UCSF
UC San Francisco Previously Published Works

Title

K-Bayes Reconstruction for Perfusion MRI II: Modeling and Technical Development

Permalink

Journal

ISSN

Authors

Kornak, John
Young, Karl

Publication Date

2010-08-01

DOI

Copyright Information

Peer reviewed

Journal of Digital Imaging

K-Bayes Reconstruction for Perfusion MRI II: Modeling and Technical Development

John Kornak[1,2] and Karl Young[3,4]

Despite the continued spread of magnetic resonance imaging (MRI) methods in scientific studies and clinical diagnosis, MRI applications are mostly restricted to high-resolution modalities such as structural MRI. While perfusion MRI gives complementary information on blood flow in the brain, its reduced resolution limits its power for detecting specific disease effects on perfusion patterns. This reduced resolution is compounded by artifacts such as partial volume effects, Gibbs ringing, and aliasing, which are caused by necessarily limited k-space sampling and the subsequent use of discrete Fourier transform (DFT) reconstruction. Here, a Bayesian modeling procedure (K-Bayes) is developed for the reconstruction of perfusion MRI. The K-Bayes approach combines a process model for the MRI signal in k-space with a Markov random field prior distribution that incorporates high-resolution segmented structural MRI information. A simulation study, described in Part I (Concepts and Applications), was performed to determine qualitative and quantitative improvements in K-Bayes reconstructed images compared with those obtained via DFT. The improvements were validated using in vivo perfusion MRI data of the human brain. The K-Bayes reconstructed images were demonstrated to provide reduced bias, increased precision, greater effect sizes, and higher resolution than those obtained using DFT.

KEY WORDS: Bayesian reconstruction, K-Bayes, Markov random field, perfusion MRI, structural MRI

INTRODUCTION

T he application of discrete Fourier transform (DFT) reconstruction to limited k-space sampling of perfusion MRI relies on the unwarranted assumption that a small region of k-space can provide a complete representation of the perfusion process. Consequently, DFT reconstructions suffer from poor resolution, partial volume effects, and artifacts such as Gibbs ringing and aliasing. In this report, a new and improved reconstruction proce-dure for perfusion MRI is developed that overcomes (or at least reduces) these limitations of the DFT. Specifically, this new method (K-Bayes) models the relationship between raw k-space data obtained with perfusion MRI and physical (image) space while incorporating high-resolution anatomical information obtained from segmented structural MRI. The K-Bayes approach thereby increases the accuracy, precision, and resolution of perfusion MRI reconstructions.

k-Space and DFT Reconstruction

In all MRI modalities, data are acquired in k-space (also known as frequency or Fourier space). In generating an image of a single slice of the brain, magnetic field gradients are applied in orthogonal directions so that the frequency and phase of rotation of the nuclei in the field of view

are perturbed according to their location in image space. By recording the intensity at particular frequencies and phases, the distribution of the signal over image space can be determined from *k*-space through a 2D Fourier transform relationship. Figure 1 gives an illustration of how MRI signals are represented in *k*-space. At each point in *k*-space, there is a corresponding spatial pattern in image space. The point (0,0) corresponds to the overall mean level of the image to be reconstructed (zero frequency in *x* and *y* directions). The pattern at coordinates (0,4) describes a pattern with spatial frequency of 4 in the *y*-direction (i.e., four cycles over the field of view of the image) but with no variation in the *x*-direction. At (8,0), the pattern has a frequency of 8 in the *x*-direction, but no variation in the *y*-direction. At (8,4), there are eight cycles in the *x*-direction and four in the *y*-direction, leading to diagonal bands. The relative magnitude of the signal at each position in *k*-space represents the magnitude of the signal in image space at that frequency.

The signal is fully represented in *k*-space by a complex signal consisting of separate real and imaginary components. The distribution of the signal between the real and imaginary parts at a particular point in *k*-space represents the phase of the signal; the phase describes the relative positioning of the bands. This is illustrated in Figure 2, which displays signal at a frequency corresponding to *k*-space point (0,4). The imaginary part shown in panel (b) has a 90° phase offset (in the vertical (*y*)-direction) to the real part shown in panel (a).
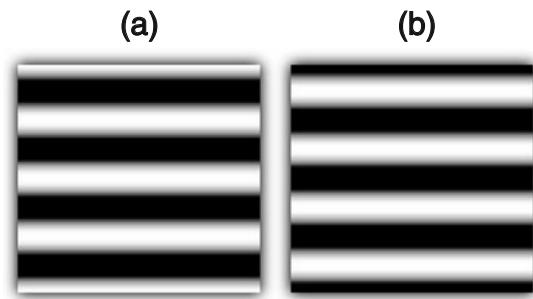
Fig 2. **a** Purely real and **b** purely imaginary (90° phase offset) components of a signal at *k*-space point **(0,4)**. *Light areas correspond to high values.*

While DFT offers a convenient approach to reconstruct images in image space from data in *k*-space, the phenomena we observe in perfusion MRI are considerably less sensitive than those of structural MRI. Hence, only a limited (central) region of *k*-space (i.e., in the range 32×32 to 64×64 compared with from 128×128 to 512×512 for structural MRI) can be sampled at a reasonable signal-to-noise ratio and in reasonable imaging time (i.e., <15 min). For a constant field of view and DFT reconstruction, this implies reduced image resolution for perfusion MRI.

Furthermore, when only a limited central region of *k*-space can be sampled, such as is the case for perfusion MRI, DFT can cause problems because it relies on the *k*-space data fully representing the image to be reconstructed (for optimality). Because there is considerable information about perfusion MRI at frequencies higher than those acquired, using DFT in this setting can cause image artifacts, including Gibbs ringing (observed as bright rings in the image), aliasing (seen as ghosting in the image), and partial volume effects (caused by the averaging of signal over large pixels composed of multiple tissue types, shrinking the signal towards a global mean measurement). Furthermore, DFT reconstruction contains no intrinsic noise reduction mechanism and provides no means for using external anatomical information from structural MRI. Anatomical information can improve perfusion MRI reconstruction by providing appropriate constraints on the perfusion process. The limited quality of DFT reconstruction for perfusion MRI can be seen in the example of Figure 3.

K-Bayes avoids these limitations of DFT by modeling the reconstruction process of the physiologic data at the higher resolution of structural MRI. Two factors lead to increased resolution and improved image quality of K-Bayes reconstruction:
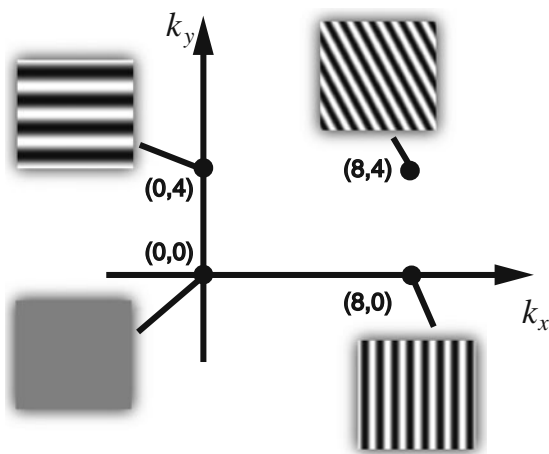

Fig 1. Illustration of *k*-space representation. Each *box* displays an image corresponding to a signal at the associated point in *k*-space.
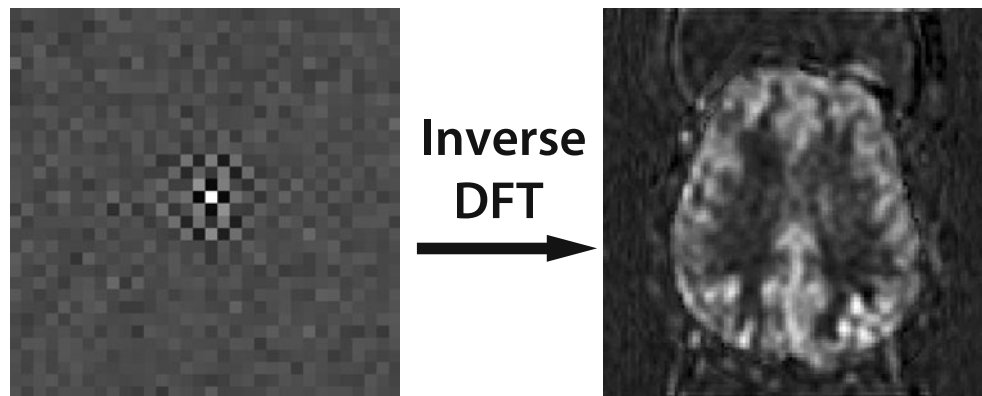
Fig 3. Schematic showing the limitations of DFT reconstruction for perfusion MRI. DFT-based reconstruction contains Gibbs ringing, aliasing, partial voluming, and high noise level.

(1) K-Bayes avoids the assumption that the limited k-space data fully represent the physiologic MR process and (2) K-Bayes incorporates a mechanism that allows integration of anatomical prior information.

METHODS

K-Bayes aims to overcome the inherent limitations of DFT by defining a k-space modeling procedure that incorporates high-resolution anatomical prior information to provide constraints for the reconstructed physiological MR process.

A schematic comparing the DFT and K-Bayes procedures is displayed in Figure 4. K-Bayes utilizes the general Bayesian modeling formulation composed of a prior and a likelihood distribution. (See Gelman[1]
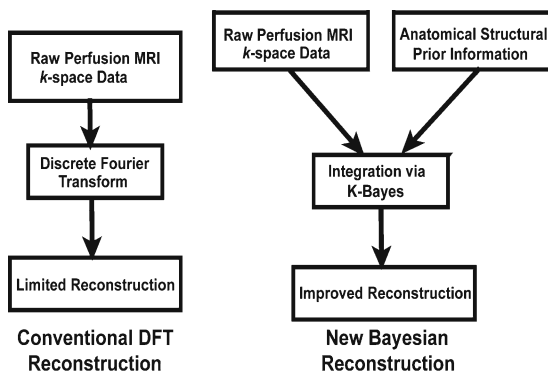

Fig 4. Differences between standard DFT and the proposed K-Bayes reconstruction procedure. K-Bayes improves reconstruction quality by (1) utilizing high-resolution anatomical prior information from structural MRI and (2) modeling the relationship between the limited perfusion MRI sample of k-space and the high-resolution image to be reconstructed.

for a description of Bayesian modeling and Winkler[2] for a description of Bayesian image analysis with Markov random field prior distributions.) Anatomical information from structural MRI drives the prior model, while the likelihood describes the information in the k-space perfusion MRI data. The information from the prior and likelihood are combined though Bayes' Theorem to generate a posterior distribution of the perfusion MRI process at high spatial resolution.

The k-space signal model forms the basis of the likelihood component of K-Bayes and relates the k-space data generation process to the parameters to be estimated. The parameters to be estimated are the voxel values of the true spatial perfusion map, i.e., what would be obtained with perfect measurement.

The other component of the Bayesian model, the prior distribution, describes prior knowledge about the true perfusion map. For perfusion MRI, the prior knowledge utilized by K-Bayes describes the behavior of perfusion with respect to different tissue types and the boundaries between them (as defined by tissue segmented structural MRI). The product of the prior and the likelihood is proportional to the posterior distribution, which needs to be summarized in order to provide an estimate for the high-resolution reconstructed map of perfusion MRI. For this purpose, the reconstruction is chosen as the image corresponding to the maximum of the posterior distribution. This reconstruction is known as the maximum a posteriori (MAP) estimate.[2] Other possible choices for reconstruction estimates such as the mean of the posterior distribution are typically more expensive computationally, requiring stochastic estimation methods. Because the MAP estimate cannot be calculated analytically, we formulate an

iterative algorithm based on the Expectation-Maximization (EM) approach.[3–5] More precisely, we use a Generalized EM approach,[3] though in common with much of the engineering literature, we continue to refer to it as an EM algorithm.

k-Space Signal Modeling and the Likelihood

A *signal model* is developed for K-Bayes that relates k-space perfusion data (the difference signal between labeled and unlabeled signals) to a voxel map in image space. The image space voxel map is defined at higher resolution than DFT permits (other than assuming, generally with no justification, that a particular form of interpolation is appropriate). The k-space model increases resolution by relating the observed signal from the acquired set of k-space points to the set of voxel intensities at the higher resolution of structural MRI. Development of the model is initially provided in 2D rather than 3D to aid the visualization of key concepts. The extension to 3D is straightforward, though the exact 3D model depends on (a) whether a set of 2D 'slices' or a single 3D volume is acquired and (b) the resolution between-slice relative to within-slice.

Formally, the k-space signal model is incorporated within the likelihood model part of the K-Bayes procedure. The raw k-space data are modeled as consisting of the true signal plus complex Gaussian noise (the errors are modeled as having independent Gaussian distributions with zero mean and constant variance in each of the real and imaginary parts). This independent Gaussian assumption implies a uniformly random phase for the noise. The k-space signal process is modeled as a continuous Fourier transformation of the true perfusion map being imaged. The true perfusion map is modeled as a piece-wise constant (uniform within each voxel) process at the (high) resolution of the structural MRI.

The 2D k-space perfusion signal model for K-Bayes is

$$d[k_x, k_y] = s[k_x, k_y] + \varepsilon[k_x, k_y] \qquad (1)$$

where $d[k_x, k_y]$ is the complex k-space data at k-space position (k_x, k_y), $k_x = -\frac{K_x}{2}, \ldots, \frac{K_x}{2} - 1$, $k_y = -\frac{K_y}{2}, \ldots, \frac{K_y}{2} - 1$. The error/noise component is modeled as isotropic complex Gaussian noise: $\varepsilon[k_x, k_y] \sim CN(0, \sigma^2 I_2)$, where $CN(\mu, V)$ is used to denote a bivariate Gaussian distribution with the first component representing the real and the second the imaginary part of a complex random variable with

vector mean μ and covariance matrix V. The independence imposed by the covariance matrix $\sigma^2 I_2$ provides for an isotropic noise distribution in the complex plane. The term $s[k_x, k_y]$ denotes the 'true' (noise-free) k-space complex signal, modeled as

$$s[k_x, k_y] = \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} \int_{y[p]-1/2}^{y[p]+1/2} \int_{x[p]-1/2}^{x[p]+1/2}$$
$$\times A[x, y] \exp\{-2\pi i[k_x x/P + k_y y/Q]\} dx dy \qquad (2)$$

where $p = 0, \ldots, P-1$ and $q = 0, \ldots, Q-1$ are the co-ordinate values (p, q) of the high-resolution structural MRI voxel locations. $A[x, y]$ is the signal amplitude at point (x, y) in image space, i.e., unknown perfusion intensity map.

Making the assumption that the signal is constant within each voxel in image space (at the higher resolution of the structural MRI), the analytic solution to the integral of Eq. (2) is:

$$s[k_x, k_y] = \text{sinc}(\pi k_x/P) \, \text{sinc}(\pi k_y/Q) \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} A[p, q]$$
$$\times \exp\{-2\pi i[k_x p/P + k_y q/Q]\} \qquad (3)$$

where $\text{sinc}(x) = \sin(x)/x$. The likelihood is proportional to the distribution of the data d (the set of $d[k_x, k_y]$ at all k-space points) given the true signal s ($s[k_x, k_y]$ at all k-space points). The complex Gaussian noise process combined with the true signal model leads to the following likelihood:

$$\pi(d|A) \propto \exp\left\{-\frac{1}{2\sigma^2} \sum_{(k_x, k_y)} |d[k_x, k_y] - s[k_x, k_y]|^2\right\}$$

$$= \exp\left\{-\frac{1}{2\sigma^2} \sum_{(k_x, k_y)} |d[k_x, k_y] - \text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q) \times \right.$$
$$\left. \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} A[p, q] \exp\{-2\pi i[k_x p/P + k_y q/Q]\}|^2\right\} \qquad (4)$$

where $\pi(\bullet)$ is used throughout this document to denote a generic probability density function. Notice that the likelihood is conditional on the true image $A = \{A[p, q] : p = 0, \ldots, P-1, q = 0, \ldots, Q-1\}$ through the $s[k_x, k_y]$ terms. That is, the likelihood describes the data generation process given the true underlying physiologic process via the true k-space signal process.

Incorporating Anatomical Information from Structural MRI (The Prior Distribution)

To optimally incorporate the anatomical prior information, a prior distribution is adopted for K-Bayes that is an adaptation of Markov random field (MRF) models.[2,6–9] MRF models are a commonly used mechanism to impose smoothness constraints on an image or map. The smoothing is probabilistic since it is incorporated via a prior probability distribution. The approach is to model neighboring voxels with higher probability of having more similar values than voxels farther apart. When used as priors for image reconstruction, MRFs yield smoother and less noisy results than reconstructions based on the data alone. The Bayesian paradigm essentially multiplies the prior and the likelihood, and the resulting posterior distribution is a compromise of the two. Therefore, a reconstruction based on Bayesian models with MRF priors is capable of compromising between real image characteristics in the data and the desired property of having generally smooth neighborhoods in the image.

The prior model for K-Bayes extends the standard MRF prior model by incorporating tissue segmented anatomical information from structural MRI. Unlike a standard MRF model, the K-Bayes prior does not define the level of probabilistic smoothing between neighboring voxels to be constant across the reconstructed perfusion map. The K-Bayes prior is designed so that the level of probabilistic smoothing of the perfusion map is only constant within homogeneous regions of tissue. Sharper changes in levels have higher probability of occurring at identified boundaries of gray and white matter and at edges of the brain or ventricles. Regions outside the brain or areas of CSF should not have a detectable perfusion signal, and the prior defines these areas to have zero signal. The K-Bayes prior is mathematically represented as an adaptation of standard MRF models and is defined at the spatial resolution of structural MRI as follows:

$$\pi(A) \propto \exp\left\{ -\frac{1}{2}\left[\sum_{\langle (p,q),(p',q')\rangle} \left\{ \left(\frac{1}{\tau_B^2} I_B[(p,q),(p',q')] + \frac{1}{\tau_G^2} I_G[(p,q),(p',q')] + \frac{1}{\tau_W^2} I_W[(p,q),(p',q')] \right) (A[p,q] - A[p',q'])^2 \right\} \right] \right\} \tag{5}$$

with the added constraint that A is assumed zero everywhere both outside the brain and within regions consisting of CSF. The sum in Eq. (5) over $\langle (p,q), (p',q')\rangle$ is over all pairs of 'neighboring' voxels (p,q), (p',q') and is only over first-order neighbors, i.e., where only horizontally and vertically adjacent pairs are considered to be neighbors, though this can be generalized to higher orders (and in the language of the MRF literature, "arbitrary cliques") in a straightforward manner. The functions $I_G[(p,q),(p',q')]$, $I_W[(p,q),(p',q')]$ and $I_B[(p,q),(p',q')]$ are indicator functions of matching pairs of gray matter, white matter and brain tissue (gray or white) voxels, respectively. For example, $I_G[(p,q), (p',q')] = 1$ if both (p, q) and $(p', q')(p',q')$ are gray matter voxels and 0 otherwise. The parameters τ_B^2, τ_G^2, and τ_W^2 control the level of probabilistic smoothing between neighbors (probabilistic because it is determined by a probability model describing the similarity of neighboring values); τ_B^2 defines the a priori expected level of smoothness of non-matching tissue type neighbors and τ_G^2, τ_W^2 describe the extra components of smoothness for gray matter and white matter neighboring pairs, respectively (lower values provide increased smoothness). This prior distribution model can be considered as an extension of the intrinsic pairwise Gaussian difference prior,[8,9] extended such that the parameter(s) for the pairwise smoothness level depends on the tissue types of neighboring voxels.

This prior model allows for different levels of Markov smoothing within different types of tissue and also across tissue boundaries. A decisive benefit of this prior model is its great flexibility to accommodate the differing prior expectations of the constraints in different tissues. For example, perfusion has been observed to bleed across tissue boundaries of gray and white matter. The parameter τ_B^2 controls the smoothness of transitions across gray/white boundaries such that the cross-boundary

bleeding can be appropriately incorporated. This model can be interpreted as an 'adaptive MRF' prior where the relative levels of desired smoothness within each tissue type and across boundaries are chosen to reflect the way that they constrain the underlying perfusion process being imaged.

The values used for σ^2, τ_B^2, τ_G^2, and τ_W^2 in the 2D simulation of Paper I were $\sigma^2 = 1.0$, $\tau_B^2 = 40.0$, $\tau_G^2 = 1.0$, and $\tau_W^2 = 5.0$. These parameter values were obtained via manual calibration, though it was found that the resulting reconstructions were reasonably robust to their specification. The increased strength of a priori smoothness in gray over white matter was necessary because there are many gray matter voxels that are not surrounded by other gray matter voxels (due to the thin "ribbon-like" nature of the cortex). Therefore, to ensure similar smoothness levels in gray and white tissue, the gray matter smoothness parameter needs to be smaller than that for white.

Posterior Optimization

In order to provide a function from which a reconstruction of the perfusion map can be obtained, the likelihood (4) and prior (5) components are combined via Bayes' Theorem to produce the posterior distribution:

$$\pi(A|d) \propto \pi(d|A)\pi(A) \qquad (6)$$

To obtain an estimate (reconstruction) of the true physiologic map, the posterior distribution is maximized with respect to the set of perfusion MRI voxel intensities, A. This maximum a posteriori (MAP) estimator is only one of many possible estimators of the true physiologic spatial map that can be obtained from the posterior distribution. However, it is a common choice for Bayesian image analysis because (a) it is easy to understand and (b) there is known methodology for obtaining it.[2] An implementation of the Expectation-Maximization (EM) algorithm[3] is here formulated to compute the K-Bayes MAP estimate. The implementation is similar in approach to those in Green[4] and Miller et al.,[10] where a set of unobservable (latent) variables is defined that can

be considered as missing data. For the K-Bayes EM reconstruction, the latent variables are defined to be the component of the data at k-space location (k_x, k_y), i.e., $d[k_x, k_y]$, relative to the physiologic signal at voxel (p, q). These terms are:

$$z[k_x, k_y, p, q] = \text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q)A[p,q]$$
$$\times \exp\{-2\pi i[k_x p/P + k_y q/Q]\}$$
$$+ \varepsilon'[k_x, k_y, p, q]$$

where $\varepsilon' \sim CN(0, \sigma^2/nI_2)$ and n is the total number of voxels. The implicit assumption here is that the error contribution from each voxel (p, q) to $d[k_x, k_y]$ (i.e., $\varepsilon'[k_x, k_y, p, q]$) is independently and identically distributed (i.i.d.) for all (p, q) and (k_x, k_y).

The individual z-variables are not defined in a rigorous sense because of the uncertainty principle, i.e., it is impossible to specify a point precisely in both spatial frequency and image space. However, because only linear combinations of the z-variables are utilized (which are well defined) within the maximization step of the EM algorithm, the methodology remains valid at all times.

The EM procedure for K-Bayes takes the following steps:

1. A set of start values is determined for the perfusion MRI map based on zero-filled DFT (zDFT) reconstruction. Zero-filling involves extrapolating k-space with zeros in each dimension. This corresponds to sinc interpolation in standard space. The zero-filling is performed to a level such that the resolution of the zDFT'd data matches that of the structural MRI dataset.

Steps 2 through 4 are then iteratively performed.

2. The expectation of each z-variable is evaluated given the current values of the perfusion MRI amplitudes A (the expectation or E-step).

3. The conditional posterior distribution for each $A[p,q]$ is maximized given the current estimates for the set z-variables (the maximization or M-step) to give new values for A.

4. The algorithm is tested for convergence and stopped if the total change in the perfusion reconstruction, A, between iterations is less than some pre-specified tolerance.

5. Once convergence has been achieved, the current values for the image, A, form the reconstructed image.

Derivation of K-Bayes E and M steps

For the purpose of clarity, the EM algorithm derivation is presented only for 2D K-Bayes. The adaptation to 3D is relatively straightforward.

Proposition: *The j-th step of the K-Bayes EM algorithm (based on Eqs. (4), (5), and (6)) has E-step:*

$$
z[k_x, k_y, p, q]^j = \frac{1}{n}\left[d[k_x, k_y] - \mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q) \sum_{(p',q')\neq(p,q)} A[p',q']^{j-1} \exp\left\{-2\pi i\left[k_x p'/P + k_y q'/Q\right]\right\} \right]
$$
$$
+ \frac{n-1}{n} z[k_x, k_y, p, q]^{j-1}
$$

where $n=PQ$ is the number of voxels in the image, and M-step:

$$
A[p,q]^j = \frac{\frac{1}{\sigma^2}\sum_{k_x,k_y}\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\left(\begin{array}{l}\mathrm{Re}\left[z[k_x,k_y,p,q]^j\right]\cos\left(2\pi\left[k_x p/P + k_y q/Q\right]\right)-\\ \mathrm{Im}\left[z[k_x,k_y,p,q]^j\right]\sin\left(2\pi\left[k_x p/P + k_y q/Q\right]\right)\end{array}\right) + \left\langle \sum_{(p',q')\in\delta(p,q)}\right\rangle \left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')] + \frac{1}{\tau_G^2}I_G[(p,q),(p',q')] + \frac{1}{\tau_W^2}I_W[(p,q),(p',q')]\right)A[p',q']^{j*}}{\frac{1}{\sigma^2}\sum_{k_x,k_y}\left(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\right)^2 + \left\langle \sum_{(p',q')\in\delta(p,q)}\right\rangle \left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')] + \frac{1}{\tau_G^2}I_G[(p,q),(p',q')] + \frac{1}{\tau_W^2}I_W[(p,q),(p',q')]\right)}
$$

where j^* can be taken as $j-1$, or for potentially improved performance, the latest available update (either $j-1$ or j).

Proof E-step

The E-step update for each missing data point $z[k_x, k_y, p, q]$ is determined by its expectation given the observed data, d, and the perfusion map, i.e., $E[z[k_x, k_y, p, q]|d, A] = E[z[k_x, k_y, p, q]| d[k_x, k_y], A]$.

Bayes' Theorem gives

$$
\pi\big(z[k_x, k_y, p, q]\,\big|\,d[k_x, k_y], A\big)
$$
$$
= \frac{\pi\big(d[k_x, k_y]\,\big|\,z[k_x, k_y, p, q], A\big)\pi\big(z[k_x, k_y, p, q]\,\big|\,A\big)}{\pi\big(d[k_x, k_y]\,\big|\,A\big)}.
$$

Furthermore,

$$
d[k_x, k_y]\,\big|\,z[k_x, k_y, p, q], A \sim
$$
$$
CN\left(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q) \sum_{(p',q')\neq(p,q)} A[p',q'] \exp\left\{-2\pi i\left[k_x p'/P + k_y q'/Q\right]\right\} + z[k_x, k_y, p, q], \frac{\sigma^2(n-1)}{n}I_2 \right)
$$

Therefore,

$$\pi\big(z[k_x,k_y,p,q]\big|d[k_x,k_y],A\big) = \frac{n}{\sigma\sqrt{2\pi(n-1)}}$$

$$\exp\left\{-\frac{n}{2\sigma^2(n-1)}\left|\begin{array}{c}z[k_x,k_y,p,q]-d[k_x,k_y]+\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\big(\pi k_y/Q\big)\\ \sum_{(p',q')\neq(p,q)}A[p',q']\exp\{-2\pi i[k_xp'/P+k_yq'/Q]\}\end{array}\right|^2\right\}$$

$$\exp\left\{-\frac{n}{2\sigma^2}\big|z[k_x,k_y,p,q]-\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\big(\pi k_y/Q\big)A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}\big|^2\right\}$$

$$\Big/\exp\left\{-\frac{1}{2\sigma^2}\left|d[k_x,k_y]-\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\big(\pi k_y/Q\big)\sum_{p=0}^{P-1}\sum_{q=0}^{Q-1}A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}\right|^2\right\}$$

$$\propto\exp\left\{-\left[\begin{array}{c}\left[\frac{n}{2\sigma^2(n-1)}\big|z[k_x,k_y,p,q]\big|^2-\right.\\ \left[\begin{array}{c}z[k_x,k_y,p,q]\left\{\overline{\left(\begin{array}{c}d[k_x,k_y]-\\ \mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\sum_{(p',q')\neq(p,q)}A[p',q']\exp\{-2\pi i[k_xp'/P+k_yq'/Q]\}\end{array}\right)}\right\}\\ +\overline{z[k_x,k_y,p,q]}\left\{\left(\begin{array}{c}d[k_x,k_y]-\\ \mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\sum_{(p',q')\neq(p,q)}A[p',q']\exp\{-2\pi i[k_xp'/P+k_yq'/Q]\}\end{array}\right)\right\}\end{array}\right]\right.\\ +\frac{n}{2\sigma^2}\big|z[k_x,k_y,p,q]\big|^2-\\ \left[\begin{array}{c}z[k_x,k_y,p,q]\left(\overline{\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}}\right)\\ +\overline{z[k_x,k_y,p,q]}\big(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}\big)\end{array}\right]\end{array}\right]\right\}$$

Completing the square, this leads to

$$\pi\big(z[k_x,k_y,p,q]\big|d[k_x,k_y],A\big) \propto$$

$$\exp\left\{-\frac{n^2}{2\sigma^2(n-1)}\left|\begin{array}{c}z[k_x,k_y,p,q]-\\ \left(\begin{array}{c}\frac{1}{n}\left[d[k_x,k_y]-\mathrm{sin\,c}(\pi k_x/P)\mathrm{sin\,c}\big(\pi k_y/Q\big)\sum_{(p',q')\neq(p,q)}A[p',q']\exp\{-2\pi i[k_x,p'/P+k_yq'/Q]\}\right]\\ +\frac{n-1}{n}\mathrm{sin\,c}(\pi k_x/P)\mathrm{sin\,c}(\pi k_y/Q)A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}\end{array}\right)\end{array}\right|^2\right\}.$$

Therefore, $\pi\big(z[k_x,k_y,p,q]\big|d[k_x,k_y],A\big)$ is complex Gaussian with mean

$$\frac{1}{n}\left[d[k_x,k_y]-\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\big(\pi k_y/Q\big)\sum_{(p',q')\neq(p,q)}A[p',q']\exp\{-2\pi i[k_xp'/P+k_yq'/Q]\}\right]$$

$$+\frac{n-1}{n}\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\big(\pi k_y/Q\big)A[p,q]\exp\{-2\pi i[k_xp/P+k_yq/Q]\}$$

and the j-th E-step update for $z[k_x, k_y, p, q]$ is,

$$
z[k_x, k_y, p, q]^j = \frac{1}{n} \left[d[k_x, k_y] - \text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q) \sum_{(p',q') \neq (p,q)} A[p', q']^{j-1} \exp\{-2\pi i [k_x p'/P + k_y q'/Q]\} \right]
$$
$$
+ \frac{n-1}{n} z[k_x, k_y, p, q]^{j-1}
$$

M-step

To generate the M-step, the posterior distribution, or equivalently the *log-posterior* distribution $\log(\pi(A[p, q]|d, z))$, is maximized with respect to each $A[p,q]$. Note that the posterior distribution is conditional on knowing the full set of latent (or missing) variables, i.e., z.

Differentiating the log-posterior distribution with respect to $A[p,q]$ and setting equal to zero gives

$$
\frac{1}{\sigma^2} \left\{
\begin{array}{l}
\sum_{k_x, k_y} \left(\text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q)A[p,q]Re\left[\exp\{-2\pi i[k_x p/P + k_y q/Q]\}\right] - Re[z[k_x, k_y, p, q]] \right) \\
\text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q)\text{Re}\left[\exp\{-2\pi i[k_x p/P + k_y q/Q]\}\right] + \\
\sum_{k_x, k_y} \left(\text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q)A[p,q]\text{Im}\left[\exp\{-2\pi i[k_x p/P + k_y q/Q]\}\right] - \text{Im}[z[k_x, k_y, p, q]] \right) \\
\text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q)\text{Im}\left[\exp\{-2\pi i[k_x p/P + k_y q/Q]\}\right]
\end{array}
\right\} +
$$
$$
\sum_{\langle (p',q')) \in \delta(p,q) \rangle} \left\{ \left(\frac{1}{\tau_B^2} I_B[(p,q), (p', q')] + \frac{1}{\tau_G^2} I_G[(p,q), (p', q')] + \frac{1}{\tau_W^2} I_W[(p,q), (p', q')] \right) (A[p,q] - A[p', q']) \right\} = 0
$$

$$(7)$$

where $\delta(p,q)$ is the set of neighboring voxels to location (p,q). Rearranging Eq. (7) in terms of $A[p,q]$, leads to

$$
A[p,q] = \frac{\frac{1}{\sigma^2} \sum_{k_x, k_y} \text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q) \left(\begin{array}{l} \text{Re}[z[k_x, k_y, p, q]] \cos(2\pi[k_x p/P + k_y q/Q]) - \\ \text{Im}[z[k_x, k_y, p, q]] \sin(2\pi[k_x p/P + k_y q/Q]) \end{array} \right) + \sum_{\langle (p',q') \in \delta(p,q) \rangle} \left(\frac{1}{\tau_B^2} I_B[(p,q), (p', q')] + \frac{1}{\tau_G^2} I_G[(p,q), (p', q')] + \frac{1}{\tau_W^2} I_W[(p,q), (p', q')] \right) A[p', q']}{\frac{1}{\sigma^2} \sum_{k_x, k_y} \left(\text{sinc}(\pi k_x/P)\text{sinc}(\pi k_y/Q) \right)^2 + \sum_{\langle (p',q') \in \delta(p,q) \rangle} \left(\frac{1}{\tau_B^2} I_B[(p,q), (p', q')] + \frac{1}{\tau_G^2} I_G[(p,q), (p', q')] + \frac{1}{\tau_W^2} I_W[(p,q), (p', q')] \right)}
$$

$$(8)$$

Differentiating Eq. (7) again we get that the second differential of the log-posterior is

$$
\frac{1}{\sigma^2}
\left\{
\begin{array}{l}
\sum_{k_x,k_y}\left(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\left(\pi k_y/Q\right)\mathrm{Re}\left[\exp\left\{-2\pi i\left[k_x p/P+k_y q/Q\right]\right\}\right]\right)^2 \\
\sum_{k_x,k_y}\left(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}\left(\pi k_y/Q\right)\mathrm{Im}\left[\exp\left\{-2\pi i\left[k_x p/P+k_y q/Q\right]\right\}\right]\right)^2
\end{array}
\right\} +
$$
$$
\sum_{\langle(p',q')\in\delta(p,q)\rangle}\left\{\left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')]+\frac{1}{\tau_G^2}I_G[(p,q),(p',q')]+\frac{1}{\tau_W^2}I_W[(p,q),(p',q')]\right)\right\} > 0 \tag{9}
$$

Because the second differential is positive, this confirms that the solution given by Eq. (8) is a local maximum and therefore the *j*-th M-step update for $A[p,q]$ is:

$$
A[p,q]^j = \frac{\frac{1}{\sigma^2}\sum_{k_x,k_y}\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\left(\begin{array}{l}\mathrm{Re}\left[z[k_x,k_y,p,q]^j\right]\cos\left(2\pi\left[k_x p/P+k_y q/Q\right]\right)-\\ \mathrm{Im}\left[z[k_x,k_y,p,q]^j\right]\sin\left(2\pi\left[k_x p/P+k_y q/Q\right]\right)\end{array}\right)+ \sum_{\langle(p',q')\in\delta(p,q)\rangle}\left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')]+\frac{1}{\tau_G^2}I_G[(p,q),(p',q')]+\frac{1}{\tau_W^2}I_W[(p,q),(p',q')]\right)A[p',q']^{j*}}{\frac{1}{\sigma^2}\sum_{k_x,k_y}\left(\mathrm{sinc}(\pi k_x/P)\mathrm{sinc}(\pi k_y/Q)\right)^2+\sum_{\langle(p',q')\in\delta(p,q)\rangle}\left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')]+\frac{1}{\tau_G^2}I_G[(p,q),(p',q')]+\frac{1}{\tau_W^2}I_W[(p,q),(p',q')]\right)}
$$

$$\tag{10}$$

$\square$

2D versus 3D Reconstruction

One can use K-Bayes as either a 2D or 3D reconstruction approach. The choice depends on a number of factors. MRI data are commonly recorded as a set of 2D slices rather than encoded in 3D *k*-space. Moreover, the between-slice thickness of perfusion MRI is typically larger than the within-slice resolution. When the difference between within- and between-slice resolution is large, 2D reconstruction may be preferred, particularly if reconstruction speed is important. However, when the resolution is closer to isotropic (equal in all dimensions) or the data are recorded in 3D *k*-space, then a 3D reconstruction procedure is likely to produce improved results over 2D. The improvement is acquired from the utilization of cross-slice structural information.

3D K-Bayes

The extension to 3D requires different formulations depending on whether the 3D data are

acquired as a set of slices (multi-slice) or as a single 3D k-space acquisition. For multi-slice data, the model of Eq. (3) is extended to:

$$s[k_x, k_y, w] = \frac{\sin(\pi k_x/P)}{\pi k_x/P} \frac{\sin(\pi k_y/Q)}{\pi k_y/Q} \sum_{r=wc}^{(w+1)c-1} \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1}$$
$$A[p,q,r] \exp\{-2\pi i[k_x p/P + k_y q/Q]\}.$$
(11)

where $w \in \{0, \dots, W-1\}$ is the slice index of the perfusion MRI (W slices total); $r \in \{0, \dots, R-1\}$ is the slice index of the higher resolution structural MRI (R slices); and c is the number of high-resolution structural MRI slices corres-

ponding to each low-resolution perfusion MRI slice ($c=4$ for the 3D simulations in Paper 1).

When 3D acquisition sequences are used, i.e., when data are acquired in 3D k-space, the model of Eq. (3) extends to:

$$s[k_x, k_y, k_z] = \frac{\sin(\pi k_x/P)}{\pi k_x/P} \frac{\sin(\pi k_y/Q)}{\pi k_y/Q} \frac{\sin(\pi k_z/R)}{\pi k_z/R}$$
$$\sum_{r=0}^{R-1} \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} A[p,q,r]$$
$$\exp\{-2\pi i[k_x p/P + k_y q/Q + k_z r/R]\}.$$
(12)

In both cases the prior distribution is extended to the third dimension as:

$$\pi(A) \propto \exp\left\{-\frac{1}{2}\left[\sum_{\langle(p,q,r),(p',q',r')\rangle}\left\{\left(\frac{1}{\tau_B^2}I_B[(p,q),(p',q')] + \frac{1}{\tau_G^2}I_G[(p,q,r),(p',q',r')] + \frac{1}{\tau_W^2}I_W[(p,q,r),(p',q',r')]\right)(A[p,q,r] - A[p',q',r'])^2\right\}\right]\right\}.$$
(13)

This is essentially the same form as in the 2D reconstruction except that the Markov property is applied equivalently in the third dimension, i.e., the neighborhood is defined for (p,q,r) and not just (p,q). Note that adjustment of the prior parameters is required to accommodate the different structure of the neighborhoods. The values used for σ^2, τ_B^2, τ_G^2, and τ_W^2 in the 3D simulation study of Paper I were changed to $\sigma^2=1.0$, $\tau_B^2 = 100.0$, $\tau_G^2 = 4.0$, and $\tau_W^2 = 15.0$.

Moreover, the model needs to be modified when the structural MRI is non-isotropic, i.e., when the between-slice resolution is different to the within-slice resolution. The values of the smoothing parameters τ_B^2, τ_G^2, and τ_W^2 need to be adjusted for neighboring pairs of voxels in the between-slice direction to ensure isotropic behavior of the prior distribution, i.e., such that the auto-correlation function is the same in all spatial directions. Theory for how to adjust these parameters is not resolved. Therefore, we recommend that, in the case of non-isotropic voxels, the structural MRI be re-sliced to have matching resolution in all dimensions prior to segmentation. Alternatively, if re-slicing is undesirable, 2D reconstruction might be used for each slice.

DISCUSSION

The K-Bayes formulation developed in this report overcomes some of the inherent limitations of conventional DFT reconstruction procedures for perfusion MRI. K-Bayes improves on DFT by combining raw k-space data modeling with high-resolution anatomical information from structural MRI. The K-Bayes method developed here and applied in Paper I demonstrates improved spatial resolution, accuracy, and precision over DFT that should boost diagnostic sensitivity and specificity. The increased information on affected areas of the brain afforded by K-Bayes reconstruction should provide extra power for the study and clinical management of neurodegenerative (and other brain) diseases.

ACKNOWLEDGEMENTS

OPEN ACCESS

REFERENCES

1. Gelman A: Bayesian Data Analysis, Boca Raton: Chapman & Hall/CRC, 2004

2. Winkler G: Image Analysis, Random Fields, and Markov Chain Monte Carlo Methods: A Mathematical Introduction, New York: Springer, 2003

3. Dempster AP, Laird NM, Rubin DB: Maximum likelihood from incomplete data via the EM algorithm (with discussion). J R Stat Soc Ser B 39:1–38, 1977

4. Green PJ: Bayesian reconstructions from emission tomography data using a modified EM algorithm. IEEE Trans Med Imag 9:84–93, 1990

5. Zhang Y, Brady M, Smith S: Segmentation of brain MR images through a hidden Markov random field model and the expectation maximization algorithm. IEEE Trans Med Imag 20:45–57, 2001

6. Besag JE: Spatial interaction and the statistical analysis of lattice systems (with discussion). J R Stat Soc Ser B 36:192–236, 1974

7. Besag JE: Statistical analysis of non-lattice data. The Statistician 24:179–195, 1975

8. Besag JE: Towards Bayesian image analysis. J Appl Stat 16:395–407, 1989

9. Besag JE, Kooperberg C: On conditional and intrinsic autoregressions. Biometrika 82:733–746, 1995

10. Miller IM, Schaewe TJ, Cohen SB, Ackerman JJH: Model-based maximum-likelihood estimation for phase- and frequency-encoded magnetic-resonance-imaging data. J Magn Reson Ser B 107:10–221, 1995