Steady-state model of the radio-pharmaceutical uptake for MR-PET

Stefano Pedemonte[1], Brian F. Hutton[2], Simon Arridge[1], and Sebastien Ourselin[1]

[1] The Centre for Medial Image Computing, UCL, London, United Kingdom
[2] Institute of Nuclear Medicine, UCL Hospitals NHS Trust, London, United Kingdom

Abstract. This work explores a fully-automated algorithm for estimation of the uptake of radio-pharmaceutical in brain MR-PET imaging. The algorithm is based on a model of the pharmaceutical uptake coupled with probabilistic models of the PET and MR acquisition systems. In contrast to algorithms that attempt to correct for the Partial Volume Effect (PVE), the problem is tackled here in the reconstruction by means of a probabilistic model of the pharmaceutical uptake. We make use of Hybrid Bayesian Networks to describe the joint probabilistic model and to obtain an efficient optimisation algorithm. We describe solutions adopted in order to mitigate the effect of local maxima and to reduce the sensitivity to the initialisation of the parameters, rendering the algorithm fully automatic. The algorithm is evaluated on simulated MR-PET data and on the reconstruction of clinical PET FDG acquisitions.

1 Introduction

Uncertainty in Emission Tomography is dominated by photon count statistics. It is therefore essential to adopt a probabilistic model of the emission and interaction of the Gamma photons in order to use optimally the information at hand for the quantification of the uptake of the radio-pharmaceutical. Given a generative probabilistic model of the emission imaging system (outlined here in Sec. 2.1), the spatial density of radio pharmaceutical can be estimated by Maximum Likelihood (ML) [1]. However, due to the acquisition being photon-limited, the information about the pharmaceutical density is scarce, thus determining an infinity of equally likely solutions (ill-posedness). Furthermore, optimisation of the likelihood can only be treated currently with greedy optimisation algorithms because of the high dimensionality of the unknown pharmaceutical density, and yet the unknowns (pharmaceutical density in each voxel) are strongly correlated due to the measurement of line integrals by the emission imaging system. For these reasons, algorithms for the optimisation of the likelihood are extremely sensitive to local maxima and to the initial parameterisation; besides, due to ill-posedness, unconstrained optimisation of the likelihood does not converge, posing the unsolved problem of establishing a stopping criterion [2]. A smoothing prior is typically adopted in order to obtain a convergent algorithm; PET images obtained under the assumption of smoothness, however, when overlaid

on an intra-subject MR image, present the problem of *partial volume* or *spill out* effect: an observer can distinguish certain regions in common between the two images, but the PET image is smoother than the MR image, giving the observer the impression of spill out of the pharmaceutical. The partial volume effect, which is nothing else but the *bias* of the estimate of the pharmaceutical density, highlights how it is problematic to quantify the uptake of pharmaceutical in a region of interest obtained by segmentation of the MR image. Partial volume *correction* algorithms attempt to deconvolve the PET image in order to estimate the dose in one or more regions of interest by introducing certain assumptions about uniformity of the uptake in each region. However, since partial volume is due to the photon count statistics, the problem cannot be solved in the domain of the PET image for the reason that the PET reconstruction is a point estimate of the high dimensional likelihood function that describes the uncertainty of the measurement. Though a number of partial volume *correction* algorithms have been proposed and are sometimes considered state of the art, concerns about the efficacy of such algorithms are starting to be raised [3].
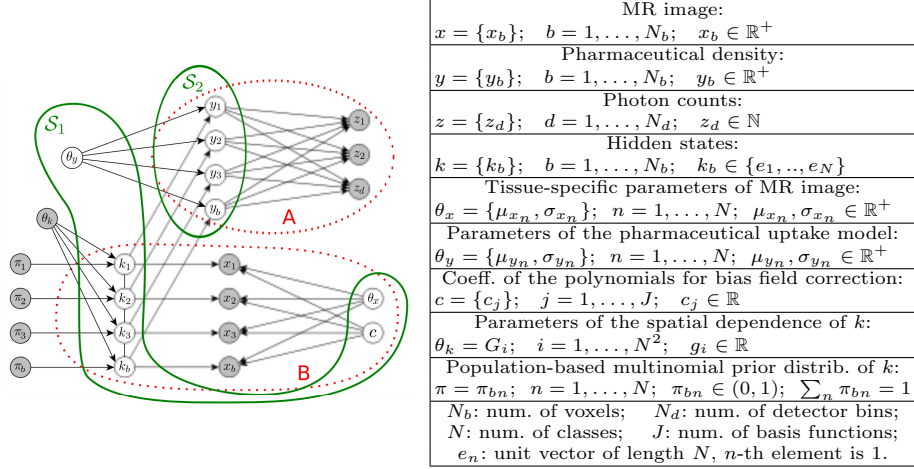
In this paper we extend the work proposed by [4], which replaces the assumption of smoothness with a parametric model of the pharmaceutical uptake. Here we extend and improve on both the stability and convergence of the pharmaceutical uptake model of [4], making it suitable for clinical usage. Full automation is obtained by means of integration of population-based prior information, by the use of a local contextual model and by improving the robustness to intensity inhomogeneity of the MR image. To the best of our knowledge, this paper presents the first algorithm that combines models of the MRI and PET acquisition systems, pharmaceutical uptake, MRI intensity inhomogeneity correction and local contextual information in a unified framework.

2   Generative Model of the Pharmaceutical Uptake in Brain Tissue

This section will first introduce the Poisson model of the PET acquisition system (Sec. 2.1), followed by the hidden-state model of the pharmaceutical uptake (Sec. 2.2) and the hidden-state model of the MR acquisition system (Sec. 2.3). Sec. 2.4 will describe the joint model and Sec. 2.5 the algorithm for optimisation of the joint probability distribution.

2.1   Model of the PET acquisition system

The rate of emission of Gamma rays is a continuous function over the spatial domain of the patient's body and proportional to the local density of radiopharmaceutical. Assuming that photon counts z_d in the detector bins indexed by $d = 1, \ldots, N_d$, are caused only by the radiation emitted by the radio pharmaceutical and approximating the continuous emission rate by a discrete set of point sources $y = y_b, b \in \{1, \ldots, N_b\}$ displaced on a regular grid of voxels, the model of the imaging system is expressed graphically by Fig. 1-A. Letting p_{bd} be

MR image:
$x = \{x_b\};\quad b = 1,\ldots,N_b;\quad x_b \in \mathbb{R}^+$
Pharmaceutical density:
$y = \{y_b\};\quad b = 1,\ldots,N_b;\quad y_b \in \mathbb{R}^+$
Photon counts:
$z = \{z_d\};\quad d = 1,\ldots,N_d;\quad z_d \in \mathbb{N}$
Hidden states:
$k = \{k_b\};\quad b = 1,\ldots,N_b;\quad k_b \in \{e_1,..,e_N\}$
Tissue-specific parameters of MR image:
$\theta_x = \{\mu_{x_n}, \sigma_{x_n}\};\ n = 1,\ldots,N;\ \mu_{x_n},\sigma_{x_n} \in \mathbb{R}^+$
Parameters of the pharmaceutical uptake model:
$\theta_y = \{\mu_{y_n}, \sigma_{y_n}\};\ n = 1,\ldots,N;\ \mu_{y_n},\sigma_{y_n} \in \mathbb{R}^+$
Coeff. of the polynomials for bias field correction:
$c = \{c_j\};\quad j = 1,\ldots,J;\quad c_j \in \mathbb{R}$
Parameters of the spatial dependence of k:
$\theta_k = G_i;\quad i = 1,\ldots,N^2;\quad g_i \in \mathbb{R}$
Population-based multinomial prior distrib. of k:
$\pi = \pi_{bn};\ n = 1,\ldots,N;\ \pi_{bn} \in (0,1);\ \sum_n \pi_{bn} = 1$
N_b: num. of voxels; N_d: num. of detector bins;
N: num. of classes; J: num. of basis functions;
e_n: unit vector of length N, n-th element is 1.

Fig. 1. Hybrid Bayesian Network model of the pharmaceutical uptake for the MR-PET imaging system. Observed and assumed quantities are shaded. The MR image intensity x_b and the pharmaceutical density y_b are assumed to be independent conditionally to the hidden tissue state k_b. Hidden states are drawn from a MRF with first order neighbourhood structure and have a spatially dependent multinomial prior probability distribution π_b, obtained from population studies. A: PET acquisition system; B: MR imaging system. The optimisation algorithm iteratively updates variables $\mathcal{S}_1$ and $\mathcal{S}_2$.

the probability that a photon emitted in b is detected in d, by the sum and thinning properties of the Poisson distribution, denoted by $\mathcal{P}$, and observing (from the d-separation property of the graph) that $z_{d'} \perp z_d|y, \forall\, d' \neq d$, the probability to observe photon counts z when activity is y, is expressed by (see [1]):

$$p(z|y) = \prod_{d=1}^{N_d} \mathcal{P}\left(\sum_{b=1}^{N_b} p_{bd}y_b, z_d\right) \tag{1}$$

The system matrix p_{bd} encompasses the characteristics of the imaging system (location, size, sensitivity, spatial resolution of the Gamma detectors, eventual collimators) and the attenuation of Gamma radiation through the patient, which can be approximated with a simple generic phantom of the brain or estimated from the MR image for the patient under examination [5]. Patient specific estimation of the attenuation coefficients from the MR image, however, is out of the scope of this paper.

2.2   Steady-state model of the pharmaceutical uptake

It is assumed that there exist a finite number of tissue types and that, in a given type of tissue, the pharmaceutical uptake is to some extent predictable. The expectation and the extent of variation of the uptake within each type of tissue are captured by a parametric finite mixture model.

The hidden tissue states $k = \{k_1,\ldots,k_{N_b}\}$ are modelled as the realisation of a random process with parametric probability distribution $p(k|\theta_k)$, where

$k_b = e_n$ for some value $n, 1 \leq n \leq N$, with N being the number of tissue types and e_n a unit vector of length N with n-th component equal to 1. The density of radio-pharmaceutical y_b in a voxel b belonging to class n is assumed to be normally distributed around a certain mean μ_{y_n}, with variance $\sigma^2_{y_n}$, grouped in $\theta_{y_n} = \{\mu_{y_n}, \sigma^2_{y_n}\}$:

$$p(y_b|k_b = e_n) = \mathcal{G}\left(y_b; \mu_{y_n}, \sigma_{y_n}\right) \tag{2}$$

$\theta_y = \{\theta_{y_1}, \ldots, \theta_{y_N}\}$ are the parameters of the pharmaceutical uptake model, describing the expectation and the extent of variation of the uptake in each type of tissue.

2.3   Model of the MR acquisition system

The MR imaging system is described by a parametric voxel-based finite mixture model commonly employed for classification of tissue types by means of MR images [6]. It relates the observed image intensities to the underlying finite hidden tissue states.

The log intensity x_b of a voxel b that belongs to class n is assumed to be normally distributed around a certain mean μ_{x_n}, with variance $\sigma^2_{x_n}$. The smoothly varying bias due to inhomogeneity of the magnetic field is modelled by a linear combination of J polynomial basis functions $\phi_j(X)$, where X denotes the 3D coordinates. Intensity is *log* transformed, as suggested in [6], in order to treat the bias field, which multiplicates the intensity, as an additive term:

$$p(x_b|k_b = e_n) = \mathcal{G}\left(x_b - \sum_{j=1}^{J} c_j \phi_j(X_b); \mu_{x_n}, \sigma_{x_n}\right) \tag{3}$$

where X_b are the coordinates of the voxels b. The model of the MR imaging system (3) is represented by the Bayesian Network in Fig. 1-B. The set of parameters are the bias field coefficients $c = \{c_1, \ldots, c_J\}$ and the mean and spread of the (log) intensity for each class, grouped in $\theta_x = \{\mu_{x_n}, \sigma^2_{x_n}\}, n = 1, \ldots, N$.

2.4   Joint Model

The hidden tissue states are considered the unique underlying cause of mutual dependence of the two images (as expressed by the arrows of the graph in Fig. 1). The hidden state in each voxel is related to the intensity of the MR image and to the pharmaceutical density (which is also a hidden variable) by the parametric models of Sec. 2.2 and Sec. 2.3. Such models are local, expressing the probability of the hidden state in voxel b only as a function of the MR intensity and pharmaceutical uptake (which is also a hidden variable) in the voxel. The model is made more robust by adding contextual information in the form of spatial dependence of the hidden labels. Not only k_b depends on x_b and y_b, but on $k_{b'}, b' \neq b$ elsewhere. For computational convenience, the spatial dependence of k is modelled with a joint distribution of the variables k_b which factorises on the first order neighbourhood structure of the image lattice (Markov Random Field - MRF): k_b is assumed to be conditionally independent from all other states if

its 6 nearest neighbours are known. The dependence of state k_b upon its neighbours $\mathcal{N}_b$ is expressed by the following Potts model $p(k_b|\mathcal{N}_b, \theta_k)$, parametrised by θ_k, which corresponds (Markov-Gibbs equivalence) to the joint density of the hidden states $p(k|\theta_k)$:

$$p(k_b|\mathcal{N}_b, \theta_k) = \frac{e^{-k_b^T\ G_{\theta_k}\ g_b}}{\sum_n^N e^{-k_b^T\ G_{\theta_k}\ g_b}} \qquad p(k|\theta_k) = \frac{e^{-\sum_b^{N_b} -k_b^T\ G_{\theta_k}\ g_b}}{Z} \qquad (4)$$

where g_b counts the labels of each class between the neighbours, G_{θ_k} is a $[N \times N]$ matrix of parameters of the Potts model that expresses the affinity of all pairs of states and Z is the partition function, not involved in the maximisation of the joint probability.

Finally, the probability of voxel b being in state n depends a priori on the location of b. Such prior distribution is expressed by a multinomial probability distribution $\pi_{bn}, \sum_n^N \pi_{bn} = 1$. The use of the spatially varying multinomial prior probability of the hidden states is key to automating the algorithm, as otherwise the parameters of the mixture models converge to different tissue classes non predictably.

2.5   Greedy optimisation algorithm

The joint probability distribution is optimised with the Iterated Conditional Modes (ICM) optimisation algorithm, consisting of iterating the optimisation of two subsets of the unknowns: $\mathcal{S}_1$ and $\mathcal{S}_2$ in Fig. 1. Iteratively, the algorithm computes $\mathcal{S}_1$ that increases the conditional probability distribution $p(\mathcal{S}_1|\mathcal{S}_2)$, given the provisional estimate of $\mathcal{S}_2$, then it computes a new value of $\mathcal{S}_2$ that increases $p(\mathcal{S}_2|\mathcal{S}_1)$. The choice of the two subsets corresponds to the existence of Generalised Expectation Maximisation (GEM) formulations for the optimisation of $p(\mathcal{S}_1|\mathcal{S}_2)$ and $p(\mathcal{S}_2|\mathcal{S}_1)$. The GEM algorithm to update $\mathcal{S}_1$ involves 4 steps (i.i, i.ii, i.iii, i.iv); while the One Step Late EM algorithm to update $\mathcal{S}_2$, involves a single step: ii. Derivation of the update formulas of the subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ is described in [4] and [6], whose notation we maintain here; we report in the following the algorithm, which consists in iterating in order the steps i.i, i.ii, i.iii, i.iv, ii:

i.i Estimate the probability p_{bn} that tissue state in voxel b is n. This step is necessary for the successive steps and arises from the GEM formulation [6].

$$p_{bn}^{(m+1)} \equiv \frac{p(y_b|k_b = e_n; \theta_y^{(m)})p(x_b|k_b = e_n; \theta_x^{(m)})p(k_b = e_n|p_{\mathcal{N}_b}^{(m)}; \theta_k)\pi_{bn}}{\sum_{n=1}^{N} p(y_b|k_b = e_n; \theta_y^{(m)})p(x_b|k_b = e_n; \theta_x^{(m)})p(k_b = e_n|p_{\mathcal{N}_b}^{(m)}; \theta_k)\pi_{bn}} \qquad (5)$$

where the spatial dependence term $p(k_b = e_n | p_{\mathcal{N}_b}^{(m)}; \theta_k)$ is obtained approximately from the previous estimate of p_{bn} by the Mean Field approximation [6]:

$$p(k_b = e_n | p_{\mathcal{N}_b}^{(m)}; \theta_k) = \frac{e^{-U_{\theta_k}(e_n | p_{\mathcal{N}_b}^{(m)})}}{\sum_{n'=1}^{N} e^{-U\theta_k(e_{n'} | p_{\mathcal{N}_b}^{(m)})}} \qquad U_{\theta_k}(e_n | p_{\mathcal{N}_b}^{(m)}) = k_b^T \ G_{\theta_k} \ g_b^{(m)} \tag{6}$$

$g_b^{(m)}$ being the vector of length N with elements $g_{bn}^{(m)} = \sum_{\mathcal{N}_b} p_{bn}^{(m)}$.

i.ii Update the tissue-specific parameters of the MR imaging system:

$$\mu_{x_n}^{(m+1)} = \frac{1}{N_b} \frac{\sum_{b=1}^{N_b} p_{bn}^{(m+1)} x_b}{\pi_{bn}} \qquad \sigma_{x_n}^{2(m+1)} = \frac{1}{N_b} \frac{\sum_{b=1}^{N_b} p_{bn}^{(m+1)} (\mu_{x_n}^{(m+1)} - x_b)^2}{\pi_{bn}} \tag{7}$$

i.iii Update the parameters of the pharmaceutical uptake model:

$$\mu_{y_n}^{(m+1)} = \frac{1}{N_b} \frac{\sum_{b=1}^{N_b} p_{bn}^{(m+1)} y_b}{\pi_{bn}} \qquad \sigma_{y_n}^{2(m+1)} = \frac{1}{N_b} \frac{\sum_{b=1}^{N_b} p_{bn}^{(m+1)} (\mu_{y_n}^{(m+1)} - y_b)^2}{\pi_{bn}} \tag{8}$$

i.iv Update the bias field parameters:

$$\begin{bmatrix} c_1^{(m+1)} \\ c_2^{(m+1)} \\ \dots \end{bmatrix} = \left(A^T \ diag \ w_b^{(m+1)} \ A \right)^{-1} A^T \ diag \ w_b^{(m+1)} \begin{bmatrix} y_1 - \tilde{y}_1^{(m+1)} \\ y_2 - \tilde{y}_2^{(m+1)} \\ \dots \end{bmatrix} \tag{9}$$

where A is the geometrical matrix of the bias field model, each of its columns evaluating the polynomial basis function ϕ_j at voxel coordinates X_b and
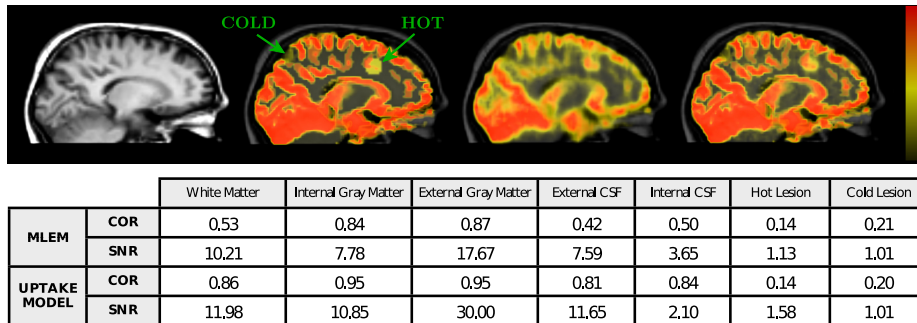
$$w_b^{(m+1)} = \sum_{n=1}^{N} w_{bn}^{(m+1)} \qquad w_{bn}^{(m+1)} = \frac{p_{bn}^{(m+1)}}{\sigma_{x_n}^{2(m+1)}} \qquad \tilde{y}_b^{(m+1)} = \frac{\sum_{n=1}^{N} w_{bn}^{(m+1)} \mu_n^{(m+1)}}{\sum_{n=1}^{N} w_{bn}^{(m+1)}}$$

ii Update the estimate of the pharmaceutical density (One Step Late EM algorithm for the Poisson model - see [1]):

$$y_b^{(m+1)} = y_b^{(m)} \frac{1}{\sum_{d=1}^{N_d} p_{bd} + \sum_{n=1}^{N} p_{bn}^{(m+1)} \frac{y_b^{(m)} - \mu_n^{(m+1)}}{\sigma_{x_n}^2}} \sum_{d=1}^{N_d} \frac{p_{bd} \ z_d}{\sum_{b'=1}^{N_b} p_{b'd} \ y_{b'}^{(m)}} \tag{10}$$

3   Results

Synthetic MR-PET FDG: Uptake of FDG was simulated by assigning typical average values observed in PET FDG scans to the BrainWeb (`http://mouldy.bic.mni.mcgill.ca/brainweb/`) ground truth tissue model: 4 parts uptake in the gray matter, 1 part in white matter, 1 part in CSF. One hot and one cold spherical lesions of 14 mm diameter were simulated by augmenting the activity by 30% in the hot lesion and reducing it by 50% in the cold lesion. 30 PET scans

		White Matter	Internal Gray Matter	External Gray Matter	External CSF	Internal CSF	Hot Lesion	Cold Lesion
MLEM	**COR**	0.53	0.84	0.87	0.42	0.50	0.14	0.21
	SNR	10.21	7.78	17.67	7.59	3.65	1.13	1.01
UPTAKE MODEL	**COR**	0.86	0.95	0.95	0.81	0.84	0.14	0.20
	SNR	11.98	10.85	30.00	11.65	2.10	1.58	1.01

Fig. 2. Synthetic MR-PET FDG imaging data (see-through volume-rendering of 5 mm-thick sagittal slices). From left to right: T1-weighted MR image, activity phantom, MLEM reconstruction, reconstruction with the proposed pharmaceutical uptake model using 6 classes from the MNI-152 brain statistical atlas.

with $150M$ counts were simulated with the Simind simulator and adding Poisson noise. The 30 sinograms were then pre-corrected for attenuation and scatter by energy thresholding and reconstructed with MLEM [1] and with the proposed algorithm. The number of classes N was set to 6: White Matter, Internal Gray Matter, External Gray Matter, External CSF, Internal CSF, everything else. The statistical atlas for the 6 classes was obtained from the MNI-152 atlas by splitting manually internal and external regions for the Gray Matter and CSF classes. The off-diagonal elements of the $N \times N$ parameter matrix θ_k were set to 0.1 and the diagonal elements to 0, penalising equally all transitions except for the transition to the identical class (more sophisticated parameter selection criteria are reported in [6]). The order of the polynomial for bias field correction J was set to 4. For simplicity both MLEM (non-convergent) and the proposed (convergent) algorithm were terminated after 50 iterations. Fig. 2 reports a sagittal slice of the reconstructions obtained for one of the 30 noise instances. The table in Fig. 2 reports the coefficient of recovery (COR) and signal to noise ratio (SNR) of the mean uptake in 8 regions of interest including the two lesions. The results highlight that the pharmaceutical uptake model yields measurements with less noise and less bias when the uptake is consistent with the model, producing remarkably truthful estimates. It's even more remarkable that the lesions are well reconstructed even though they are outliers of the model (there isn't a class that captures the lesions), presenting overall improved SNR. This fact can be explained with the strong inter-correlation of the unknowns in emission imaging due to the line-integral measurements: improving the estimate everywhere outside of the lesion already improves the estimate in the lesion.

Clinical data: The algorithm has been applied to clinical PET FDG using the same 6 classes that were employed in the synthetic PET FDG study. The emission data and the MR images were acquired on separate machines: the pharmaceutical density was estimated initially with MLEM, aligned with the MR image by rigid registration with Normalised Mutual Information cost function
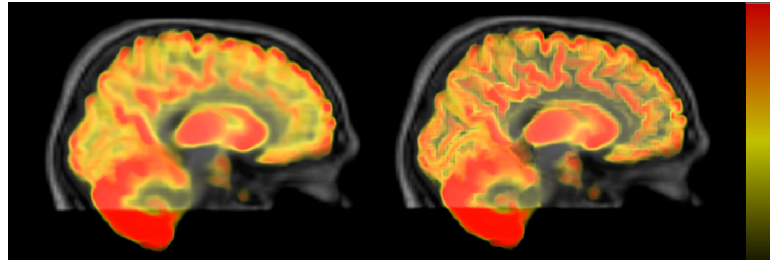
Fig. 3. PET FDG reconstructions obtained with MLEM (left) and with the proposed pharmaceutical uptake model (right) using 6 classes from the MNI-152 brain statistical atlas. See-through volume-rendering of 5 mm-thick sagittal slices.

and reconstructed again with the uptake model. Reconstructions are reported in Fig. 3.

4   Conclusion

In this paper we have extended the pharmaceutical uptake model presented by [4], making it fully automated. Robust estimation is achieved by 1) adopting a population-based statistical atlas to initialise and drive the optimisation of the parameters of the pharmaceutical uptake model and of the MR acquisition system model; 2) adding contextual information in the form of a Markov Random Field over the hidden tissue labels; 3) capturing the MR image bias field. We have evaluated the algorithm on a synthetic study, showing that it improves the quantification of the pharmaceutical uptake when the simulated data reflects the assumptions of the model and, remarkably, that the uptake estimate may improve also in regions that do not obey to the model. Validation with real data remains an open problem as it would require large sets of imaging data, possibly labelled with long term clinical outcome for specific imaging tasks. To this extent, automation of the reconstruction algorithm is crucial.

The source code and data for the synthetic experiment will be made available on the author's website at the time of publication.

References

1. Green, P.G., Bayesian Reconstructions From Emission Tomography Data Using a Modified EM Algorithm. IEEE Trans. on Med. Imag. 9(1), 84–93, 1990.
2. Nuyts, J., Baete, K., Bequ, D., Dupont, P., Comparison between MAP and post-processed ML for image reconstruction in emission tomography when anatomical knowledge is available. IEEE Trans. on Med. Imag. 24(5), 667–675, 2005.
3. Moghbel, M.C., Saboury, B., Basu, S., Metzler, S.D., Torigian, D.A., Langstrom, B., Alavi, A., Amyloid-β imaging with PET in Alzheimers disease: is it feasible with current radiotracers and technologies? European J. of Nucl. Med. and Mol. Imag. 39(2), 202–208, 2012.

4. Pedemonte, S., Bousse, A., Hutton, B.F., Arridge, S.R., Ourselin, S., 4-D generative model for PET/MRI reconstruction. MICCAI, 14(Pt 1), 581–588, September 2011.
5. Hofmann, M., Pichler, B., Scholkopf, B., Beyer, T., Towards quantitative PET/MRI: a review of MR-based attenuation correction techniques. Eur. J. Nucl. Med. Mol. Imag. 36(Suppl 1), 93–104, Mar 2009.
6. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., Automated Model-Based Bias Field Correction of MR Images of the Brain. IEEE Trans. on Med. Imag. 18(10), 885–896, October 1999.