathbf{y})^{-1}, \beta = \text{Var}(\mathbf{y})^{-1}$, where n is the number of basis functions). Incorrect *a priori* values for the initial parameters cause slow down in convergence for the algorithm. For the first step the weight parameter posterior distribution is obtained by calculating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ from the current values of the hyper-parameters. For the second step the hyper-parameter posterior

$$p(\alpha, \beta | \mathbf{y}) \propto p(\mathbf{y} | \alpha, \beta) p(\alpha) p(\beta) \quad (23)$$

is maximized with respect to the hyper-parameters α and β by differentiating

$$p(\mathbf{y} | \alpha, \beta) = \int p(\mathbf{y} | \mathbf{w}, \beta) p(\mathbf{w} | \alpha) d\mathbf{w} \quad (24)$$

with respect to $\log(\alpha)$ and $\log(\beta)$, as the prior distributions of the hyper-parameters α_i and β are assumed to be log-uniform, as mentioned above. This yields the update equations

$$\alpha_i^{\text{new}} = \frac{\gamma_i}{\mu_i^2} \quad (25)$$

$$\beta^{\text{new}} = \frac{m - \sum_i \gamma_i}{\|\mathbf{y} - \Phi \boldsymbol{\mu}\|^2} \quad (26)$$

$$\gamma_i \equiv 1 - \alpha_i \Sigma_{ii}. \quad (27)$$

where m is the number of time points and intuitively γ_i gives an indication of the relevance of that basis function to the data such

that it can be used in the calculation β as a surrogate for degrees-of-freedom (see [7, Appendix A] for derivation details and the formal meaning of γ_i). Incidentally it is exactly because $p(\mathbf{y})$, the normalizing factor for (23), cannot be calculated that we have to use (22) to approximate the hyper-parameter posterior.

The two steps of calculating $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ and updating $\boldsymbol{\alpha}, \beta$ are iterated in turn until the hyper-parameters converge. ARD is achieved procedurally by setting a large number as the threshold for α_i (e.g., 10^{12}). Once an α_i exceeds the threshold, it is logically set to infinity and the corresponding basis discarded in further iterations, effectively reducing the computation time of subsequent iterations.

On the surface it appears that SBL uses twice as many parameters as non-Bayesian methods, but from the Bayesian perspective the hyper-parameters should be correctly integrated out at the end and the correct parameter posterior obtained, thus overfitting should not be a concern [7]. The SBL method calculates an approximation for $p(\boldsymbol{\alpha}, \beta | \mathbf{y})$, which is then used to calculate the weight posterior

$$p(\mathbf{w} | \mathbf{y}) = \int p(\mathbf{w} | \boldsymbol{\alpha}, \beta) p(\boldsymbol{\alpha}, \beta | \mathbf{y}) d\boldsymbol{\alpha} d\beta \quad (28)$$

hence in this sense the hyper-parameters $\boldsymbol{\alpha}$ and β are “integrated out” of the final output, which consists of the posterior distribution of the weights. The fact that SBL uses the mode of $p(\boldsymbol{\alpha}, \beta | \mathbf{y})$ to approximate the whole distribution makes the integration less apparent. This also implies that only $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ should be considered the output of SBL, and it is only through the choice of approximating the hyper-parameter posterior by a delta function that $\boldsymbol{\alpha}$ and β converge to a single value.

The error estimate from the model is calculated from the posterior distribution through the parameter $\boldsymbol{\Sigma}$. As will be discussed later, this does not yield the same parameter errors as the bootstrap which is used for error calculation in both NNLS and basis pursuit, but does provide useful information as to the error associated with the SBL model assumptions and the data itself.

In the presented method, SBL is used in place of other solvers for the linear regression problem. For the blood/plasma input models where the weights are constrained to be nonnegative, each negative element in $\boldsymbol{\mu}$ of SBL is multiplied by a small negative value (e.g., -0.01) at every iteration to ensure the chosen weight will be positive. This is equivalent to restarting the iteration with positive weight mean, assuming that the “correct” fitting will involve only positive components. This step is heuristic as convergence is no longer absolutely guaranteed when restarting is used, as it might be possible to encounter a loop of restarting and becoming negative. However, when signal is present and given the nature of the priors to favour sparsity, this is unlikely to happen. Indeed in the simulated and real data analysis, convergence always occurred.

IV. DATA ANALYSIS

First the presented method is compared against previous methods using 1-D data simulated from real input functions, then analysis of real PET data is performed.

A. One-Dimensional Data Simulations

In this section simulated data derived from a real blood/plasma input function is used to test the three methods. First, noise level performance is assessed and then two sets of 1-D time activity curve (TAC) data are generated using the one-tissue compartment model and two-tissue compartment model respectively. Appropriate ranges of values (based on actual analysis results of the corresponding PET data) for basis function weights w and exponential coefficients ν are used in these simulations, to assess parameter dependent bias for the methods. While of interest is the relative performances of the data-driven methods, the true model is also fitted to the data using nonlinear least squares (NLLS), using a Gauss-Newton method, for comparison (see [17] for more details on NLLS).

1) *Relative Noise Level Performance*: First, the two-tissue compartmental model plasma input and reference tissue input simulations that were performed in [2] are carried out for the three methods. Here, a different plasma input function is used to generate the data, and a larger number of time frames (32) are also generated than the corresponding analysis in [2]. This was to allow direct comparison with the measured [^{11}C]-diprenorphine data that is analyzed in Section IV-B. However, the same parameter values are used. Thus, following [2], the rate constants $K_1 = 6.7 \times 10^{-3}$ (mL plasma) $\cdot\text{sec}^{-1} \cdot$ (mL tissue) $^{-1}$, $k_2 = 3.3 \times 10^{-3} \text{ sec}^{-1}$, $k_3 = 6.7 \times 10^{-3} \text{ sec}^{-1}$ and $k_4 = 1.7 \times 10^{-3} \text{ sec}^{-1}$ for the target tissue and $K'_1 = 6.7 \times 10^{-3}$ (mL plasma) $\cdot\text{sec}^{-1} \cdot$ (mL tissue) $^{-1}$, $k'_2 = 6.7 \times 10^{-3} \text{ sec}^{-1}$, $k_5 = 0.017 \text{ sec}^{-1}$ and $k_6 = 0.017 \text{ sec}^{-1}$ for the reference tissue. These correspond to weight and exponential coefficient pairs of ($1.8 \times 10^{-3} \text{ sec}^{-1}$, 0.011 sec^{-1}) and ($4.9 \times 10^{-3} \text{ sec}^{-1}$, $5 \times 10^{-4} \text{ sec}^{-1}$) for the two components in the plasma case, and there will be three components in the reference tissue case. The volume of distribution for the target region is 10 and the binding potential for the target region, given the reference region, is 4. As in [2], the noise added to the data was proportional to the true data divided by the frame duration, with the proportionality varying from 0.01 to 5.12, giving noise levels from lower than typical ROI analysis to higher than voxel level analysis. We generated 2000 realizations for each noise level. In addition, the reference tissue input function simulation from [2], with the same noise levels and number of realizations was performed, and the results from basis pursuit and the presented method compared.

In Figs. 1 and 2, graphs of bias and model order histograms are presented. The bias graphs show the mean result over the 2000 realizations along with standard error bars. Percentage bias can be derived by dividing by the true value. Histograms of model order reflect the number of times a model with a particular order was chosen across the 2000 realizations. The values of the color bar refer to the proportion falling into each box and thus the histograms should be compared down each column with separate columns (indicating various noise levels) being independent.

As can be seen in Fig. 1, there is positive bias for the NNLS algorithm at most noise levels, as is well known for the algorithm. It is also seen that the basis pursuit algorithm has positive bias at higher noise levels for this set of parameter values, although

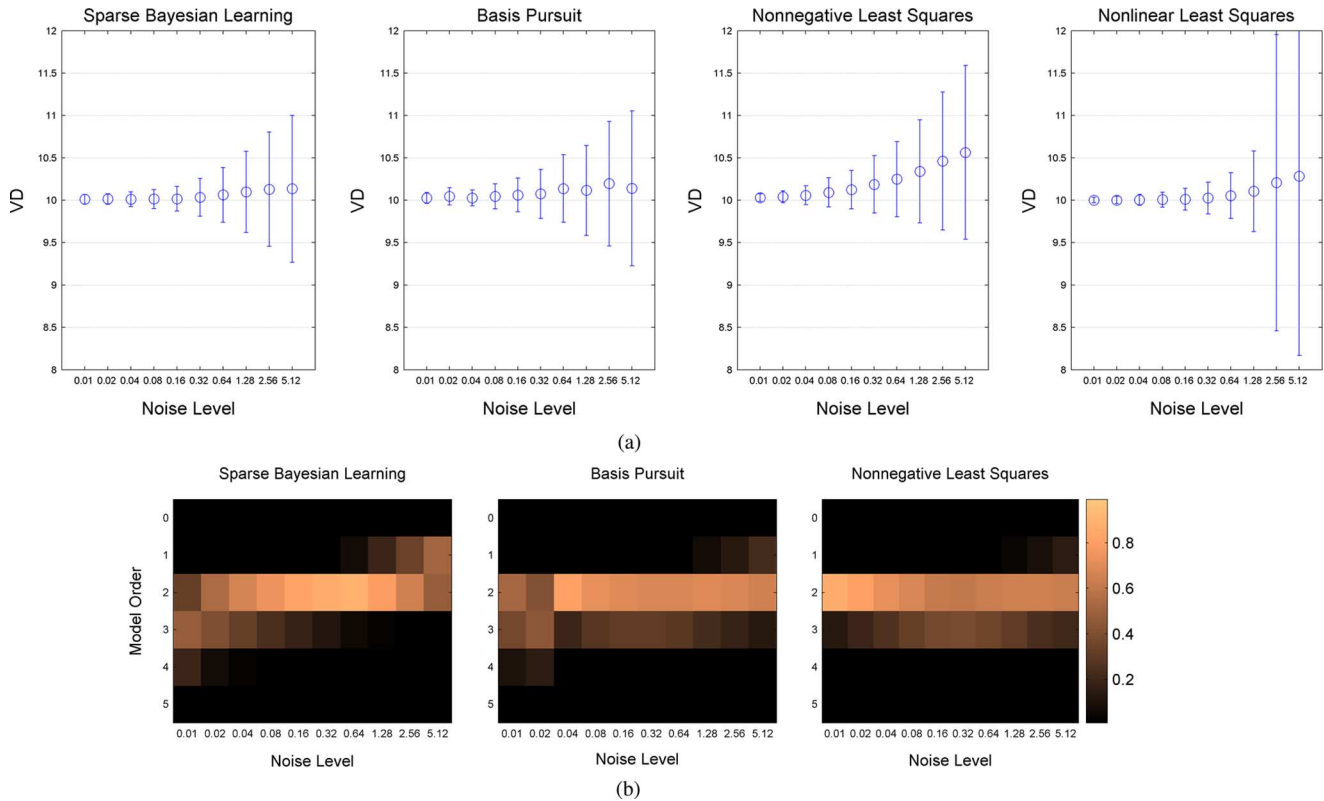


Fig. 1. Noise level simulations for two-tissue compartments using a plasma input function. (a) Volume of distribution V_T estimates and error bars. (b) Model order histograms.

it has little bias for low noise levels. The findings are in agreement with the findings in [2] although slightly different MSEs are found (see Table I), due to the different input function and also the different number of frames used. SBL does not contain as much bias as the other methods are subject to at higher noise levels, and the variance of the parameter estimates is slightly lower for SBL than for basis pursuit. This results in large reductions in MSE for higher noise levels using SBL. Thus, while both basis pursuit and SBL are essentially unbiased for ROI analysis (lower noise levels), SBL should be preferred at higher noise levels. In comparison to the fixed model estimated with NLLS, SBL performed well and is not subject to the high variance that can affect NLLS at higher noise levels.

The model orders are more stable for SBL at higher noise levels, although slightly less stable at very low noise levels, given that noise is assumed to be in the system for SBL. The model order estimate for NLLS is not given as it is fixed at the true model. It was also found that the value of β^{-1} corresponded well with the variance of the noise added to the data. However, the value was slightly biased downwards, about 10% lower than the true value.

In Fig. 2, there is a positive bias for basis pursuit when the noise level starts to get higher, as was seen in the original paper [2]. However, again this bias is less using the SBL algorithm. In addition, larger variances result from using basis pursuit, and as such the MSE (see Table II) is greatly reduced using SBL. Again SBL gives good performance relative to NLLS without the high variance that arises at higher noise levels.

2) *Parameter Dependent Bias—One-Tissue Compartment Model:* In the next two subsections, the plasma input function

will be considered so a comparison can be made between all three methods. While from the simulations in Section IV-A1, it can be seen that only the presented method is essentially unbiased for those parameter values, it is of interest to determine if the amount of bias present is dependent on the parameter values themselves. The noise level used in these simulations is comparable to voxel noise levels in the real data (and corresponds to a value of 1 on the noise simulation graphs in Section IV-A1).

For the first simulation, the tissue is assumed to contain only one tissue compartment, and the corresponding linear model is then set to have one component with w varying from 0.0001 sec^{-1} to 0.02 sec^{-1} , and ν varying from 0.0008 sec^{-1} to 0.012 sec^{-1} . For each combination of w and ν values 1000 realizations are generated, and the noise level is scaled to the data value at each time frame and inversely scaled to the frame length, specifically the noise variance at each time frame equals the data value at that time point divided by the frame length, which roughly corresponds to residuals observed in real data analysis. A total number of $30 \times 40 \times 1000$ time activity curves are thus generated. Each TAC is then analyzed by all three methods using the original input function, the analysis results (macro-parameters) for each combination of w and ν values are averaged over 1000 trials, and presented on a plane with w and ν axis. The data is also analyzed against NLLS using the true model.

Fig. 3(a) shows the absolute average V_T bias divided by the true V_T value (absolute percentage bias) at each point for all three methods, truncated at 3%. The NNLS algorithm has the highest positive bias overall, whereas the basis pursuit method has both positive and negative bias. The method of SBL has the lowest absolute bias over the whole range. It can be seen

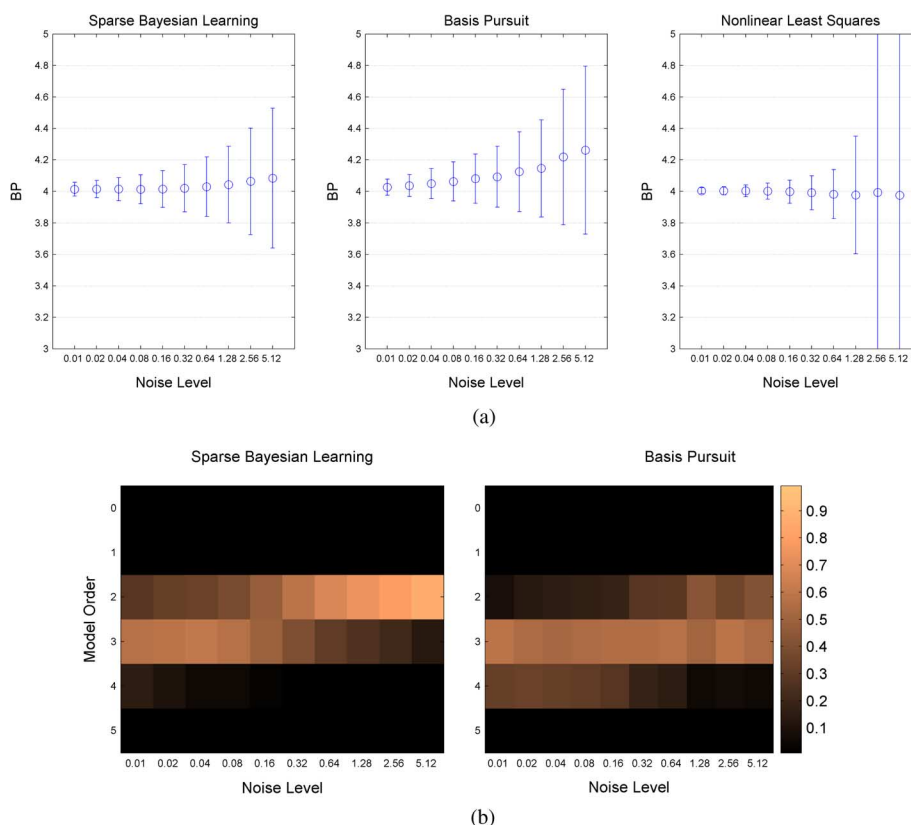


Fig. 2. Noise level simulations for two-tissue compartments using a reference tissue input function. (a) Binding potential BP_{ND} estimates and error bars. (b) Model order histograms.

TABLE I
MSE FOR THE NOISE LEVEL SIMULATIONS USING A PLASMA INPUT FUNCTION. (NLLS USING THE TRUE MODEL IS INCLUDED FOR COMPARISON)

Noise Level	0.01	0.02	0.04	0.08	0.16	0.32	0.64	1.28	2.56	5.12
SBL	0.003	0.004	0.008	0.012	0.021	0.050	0.108	0.239	0.471	0.771
Basis Pursuit	0.005	0.012	0.009	0.024	0.043	0.089	0.178	0.296	0.578	0.856
NNLS	0.003	0.006	0.016	0.038	0.067	0.148	0.259	0.486	0.877	1.370
NLLS	0.001	0.002	0.004	0.008	0.017	0.035	0.076	0.238	3.096	4.562

though that the bias is dependent on the parameters and is not the same across the image for both NNLS and basis pursuit, and to a lesser extent SBL. It is interesting to note that the performance of NLLS and SBL are very similar in these examples, indicating that SBL performs well even against a method where the true model is known.

Fig. 3(b) shows the average model order obtained from each method. SBL has the lowest and most uniform model order compared to the other two methods, and is closest to model order 1, which is the correct model order. Nonnegative least square has the most spurious components, which is related to its high positive bias in those regions, whereas basis pursuit also slightly overestimates the model order.

Table III gives summary statistics for the overall percentage bias in the simulations. SBL can be seen to be the least biased both in terms of percentage and absolute percentage bias. In addition NNLS only contains positive bias. These simulation results show that SBL is overall the most stable method for one-tissue compartment data, with generally smaller bias in \hat{V}_T estimation and model orders that more accurately reflect the underlying data, compared to NNLS and basis pursuit.

3) *Parameter Dependent Bias—Two-Tissue Compartment Model:* A simulation for two-tissue compartment data, similar to the one-tissue compartment simulation above, was also carried out. Values of component weights $w_1 = 0.01 \text{ sec}^{-1}$ and $w_2 = 0.006 \text{ sec}^{-1}$ are chosen, and the component exponential coefficient for both ν_1, ν_2 are varied from 0.0008 sec^{-1} to 0.012 sec^{-1} , yielding $36 \times 36 \times 1000$ TAC realizations. The noise level is chosen to be the same as the one-tissue compartment simulations. Averaged macro-parameter analysis results over all trials are plotted on a plane with ν_1 and ν_2 axis.

Comparing the absolute bias relative to the true V_T value [Fig. 4(a)], truncated at 2%, SBL has the lowest bias overall. NNLS has the most bias while basis pursuit is a little biased in some parameter combinations, but overall is fairly unbiased. The bias in all cases for NNLS is parameter dependent. Basis pursuit is less parameter dependent while SBL is possibly even less so. The overall mean (standard deviation) for the biases in these simulations for the three methods are given in Table IV. It shows that again SBL has the lowest overall absolute percentage bias, but in this case, basis pursuit has both positive and negative bias resulting in low overall bias when averaged. This again

TABLE II
MSE FOR THE NOISE LEVEL SIMULATIONS USING A REFERENCE TISSUE INPUT FUNCTION. (NLLS USING THE TRUE MODEL IS INCLUDED FOR COMPARISON)

Noise Level	0.01	0.02	0.04	0.08	0.16	0.32	0.64	1.28	2.56	5.12
SBL	0.002	0.003	0.005	0.009	0.014	0.023	0.037	0.061	0.119	0.205
Basis Pursuit	0.003	0.006	0.012	0.020	0.031	0.046	0.080	0.117	0.233	0.353
NLLS	0.000	0.001	0.001	0.003	0.005	0.012	0.025	0.140	1.368	13.587

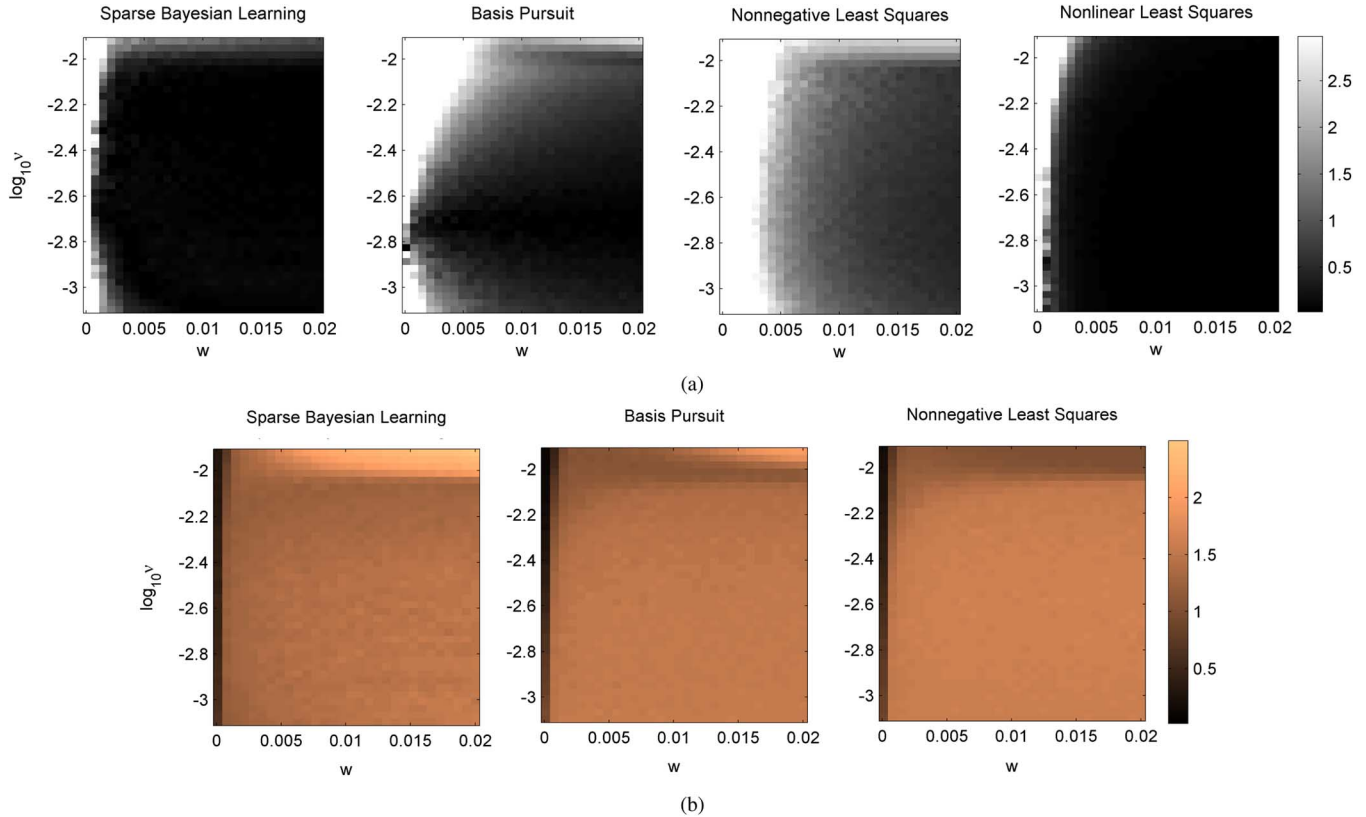


Fig. 3. Comparison of V_T absolute %bias and model order for one-tissue compartmental simulations. The model orders are the averages over the 1000 trials at each point. (a) V_T Absolute %Bias. (b) Model order.

TABLE III
OVERALL STATISTICS FOR ONE-TISSUE COMPARTMENTAL SIMULATIONS

Method	Mean % Bias (Std)	Mean Absolute % Bias (Std)
SBL	-0.681 (3.07)	0.792 (3.04)
Basis Pursuit	-1.28 (4.05)	1.68 (3.90)
NNLS	5.94 (23.8)	5.94 (23.8)
NLLS	4.91 (81.5)	6.42 (81.4)

illustrates the parameter dependency of the bias in basis pursuit. As in the one-tissue compartmental case, there is a similar performance for SBL and NLLS using the known model.

In the model order plot [Fig. 4(b)] the threshold between distinguishing the two components can be clearly seen for all three methods. Overall NNLS and basis pursuit still tend to overestimate the model order, whereas SBL has model order 1 on the diagonal and model order 2 at the corner where the two components are distinguishable.

In comparison to the noise level simulations given in Section IV-A1, the exponential coefficients lie in the lower right hand corner of the image. The bias in this region is consistent with the results from the noise level simulations with the ratio equal to 1.

In conclusion, simulation results show that SBL outperforms both other data driven methods and compares favorably with a model-driven method in terms of parameter estimation bias and model order estimation. This is true for both one-tissue and two-tissue compartment models, over a wide range of possible micro-parameters.

B. [^{11}C]-Diprenorphine Measured PET Data

Four dynamic scans from a measured [^{11}C]-diprenorphine study of normal subjects, for which an arterial input function was available, were analyzed. The subjects underwent 95-min dynamic [^{11}C]-diprenorphine PET baseline scans on the same camera. The subjects were injected with 185 MBq of [^{11}C]-diprenorphine. PET scans were acquired in 3D mode on a Siemens/CTI ECAT EXACT3D PET camera, with a spatial resolution after image reconstruction of approximately 5 mm. Data were reconstructed using the reprojection algorithm [18] 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




Fig. 4. Comparison of V_T absolute % bias and model orders for two-tissue compartmental simulations. The model orders are the averages over the 1000 trials at each point. (a) V_T Absolute %Bias. (b) Model order.

TABLE IV
OVERALL STATISTICS FOR TWO-TISSUE COMPARTMENTAL SIMULATIONS

Method	Mean % Bias (Std)	Mean Absolute % Bias (Std)
SBL	-0.0543 (0.212)	0.161 (0.149)
Basis Pursuit	-0.00415 (0.449)	0.351 (0.281)
NNLS	0.756 (0.246)	0.756 (0.246)
NLLS	0.155 (0.183)	0.179 (0.159)

on the dynamic [^{11}C]-diprenorphine PET images. The computational time of one scan for the three methods were, NNLS 36 min, basis pursuit 33 min and SBL 5 h 30 min on a single processor Celeron PC.

As measured input data was available, all three methods were compared, and in each case V_T estimated. In the case of NNLS and SBL, the voxels are all analyzed individually and independently, while for basis pursuit the regularization parameter is globally estimated, but then analysis proceeds independently at each voxel. Results from one of the scans is shown in Fig. 5. The distribution of V_T is moderately similar for each of the methods, although differences do occur. NNLS is higher than SBL, as would be predicted from the simulations due to the positive bias of NNLS. However, basis pursuit yields parameter values even higher than NNLS, which can be most easily seen around the basal ganglia.


Given the bias shown in the simulations this is surprising. However, Fig. 6 immediately reveals the reason for this. The model order of basis pursuit is almost uniformly greater than either NNLS or SBL and as such higher values of V_T result. Problems with high bias for basis pursuit have been witnessed

in other studies as well [12]. The model order is higher in the case of NNLS versus SBL, which would be expected from the simulations, and SBL has model order consistently between one and two. The results for all four of the scans analyzed were very similar, although one showed slightly less positive bias in basis pursuit than the other three (Fig. 5 is one of the scans with higher bias for basis pursuit).


As further investigation of the bias, the dynamic data was spatially smoothed, and the V_T obtained from this data compared to the spatially smoothed original V_T results (data not shown). This gives an indication of the bias in the methods, as the smoothed dynamic should be less noisy and thus the estimation of V_T be less biased, as indicated by the simulation results. It is, however, only an approximation due to partial volume effects that are present in the smoothed data. It was found that SBL outperformed NNLS in regions of high V_T (and was essentially unbiased) while NNLS was marginally better than SBL in regions with very small V_T . However, basis pursuit had much greater bias than either of the other two methods especially in regions of high signal. This was likely due to difficulty in finding a suitable unique regularization parameter.

C. [^{11}C]-Raclopride Measured PET Data


The second PET example considered a dynamic [^{11}C]-raclopride study of the D_2 -receptor distribution in the normal brain. The scanner of reference is the ECAT 953B PET camera (CTI/Siemens, Knoxville, TN) with images reconstructed using the 3-D-RP FBP algorithm [19] with a ramp filter (transaxial resolution ~ 5 mm at the center of the FOV) and corrected using



(a)



(b)



(c)

Fig. 5. $[^{11}\text{C}]$ -Diprenorphine: Image of V_T . The three methods produce moderately similar results, although higher values are associated with basis pursuit. (a) NNLS. (b) Basis pursuit. (c) SBL.

measured attenuation. Images were 128×128 matrices with 2-mm pixel size. Arterial samples were not available, and a reference region placed on the cerebellum was used as input function. Again the analysis was performed voxel-by-voxel, independently, except for the estimation of the global regularization parameter for basis pursuit. The time for computation for the two methods was 11 min for basis pursuit and 2 h 40 min for SBL again on a single processor Celeron PC.

One standard method of analysis for $[^{11}\text{C}]$ -raclopride is to use the simplified reference tissue model (SRTM) [10] which was fitted using a basis function method [16]. Fig. 7 shows the output for the SRTM, SBL, and basis pursuit. They can be seen to be very similar, even though the SRTM method has the additional assumption of a known model, which for the case of $[^{11}\text{C}]$ -raclopride is widely accepted as a good model for the data. Thus SBL is able to produce robust estimates of the parameters even




Fig. 6. $[^{11}\text{C}]$ -Diprenorphine: Image of Model order. The model order is slightly larger in the NNLS and much higher in the basis pursuit method. (a) NNLS. (b) Basis pursuit. (c) SBL.

without assuming the model. In fact Bland-Altman [20] plots of SRTM and SBL show that the methods are in agreement, while there is evidence of bias for the plot of basis pursuit against the SRTM model (data not shown). Again the model selected was slightly more sparse for SBL than for basis pursuit, but the parameter images did not display as much difference as in the $[^{11}\text{C}]$ -Diprenorphine case.

V. DISCUSSION

Through the use of SBL, a data-driven method of determining macro-parameters from PET data has been presented. This method has been shown through simulation to be less biased than previous methods, and also to generate models which are generally sparser. Indeed, using simulations, it has




Fig. 7. $[^{11}\text{C}]\text{Raclopride}$: Comparison between known compartmental structure (SRTM) and the SBL technique. Images of binding potential. (a) SRTM. (b) Basis pursuit. (c) SBL.

been shown that SBL predominantly determines the correct model order from exponential bases at higher noise levels. While the simulations were only carried out for a single scan protocol, the results likely generalize to other protocols. Basis pursuit produces less biased results than NNLS, but can still be biased at higher noise levels. In addition the bias seems to be parameter dependent for both NNLS and basis pursuit at high noise levels. It is noticeable in the real analysis that while

the three methods give similar results in terms of estimation of parameters, there are likely distinct biases in NNLS and basis pursuit, and there are also model order differences present. In addition, as seen in the simulations, the bias is parameter dependent and so bias in real data will depend on the kinetics involved. This is of importance when dealing with comparisons across subjects as the bias may then not be constant between the different scans.

The bias results for basis pursuit deserve particular mention. In the simulations for one and two-tissue compartmental model structures, it was seen that both positive and negative bias could be present, while in the real data with an available input function, the bias was likely positive. After further investigation, it appears that basis pursuit is more sensitive to the weighting of the data than either NNLS or SBL. Whether inadequate weighting either directly affects the data fit itself, or the estimation of the regularization parameter is unclear, but it is likely a combination of both. However, as a single weighting function is used for the entire image (this procedure was also used for the simulations to allow accurate comparison) and as this is unlikely to be correct for all points, it is difficult to know exactly which points will be sensitive to what bias. The resulting stability issues for basis pursuit, in the case of reference tissue models as opposed to input function models, have also been remarked upon in the recent paper by Cselenyi *et al.* [12]. They suggest that the bias issues could also be related to the suitable choice of basis range. However, it should be stressed that the simulation results indicate that the bias is only a problem at higher noise levels.

It could be argued that linear parametric data-driven methods such as the Patlak [21] and Logan [22] plots could be used instead of the kinetic based methods given here. Firstly, these methods do not provide estimates of model order which can be useful especially at lower noise levels or when examining many voxels in a region. In addition, these methods require a user determination of where linearity has been achieved, adding in an additional source of error into the analysis. It has also been shown that methods such as the Logan plot have noise dependent bias [23], which is difficult to account for and thus makes them unattractive when methods with less bias are available.

In addition to estimation of parameters, SBL allows some estimation of errors on the parameters. This is due to the posterior distributions being found through the algorithm. Normally the expectation of the distribution is used as the estimate for the parameters, but use can also be made of the variance of the distribution with knowledge of its exact nature. However, care must be taken when using these estimates. The estimates of the noise in the model (from the β parameter) generally corresponded well with the noise level in the simulations. However, it was slightly biased downwards. Two reasons can be postulated for this. First, in the simulations the simulated parameter values were fixed. SBL assumes that they come from a distribution, thus a certain amount of variance will be assigned to them as opposed to the noise variance estimate. However, it should be recognized that while β is estimated by the algorithm, the parameter variance Σ is really the output of the model, and this incorporates both the variance of the weights and of the noise together.

Another explanation for this is that the basis functions are not actually fixed but are functions of parameters themselves. In this case, there are actually more parameters to be estimated than acknowledged which will bias the error variance estimate downwards. However, given the model is unknown before the start, it is not easy to know how much of an effect this will have. It might be possible to incorporate the final number of bases chosen into the estimate but this would require fundamental re-

vision of the SBL algorithm. In addition, if error estimates of the parameters themselves are of interest, the implied nonlinearity of the model due to the basis function parameters should be accounted for. This could be approximated using methods such as those in [24] and assuming that the model was fixed after it has been chosen by the algorithm.

It should be emphasized that if parameter standard deviation plots were obtained from normal SBL procedures, would be fundamentally different from bootstrap estimates of parameter estimation errors obtained from bootstrapping basis pursuit or NNLS. The bootstrap errors are variances of the parameter estimate conditional on the observed data and its corresponding model fit. For a particular analysis method, it is essentially a property of the method conditional on the data. In contrast, SBL parameter variance is a property of the parameter itself, given all SBL model assumptions (Gaussian errors, prior distributions, fixed basis, etc.). It does not say anything about the SBL method itself, for that a bootstrap procedure for SBL is needed, but rather how dispersed the parameter, treated as a random variable, is distributed conditional on the data. By assuming Bayesian structure in the method, the output result is exactly another Bayesian statement. While the validity of both the bootstrap resampling procedure and the Bayesian structure assumptions can be disputed, nevertheless they both provide information as to how the macro-parameter estimate would vary given the same underlying model with different realizations of noise.

SBL has well-known theoretical properties relating to its convergence [25]. It has been shown to converge and also that any local minima present are related to sparse solutions. In practice, SBL converged in all cases for both the simulated and real data. However, in the case of voxels where there is little signal present, convergence is often slow. SBL, due to the need to compute a matrix inversion in every iteration, is generally slower than either the NNLS algorithm or the basis pursuit algorithm (but faster than the standard NLLS algorithm), but due to the need to perform cross validation, basis pursuit is sometimes slower than SBL when the input data is not sufficiently homogeneous (this is especially apparent in the simulations). Also a new implementation of SBL using a faster algorithm has been developed [26] and this might be adapted to increase the efficiency of the calculations.

SBL does allow for model selection even in the cases where the model underlying the data is complex. Using ARD allows for complex data to be characterized with relative simplicity from the basis by penalizing models towards sparsity. This feature is especially useful when the model is assumed to be complex overall, but the data will not support such complexity due to a small number of measurements, poor signal to noise ratio or other such factors. This type of data is often the case when dealing with biological systems, and certainly true in PET. Indeed due to the Bayesian nature of the model, it may be possible to combine multiple subjects in the same large Bayesian model. This could help alleviate many of the bias-related shortcomings that result from the current procedures, whereby kinetic modeling is done independently for each subject and the results entered into a statistical analysis to produce a p -value. The use of Bayes factors for model comparisons, along with the use of the

evidence framework, given by ARD and the type II ML procedure, for approximating the probabilities, is increasingly being advocated for use within medical analyses [27], and the framework proposed here could be extended to incorporate these ideas for PET analysis.

In conclusion, SBL provides a Bayesian interpretation of model selection in PET data-driven methodology, leading to all its inherent advantages, and presents an interesting alternative perspective to standard analysis methods.

ACKNOWLEDGMENT

The authors would like to thank Dr. A. Hammers and Dr. F. Turkheimer for supplying the measured data. The authors gratefully acknowledge the helpful comments and assistance of the Associate Editor and three anonymous referees. Their comments and suggestions greatly improved the paper.

REFERENCES

- [1] R. N. Gunn, S. R. Gunn, and V. J. Cunningham, "Positron emission tomography compartmental models," *J. Cerebral Blood Flow Metabol.*, vol. 21, no. 6, pp. 635–652, 2001.
- [2] R. N. Gunn, S. R. Gunn, F. E. Turkheimer, J. A. D. Aston, and V. J. Cunningham, "Positron emission tomography compartmental models: A basis pursuit strategy for kinetic modeling," *J. Cerebral Blood Flow Metabol.*, vol. 22, pp. 1425–1439, 2002.
- [3] V. J. Cunningham and T. Jones, "Spectral analysis of dynamic PET studies," *J. Cerebral Blood Flow Metabol.*, vol. 13, pp. 15–23, 1993.
- [4] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*. New York: Prentice-Hall, 1974.
- [5] K. Schmidt, "Which linear compartmental systems can be analyzed by spectral analysis of PET output data summed over all compartments?," *J. Cerebral Blood Flow Metabol.*, vol. 19, pp. 560–569, 1999.
- [6] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1999.
- [7] M. E. Tipping, "Sparse bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, 2001.
- [8] M. E. Tipping, *Bayesian Inference: An Introduction to Principles and Practice in Machine Learning*. New York: Springer, 2004, vol. 3176, Lecture Notes Artificial Intelligence, pp. 41–62.
- [9] D. Mackay, "Bayesian interpolation," *Neural Comput.*, vol. 4, pp. 415–447, 1992.
- [10] A. A. Lammertsma and S. P. Hume, "Simplified reference tissue model for PET receptor studies," *NeuroImage*, vol. 4, no. 3, pp. 153–158, Dec. 1996.
- [11] R. B. Innis, V. J. Cunningham, J. Delforge, M. Fujita, R. N. Gunn, J. Holden, S. Houle, S.-C. Huang, M. Ichise, H. Iida, H. Ito, Y. Kimura, R. A. Koeppe, G. M. Knudsen, J. Knuuti, A. A. Lammertsma, M. Laruelle, R. P. Maguire, M. Mintun, E. D. Morris, R. Parsey, J. Price, M. Slifstein, V. Sossi, T. Suhara, J. Votaw, D. F. Wong, and R. E. Carson, "Consensus nomenclature for in vivo imaging of reversibly binding radioligands," *J. Cerebral Blood Flow Metabol.*, vol. 27, pp. 1533–1539, 2007.
- [12] Z. Cselényi, H. Olsson, C. Halldin, B. Gulyás, and L. Farde, "A comparison of recent parametric neuroreceptor mapping approaches based on measurements with the high affinity PET radioligands [11C]FLB 457 and [11C]WAY 100635," *NeuroImage*, vol. 32, pp. 1690–1708, 2006.
- [13] F. Turkheimer, L. Sokoloff, A. Bertoldo, G. Lucignani, M. Reivich, J. L. Jaggi, and K. Schmidt, "Estimation of component and parameter distributions in spectral analysis," *J. Cerebral Blood Flow Metabol.*, vol. 18, pp. 1211–1222, 1998.
- [14] S. L. Kukreja and R. N. Gunn, "Bootstrapped DEPICT for error estimation in PET functional imaging," *NeuroImage*, vol. 21, pp. 1096–1104, 2004.
- [15] J. Berger, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. New York: Springer-Verlag, 1985.
- [16] R. N. Gunn, A. A. Lammertsma, S. P. Hume, and V. J. Cunningham, "Parametric imaging of ligand-receptor binding in PET using a simplified reference region model," *NeuroImage*, vol. 6, no. 4, pp. 279–287, Nov. 1997.
- [17] G. A. F. Seber and C. J. Wild, *Nonlinear Regression*. New York: Wiley, 1989.
- [18] P. E. Kinahan and J. G. Rogers, "Analytic 3D image reconstruction using all detected events," *IEEE Trans. Nucl. Sci.*, vol. 36, no. 1, pp. 964–968, Feb. 1989.
- [19] D. Townsend, A. Geissbuhler, M. Defrise, E. J. Hoffman, T. J. Spinks, D. L. Bailey, M.-C. Gilardi, and T. Jones, "Fully three-dimensional reconstruction for a PET camera with retractable septa," *IEEE Trans. Med. Imag.*, vol. 10, pp. 505–512, Oct. 1991.
- [20] J. M. Bland and D. G. Altman, "Statistical methods for assessing agreement between 2 methods of clinical measurement," *Lancet*, vol. 8476, pp. 307–310, 1986.
- [21] C. S. Patlak, R. G. Blasberg, and J. D. Fenstermacher, "Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data," *J. Cereb. Blood Flow Metabol.*, vol. 3, pp. 1–7, 1983.
- [22] J. Logan, J. S. Fowler, N. D. Volkow, A. P. Wolf, S. L. Dewey, D. J. Schlyer, R. R. MacGregor, R. Hitzemann, B. Bendriem, S. J. Gatley, and D. R. Christman, "Graphical analysis of reversible radioligand binding from time activity measurements applied to N-C-11-methy(-)cocaine PET studies in human subjects," *J. Cereb Blood Flow Metabol.*, vol. 10, no. 5, pp. 740–747, 1989.
- [23] M. Slifstein and M. Laruelle, "Effects of statistical noise on graphic analysis of PET neuroreceptor studies," *J. Nucl. Med.*, vol. 41, pp. 2083–2088, 2000.
- [24] J. A. D. Aston, R. N. Gunn, K. J. Worsley, Y. Ma, A. C. Evans, and A. Dagher, "A statistical method for the analysis of positron emission tomography neuroreceptor ligand data," *NeuroImage*, vol. 12, pp. 245–256, 2000.
- [25] D. P. Wipf and B. D. Rao, "Sparse bayesian learning for basis selection," *IEEE Trans. Signal Process.*, vol. 52, no. 8, pp. 2153–2164, Aug. 2004.
- [26] M. E. Tipping and A. Faul, "Fast marginal likelihood maximisation for sparse bayesian models," presented at the 9th Int. Workshop Artif. Intell. Stat., Key West, FL, 2003.
- [27] S. N. Goodman, "Toward evidence-based medical statistics. 2: The Bayes factor," *Ann. Internal Med.*, vol. 130, pp. 995–1013, 1999.