

dcemri: A Package for Medical Image Analysis

Brandon Whitcher

bjw34032@users.sourceforge.net

Volker J. Schmid

volkerschmid@users.sourceforge.net

September 4, 2009

1 Introduction

Quantitative analysis of perfusion imaging using dynamic contrast-enhanced MRI (DCE-MRI) is achieved through a series of processing steps, starting with the raw data acquired from the MRI scanner, and involves a combination of physics, mathematics, engineering and statistics. The purpose of the **dcemri** package is to provide a collection of functions that move the experimental data through all steps of the data analysis pipeline using standard data formats that may be visualized and manipulated across a wide variety of software packages.

2 Data Input/Output

2.1 Data Formats and Conversion

The **dcemri** package requires incoming data to be in the ANALYZE 7.5 or NIfTI formats. Data conversion (e.g., from DICOM to NIfTI) must be performed by the user before **dcemri** may be used to summarize the data. Several software packages allow DICOM-to-NIfTI (or ANALYZE) conversion; e.g.,

- FreeSurfer (surfer.nmr.mgh.harvard.edu)
- Xmedcon (xmedcon.sourceforge.net)
- MRIConvert (lnci.oregon.edu/~jolinda/MRIConvert)

This is by no means an exhaustive list of software available for DICOM conversion.

2.2 A Note on Axes and Orientation

The NIfTI format contains an implicit generalized spatial transformation from the data co-ordinate system (i, j, k) into a real-space “right-handed” co-ordinate system. In this real-space system, the (x, y, z) axes are *usually* set such that x increases from left to right, y increases from posterior to anterior and z increases from inferior to superior.

At this point in time the **dcmri** package cannot apply an arbitrary transform to the imaging data into (x, y, z) space – such a transform may require non-integral indices and interpolation steps. The package does accommodate straightforward transformations of imaging data; e.g., setting the x -axis to increase from right to left (neurological). Future versions of **dcmri** will attempt to address more complicated transformations.

2.3 Labelled LR Standard (MNI152) Images in NIfTI Format

The first example of reading in, and displaying, medical imaging data in NIfTI format (`avg152T1_LR_nifti.nii.gz`) was obtained from the NIfTI website (nifti.nimh.nih.gov/nifti-1/). Successful execution of the command:

```
> mni <- read.img("avg152T1_LR_nifti.nii.gz")
> dim(mni)

[1] 91 109 91 1

> X <- nrow(mni)
> Y <- ncol(mni)
> Z <- nsli(mni)
> zrange <- range(mni)
> par(mfrow = c(10, 10), mar = rep(0, 4))
> for (z in 1:Z) {
+   image(1:X, 1:Y, mni[, , z, 1], zlim = zrange, col = grey(0:64/64),
+       xlab = "", ylab = "", axes = FALSE)
+ }
```

produces a 4D array of the image data, with the default NIfTI axes, and is displayed on a 10×10 grid of images (Figure 1).

The second example of reading in, and displaying, medical imaging data in NIfTI format (`avg152T1_RL_nifti.nii`) was also obtained from the NIfTI website (nifti.nimh.nih.gov/nifti-1/). Successful execution of the command

```
> mni <- read.img("avg152T1_RL_nifti.nii.gz")
> dim(mni)

[1] 91 109 91 1

> X <- nrow(mni)
> Y <- ncol(mni)
> Z <- nsli(mni)
> zrange <- range(mni)
> par(mfrow = c(10, 10), mar = rep(0, 4))
> for (z in 1:Z) {
+   image(1:X, 1:Y, mni[, , z, 1], zlim = zrange, col = grey(0:64/64),
+       xlab = "", ylab = "", axes = FALSE)
+ }
```




Figure 1: Axial slices of MNI volume `avg152T1_LR_nifti` stored in radiological convention.

produces a 4D array of the image data that may be displayed in a 10×10 grid of images (Figure 2).




Figure 2: Axial slices of MNI volume `avg152T1_RL_nifti` stored in neurological convention.

The first image (LR) is stored in radiological convention. The second image (RL) is stored in neurological convention. Any NIfTI-1 compliant viewing software should display these images identically.

2.4 Simple Time-series or Multi-volume Image

This is an example of reading in, and displaying, a four-dimensional medical imaging data set in NIfTI format (`filtered_func_data.nii`) obtained from the NIfTI website (nifti.nimh.nih.gov/nifti-1/). Successful execution of the command

```
> ffd <- read.img("filtered_func_data.nii.gz")
> dim(ffd)
```

```
[1] 64 64 21 180
```

```

> X <- nrow(ffd)
> Y <- ncol(ffd)
> Z <- nsli(ffd)
> W <- ntim(ffd)
> zrange <- range(ffd)
> w <- 1
> par(mfrow = c(5, 5), mar = rep(0, 4))
> for (z in 1:Z) {
+   image(1:X, 1:Y, ffd[, , z, w], zlim = zrange, col = grey(0:64/64),
+       xlab = "", ylab = "", axes = FALSE)
+ }

```

produces a four-dimensional (4D) array of imaging data that may be displayed in a 5×5 grid of images (Figure 3). The first three dimensions are spatial locations of the voxel (volume element) and the fourth dimension is time.




Figure 3: Axial slices of the functional MRI “volume” `filtered_func_data` from the first acquisition.

2.5 Statistic Image

This is an example of reading in and displaying a statistical image so that it may be overlaid on the EPI (echo planar imaging) data taken from the functional MRI experiment. The original NIfTI files (`filtered_func_data.nii` and `zstat1.nii`) were obtained from the NIfTI website (nifti.nimh.nih.gov/nifti-1/). Successful execution of the command

```
> zstat1 <- read.img("zstat1.nii.gz")
> dim(zstat1)

[1] 64 64 21 1

> par(mfrow = c(5, 5), mar = rep(0, 4))
> w <- 1
> for (z in 1:Z) {
+   image(1:X, 1:Y, ffd[, , z, w], zlim = zrange, col = grey(0:64/64),
+       xlab = "", ylab = "", axes = FALSE)
+   image(1:X, 1:Y, ifelse(abs(zstat1[, , z, 1]) > 5, zstat1[, ,
+       z, 1], NA), zlim = range(zstat1), col = hotmetal(), add = TRUE)
+ }
```

produces a 4D array of parameter estimates (essentially coefficients from a linear regression performed at each voxel) that may be overlaid on the original data for anatomical reference (Figure 4).

3 Motion Correction and Co-registration

Basic motion correction within an acquisition, and co-registration between acquired series, is available using template matching (en.wikipedia.org/wiki/Template_matching). A reference volume must be pre-specified where a mask has been applied to remove all voxels that should not be included in the algorithm. Note, only three-dimensional translations are allowed and no interpolation is used (i.e., only whole-voxel translations) at this time.

4 T1 Relaxation and Gadolinium Concentration

Estimation of the tissue T1 relaxation rate is the first step in converting signal intensity, obtained in the dynamic acquisition of the DCE-MRI protocol, to contrast agent concentration. The subsequent steps provided here focus on pharmacokinetic modeling and assumes one has converted the dynamic acquisition to contrast agent concentration. Please see [Collins and Padhani \(2004\)](#) for a discussion on this point.

There are a myriad of techniques to quantify T1 using MRI. Currently curve-fitting methods for two popular acquisition schemes are available

- Inversion recovery (www.e-mri.org/mri-sequences/inversion-recovery-stir-flair.html)




Figure 4: Axial slices of the functional MRI data with with the statistical image overlayed. The test statistics were thresholded at $|Z| \geq 4$.

- Multiple flip angles ([Parker and Padhani, 2003](#))

Once the tissue T1 relaxation rate has been estimated, the dynamic acquisition is then converted to contrast agent concentration. Note, the B1 field is assumed to be constant (and accurate) when using multiple flip angles to estimate T1. At higher fields strengths (e.g., 3T) the B1 field should be estimated in order to correct the prescribed flip angles.

4.1 B1 Mapping Via the Saturated Double-Angle Method

For in vivo MRI at high field (≥ 3 Tesla) it is essential to consider the homogeneity of the active B1 field (B1+). The B1+ field is the transverse, circularly polarized component of B1 that is rotating in the same sense as the magnetization. When exciting or manipulating large collections of spins, non-uniformity in B1+ results in nonuniform treatment of spins. This leads to spatially varying image signal and image contrast and to difficulty in image interpretation and image-based quantification ([Cunningham *et al.*, 2006](#)).

The proposed method uses an adaptation of the double angle method (DAM). Such methods allow calculation of a flip-angle map, which is an indirect measure of the B1+ field. Two images are acquired: I_1 with prescribed tip α_1 and I_2 with prescribed tip $\alpha_2 = 2\alpha_1$. All other signal-affecting sequence parameters are kept constant. For each voxel, the ratio of magnitude images satisfies

$$\frac{I_2(r)}{I_1(r)} = \frac{\sin \alpha_2(r) f_2(T_1, \text{TR})}{\sin \alpha_1(r) f_1(T_1, \text{TR})}$$

where r represents spatial position and $\alpha_1(r)$ and $\alpha_2(r)$ are tip angles that vary with the spatially varying B1+ field. If the effects of T_1 and T_2 relaxation can be neglected, then the actual tip angles as a function of spatial position satisfy

$$\alpha(r) = \arccos \left(\left| \frac{I_2(r)}{2I_1(r)} \right| \right)$$

A long repetition time ($\text{TR} \leq 5T_1$) is typically used with the double-angle methods so that there is no T_1 dependence in either I_1 or I_2 (i.e., $f_1(T_1, \text{TR}) = f_2(T_1, \text{TR}) = 1.0$). Instead, the proposed method includes a magnetization-reset sequence after each data acquisition with the goal of putting the spin population in the same state regardless of whether the α_1 or α_2 excitation was used for the preceding acquisition (i.e., $f_1(T_1, \text{TR}) = f_2(T_1, \text{TR}) \neq 1.0$).

4.2 Example

Using data acquired from a T1 phantom at two flip angles, $\alpha_1 = 60^\circ$ and $\alpha_2 = 120^\circ$, we compute the multiplicative factor relative to the low flip angle using the saturated double-angle method ([Cunningham *et al.*, 2006](#)).

```
> sdam60 <- read.img("SDAM_ep2d_60deg_26slc.nii.gz")[, , 1:26, ]
> dim(sdam60)

[1] 64 64 26 10

> sdam120 <- read.img("SDAM_ep2d_120deg_26slc.nii.gz")[, , 1:26, ]
> dim(sdam120)
```



```
[1] 64 64 26 10
```

```
> sdam.image <- rowMeans(dam(sdam60, sdam120, 60), dims = 3)
> mask <- rowSums(sdam60, dims = 3) > 500
> dim(mask)
```

```
[1] 64 64 26
```

```
> zi <- 10:13
> w <- 5
> SDAM <- read.img("SDAM_smooth.nii.gz")[, , 1:26, , drop = TRUE]
> par(mfrow = c(2, 2), mar = rep(0, 4))
> for (z in zi) {
+   image(sdam120[, , z, w], col = grey(0:128/128), xlab = "", ylab = "",
+         axes = FALSE)
+   image(ifelse(mask[, , z], SDAM[, , z], 0), zlim = c(0.5, 1.5),
+         col = tim.colors(), add = TRUE)
+ }
```

Figure 5 is the estimated B1+ field (with isotropic Gaussian smoothing) for a gel-based phantom containing a variety of T1 relaxation times. The center of the phantom exhibits a flip angle $> 60^\circ$ while the flip angle rapidly becomes $< 60^\circ$ when moving away from the center in either the x , y or z dimensions. Isotropic smoothing should be applied before using this field to modify flip angles associated with additional acquisitions (e.g., in the **AnalyzeFMRI** package).




Figure 5: Estimated B1+ field (with isotropic Gaussian smoothing) using the saturated double-angle method. The colors correspond to a multiplicative factor relative to the true flip angle (60°).

Assuming the smoothed version of the B1+ field has been computed (SDAM), multiple flip-angle acquisitions can be used to estimate the T1 relaxation rate from the subject (or phantom). The multiplicative factor, derived from the saturated double-angle method, is used to produce a spatially-varying flip-angle map and input into the appropriate function.

```
> fnames <- c("fl3d_vibe-5deg", "fl3d_vibe-10deg", "fl3d_vibe-20deg",
+ "fl3d_vibe-25deg", "fl3d_vibe-15deg")
> alpha <- c(5, 10, 20, 25, 15)
> nangles <- length(alpha)
> X <- Y <- 64
> Z <- 26
> flip <- fangles <- array(0, c(X, Y, Z, nangles))
> for (w in 1:nangles) {
+   vibe <- read.img(fnames[w])
+   flip[, , , w] <- vibe
+   fangles[, , , w] <- array(alpha[w], c(X, Y, Z))
+ }
> TR <- 4.22/1000
> fanglesB1 <- fangles * array(SDAM, c(X, Y, Z, nangles))
> zi <- 10:13
> mask[, , (!1:Z %in% zi)] <- FALSE
> R1 <- R1.fast(flip, mask, fanglesB1, TR, verbose = TRUE)

Deconstructing data...
Calculating R10 and M0...
Reconstructing results...

> par(mfrow = c(2, 2), mar = rep(0, 4))
> for (z in zi) {
+   image(drop(vibe)[, , z], zlim = c(0, 1024), col = grey(0:64/64),
+         xlab = "", ylab = "", axes = FALSE)
+   image(1/drop(R1$R10)[, , z], zlim = c(0, 2.5), col = hotmetal(),
+         add = TRUE)
+ }
```

Figure 6 displays the quantitative T1 map for a gel-based phantom using information from the estimated B1+ field.

By defining regions of interest (ROIs) in

```
> pmask <- read.img("t1_phantom_mask.nii.gz")
> pmask[, , c(1:24, 26), 1] <- pmask[, , 25, 1]
> dim(pmask)
```

```
[1] 64 64 26 1
```

We may compare the “true” T1 values for each ROI with those obtained from acquiring multiple flip angles with the application of B1 mapping. Figure 7 compares T1 estimates




Figure 6: Estimated T1 relaxation rates for the phantom data acquisition. The colors range from 0-2.5 seconds.

in the 10 ROIs, defined by `pmask`, with the true T1 values (large circles). The first seven ROIs correspond to the cylinders that run around the phantom, clockwise starting from approximately one o'clock. The eighth and ninth ROIs are taken from the main compartment in the phantom; ROI eight is drawn in the middle of the phantom while ROI nine is drawn from the outside of the phantom. The final ROI is taken from the central cylinder embedded in the phantom.

4.3 Contrast Agent Concentration

The `CA.fast` function rearranges the assumed multidimensional (2D or 3D) structure of the multiple flip-angle data into a single matrix to take advantage of internal R functions instead of loops, and called `E10.1m`. Conversion of the dynamic signal intensities to contrast agent concentration is performed via

$$[\text{Gd}] = \frac{1}{r_1} \left(\frac{1}{T_1} - \frac{1}{T_{10}} \right),$$

where r_1 is the spin-lattice relaxivity constant and T_{10} is the spin-lattice relaxation time in the absence of contrast media (Buckley and Parker, 2005). For computational reasons, we follow the method of Li *et al.* (2000).

4.4 Arterial Input Function

Whereas quantitative PET studies routinely perform arterial cannulation on the subject in order to characterize the arterial input function (AIF), it has been common to use literature-based AIFs in the DCE-MRI literature. Examples include

$$C_p(t) = D \left(a_1 e^{-m_1 t} + a_2 e^{-m_2 t} \right),$$

```

> T1 <- c(0.484, 0.35, 1.07, 0.648, 0.456, 1.07, 0.66, 1.543, 1.543,
+        0.353)
> par(mfrow = c(1, 1), mar = c(5, 4, 4, 2) + 0.1)
> boxplot(split(1/drop(R1$R10), as.factor(drop(pmask))))[-1], ylim = c(0,
+        2.5), xlab = "Region of Interest", ylab = "T1 (seconds)")
> points(1:10, T1, bg = rainbow(10), pch = 21, cex = 2)

```




Figure 7: Boxplots of the estimated T1 values for the gel-based phantom, grouped by user-specified regions of interest. True T1 values are plotted as colored circles for each distinct ROI.

where $D = 0.1$ mmol/kg, $a_1 = 3.99$ kg/l, $a_2 = 4.78$ kg/l, $m_1 = 0.144$ min⁻¹ and $m_2 = 0.0111$ min⁻¹ (Weinmann *et al.*, 1984; Tofts and Kermode, 1984); or $D = 1.0$ mmol/kg, $a_1 = 2.4$ kg/l, $a_2 = 0.62$ kg/l, $m_1 = 3.0$ and $m_2 = 0.016$ (Fritz-Hansen *et al.*, 1996). There has been progress in measuring the AIF using the dynamic acquisition and fitting a parametric model to the observed data. Recent models include Parker *et al.* (2006) and Orton *et al.* (2008). **dcmri** has incorporated one of these parametric models from Orton *et al.* (2008)

$$C_p(t) = A_B t e^{-\mu_B t} + A_G (e^{-\mu_G t} + e^{-\mu_B t})$$

which can be fitted to the observed data using nonlinear regression. Using the AIF defined in Buckley (2002), we illustrate fitting a parametric model to characterize observed data. The `orton.exp.lm` function provides this capability using a common double-exponential parametric form.

```
> data("buckley")
> aifparams <- with(buckley, orton.exp.lm(time.min, input))
> fit.aif <- with(aifparams, aif.orton.exp(buckley$time.min, AB, muB,
+     AG, muG))
```

Figure 8 shows both the true AIF and the best parametric description using a least-squares fitting criterion.

5 Kinetic Parameter Estimation

The standard Kety model (Kety, 1960), a single-compartment model, or the extended Kety model, the standard Kety model with an extra “vascular” term, form the collection of basic parametric models one can apply using **dcmri**. Regardless of which parametric model is chosen for the biological system, the contrast agent concentration curve at each voxel in the region of interest (ROI) is approximated using the convolution of an arterial input function (AIF) and the compartmental model; e.g.,

$$\begin{aligned} C_t(t) &= K^{\text{trans}} [C_p(t) \otimes e^{-k_{\text{ep}} t}], \\ C_t(t) &= v_p C_P(t) + K^{\text{trans}} [C_p(t) \otimes e^{-k_{\text{ep}} t}]. \end{aligned}$$

Parameter estimation is achieved using one of two options in the current version of this software:

- Non-linear regression using non-linear least squares (Levenburg-Marquardt optimization)
- Bayesian estimation using Markov chain Monte Carlo (MCMC) (Schmid *et al.*, 2006)

Least-square estimates of the kinetic parameters K^{trans} and k_{ep} (also v_p for the extended Kety model) are provided in `dcmri.lm` while the posterior median is provided in `dcmri.bayes`. When using Bayesian estimation all samples from the joint posterior distribution are also provided, allowing one to interrogate the empirical probability density function (PDF) of the parameter estimates.

```
> with(buckley, plot(time.min, input, type = "l", lwd = 2, xlab = "Time (minutes)",  
+   ylab = ""))  
> with(buckley, lines(time.min, fit.aif, lwd = 2, col = 2))  
> legend("topright", c("Simulated AIF", "Estimated AIF"), lwd = 2,  
+   col = 1:2)
```




Figure 8: Arterial input function (AIF) from [Buckley \(2002\)](#) and the best parametric fit, using the exponential model from [Orton *et al.* \(2008\)](#).

Using the simulated breast data from [Buckley \(2002\)](#), we illustrate fitting the “extended Kety” model to the contrast agent concentration curves using the exponential model for the AIF. We use non-linear regression to fit the data on an under-sampled subset (in time) of the simulated curves.

```
> xi <- seq(5, 300, by = 5)
> img <- array(t(breast$data)[, xi], c(13, 1, 1, 60))
> time <- buckley$time.min[xi]
> aif <- buckley$input[xi]
> mask <- array(TRUE, dim(img)[1:3])
> aifparams <- orton.exp.lm(time, aif)
> fit <- dcmri.lm(img, time, mask, model = "orton.exp", aif = "user",
+   user = aifparams)
```

Figure 9 displays the 13 unique simulated curves along with the fitted curves from the compartmental model. There is decent agreement between the observed and fitted values, except for Series 6 which changes too rapidly in the beginning and cannot be explained by the parametric model.

6 Statistical Inference

No specific support is provided for hypothesis testing in **dcmri**. We recommend one uses built-in facilities in R to perform ANOVA (analysis of variance) or mixed-effects models based on statistical summaries of the kinetic parameters over the ROI per subject per visit. An alternative to this traditional approach is to analyze an entire study using a Bayesian hierarchical model ([Schmid et al., 2009](#)), available in the software project **PILFER** ([pilfer.sourceforge.net](#)).

One may also question the rationale for hypothesis testing in only one kinetic parameter. Preliminary work has been performed in looking at the joint response to treatment of both K^{trans} and k_{ep} in DCE-MRI by [O'Connor et al. \(2010\)](#).

References

- Buckley, D. L. (2002). Uncertainty in the analysis of traced kinetics using dynamic contrast-enhanced T_1 -weighted MRI. *Magnetic Resonance in Medicine* 47, 601–606.
- Buckley, D. L. and G. J. M. Parker (2005). Measuring contrast agent concentration in T_1 -weighted dynamic contrast-enhanced MRI. In A. Jackson, D. L. Buckley, and G. J. M. Parker (Eds.), *Dynamic Contrast-Enhanced Magnetic Resonance Imaging in Oncology*, pp. 69–80. Berlin: Springer.
- Collins, D. J. and A. R. Padhani (2004). Dynamic magnetic resonance imaging of tumor perfusion. *IEEE Engineering in Biology and Medicine Magazine*, 65–83.
- Cunningham, C., J. Pauly, and K. Nayak (2006). Saturated double-angle method for rapid B1+ mapping. *Magnetic Resonance in Medicine* 55, 1326–1333.

```

> par(mfrow = c(4, 4), mar = c(5, 4, 4, 2)/1.25, mex = 0.8)
> for (x in 1:nrow(img)) {
+   plot(time, img[x, 1, 1, ], ylim = range(img), xlab = "Time (minutes)",
+     ylab = "", main = paste("Series", x))
+   kinparams <- with(fit, c(vp[x, 1, 1], ktrans[x, 1, 1], kep[x,
+     1, 1]))
+   lines(time, model.orton.exp(time, aifparams[1:4], kinparams),
+     lwd = 1.25, col = 2)
+ }

```

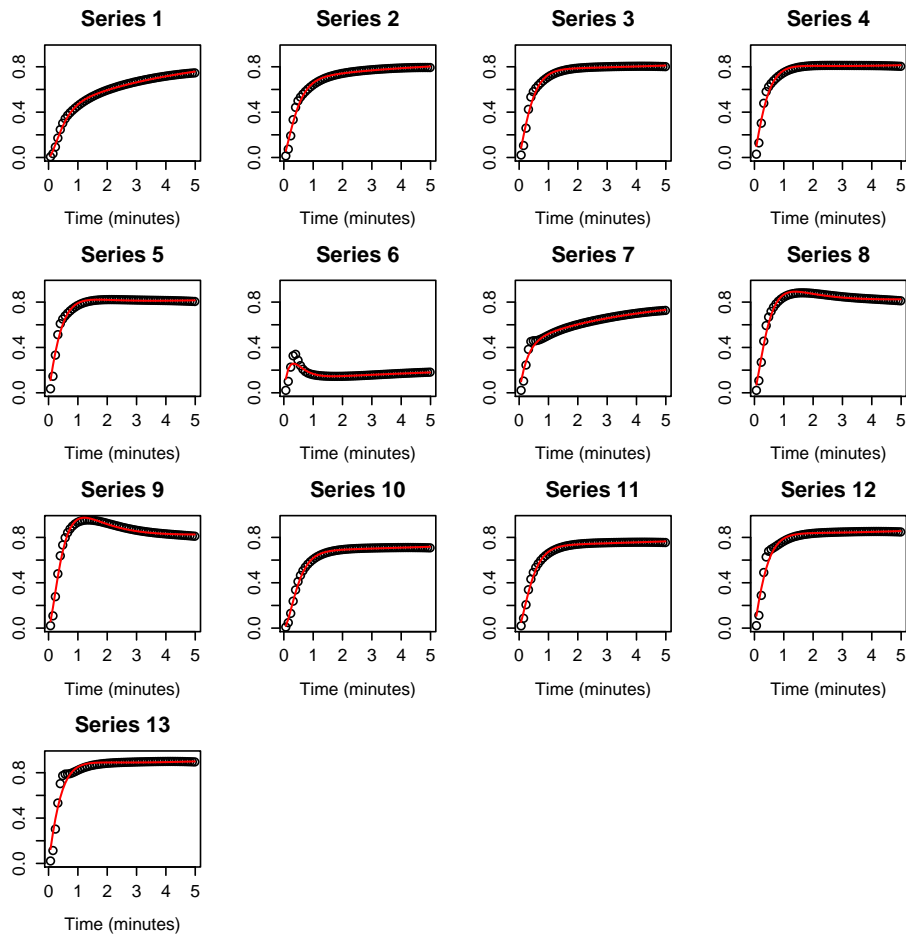


Figure 9: Simulated signal intensity curves from Buckley (2002), for breast tissue, with the best parametric fit using an exponential model for the AIF and the “extended Kety” model.

- Fritz-Hansen, T., E. Rostrup, H. B. W. Larsson, L. S ndergaard, P. Ring, and O. Henriksen (1996). Measurement of the arterial concentration of Gd-DTPA using MRI: A step toward quantitative perfusion imaging. *Magnetic Resonance in Medicine* 36, 225–231.
- Kety, S. (1960). Blood–tissue exchange methods. Theory of blood-tissue exchange and its application to measurement of blood flow. *Methods in Medical Research* 8, 223–227.
- Li, K.-L., X. P. Zhu, J. Waterton, and A. Jackson (2000). Improved 3D quantitative mapping of blood volume and endothelial permeability in brain tumors. *Journal of Magnetic Resonance Imaging* 12, 347–357.
- O’Connor, E., N. Fieller, A. Holmes, J. C. Waterton, and E. Ainscow (2010). Functional principal component analyses of biomedical images as outcome measures. *Journal of the Royal Statistical Society C (Applied Statistics)*. *in press*.
- Orton, M. R., J. A. d’Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach (2008). Computationally efficient vascular input function models for quantitative kinetic modelling using DCE-MRI. *Physics in Medicine and Biology* 53, 1225–1239.
- Parker, G. J. M. and A. R. Padhani (2003). T_1 -w DCE-MRI: T_1 -weighted dynamic contrast-enhanced MRI. In P. Tofts (Ed.), *Quantitative MRI of the Brain: Measuring Changes Caused by Disease*, Chapter 10, pp. 341–364. Chichester, UK: Wiley.
- Parker, G. J. M., C. Roberts, A. Macdonald, G. A. Buonaccorsi, S. Cheung, D. L. Buckley, A. Jackson, Y. Watson, K. Davies, and G. C. Jayson (2006). Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced MRI. *Magnetic Resonance in Medicine* 56, 993–1000.
- Schmid, V., B. Whitcher, A. R. Padhani, N. J. Taylor, and G.-Z. Yang (2006). Bayesian methods for pharmacokinetic models in dynamic contrast-enhanced magnetic resonance imaging. *IEEE Transactions on Medical Imaging* 25(12), 1627–1636.
- Schmid, V. J., B. Whitcher, A. R. Padhani, N. J. Taylor, and G.-Z. Yang (2009). A Bayesian hierarchical model for the analysis of a longitudinal dynamic contrast-enhanced MRI oncology study. *Magnetic Resonance in Medicine* 61(1), 163–174.
- Tofts, P. S. and A. G. Kermode (1984). Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts. *Magnetic Resonance in Medicine* 17(2), 357–367.
- Weinmann, H. J., M. Laniado, and W. Mutzel (1984). Pharmacokinetics of Gd-DTPA/dimeglumine after intravenous injection into healthy volunteers. *Physiological Chemistry and Physics and Medical NMR* 16, 167–172.