

Data-Driven Multi-Contrast Spectral Microstructure Imaging with InSpect

Paddy J. Slator¹, Jana Hutter^{2,3}, Razvan V. Marinescu¹, Marco Palombo¹,
Laurence H. Jackson^{2,3}, Alison Ho⁴, Lucy C. Chappell⁴, Mary Rutherford²,
Joseph V. Hajnal^{2,3}, and Daniel C. Alexander¹

¹ Centre for Medical Image Computing, Department of Computer Science, University
College London, UK

² Centre for the Developing Brain, King's College London, London, UK

³ Biomedical Engineering Department, King's College London, London, UK

⁴ Women's Health Department, King's College London, London, UK
p.slator@ucl.ac.uk

Abstract. We introduce and demonstrate an unsupervised machine learning method for spectroscopic analysis of quantitative MRI (qMRI) experiments. qMRI data can support estimation of multidimensional correlation (or single-dimensional) spectra, which allow model-free investigation of tissue properties, but this requires an ill-posed calculation. Moreover, in the vast majority of applications ground truth knowledge is unobtainable, preventing the application of supervised machine learning. Here we present a new method that addresses these limitations in a data-driven way. The algorithm simultaneously estimates a canonical basis of spectral components and voxelwise maps of their weightings, thereby pooling information across whole images to regularise the ill-posed problem. We show that our algorithm substantially outperforms current voxelwise spectral approaches. We demonstrate the method on combined diffusion-relaxometry placental MRI scans, revealing anatomically-relevant substructures, and identifying dysfunctional placentas. Our algorithm vastly reduces the data required to reliably estimate multidimensional correlation (or single-dimensional) spectra, opening up the possibility of spectroscopic imaging in a wide range of new applications.

Keywords: Quantitative MRI · Unsupervised Learning · Placenta MRI

1 Introduction

Continuum modelling is an attractive method for analysing quantitative MRI (qMRI) data. The technique assumes that spins have a distribution of values (e.g. relaxivity, diffusivity), which are quantified by a multidimensional correlation (or single-dimensional) spectrum. This approach is particularly powerful for experiments that simultaneously measure multiple MR properties, as it can resolve microstructural compartments that are indistinguishable with single contrast data. Imaging studies have recently demonstrated this in the T1-diffusion

[8, 7, 20], T2-diffusion [25, 12, 17], T1-T2-diffusion [3], T2*-diffusion [24], T1-T2 [15], and T1-T2*-diffusion [11] domains.

A general continuum model gives a Fredholm integral equation on the MR signal [4], or a Laplace transform in the specific exponential decay case. The spectrum can be estimated using regularised inversion of this integral [9], although this is highly ill-posed. Estimating spectra in each image voxel independently therefore requires unrealistically high signal-to-noise (SNR). Moreover, to derive meaningful image maps from voxelwise spectra typically involves a procedure known as spectral integration. In spectral integration, the user manually identifies regions of the spectrum (termed spectral regions of interest, sROI) that correspond to particular features of spectral components of interest. Scalar indices are then calculated by numerical integration of each reconstructed spectrum over these sROIs [16, 12, 3].

Recently, methods have been proposed for increasing the robustness of voxelwise spectral fits, utilising marginal distributions [2] or spatial regularisation [12]. These methods can improve inversion stability and give more meaningful derived spatial maps. However, inherent limitations remain. In particular, the reliance on ad-hoc choices of regularisation terms and manually defined sROIs. A recent technique automatically identifies these spectral integration regions [21], but restricts to rectangular and non-overlapping sROIs, depends on a user-defined threshold value, and still requires voxelwise spectral estimation.

Here we present a method which addresses these limitations in a data-driven way. It simultaneously estimates a canonical basis of spectral components for a whole image (or a data set comprising multiple images), and the voxelwise weighting factors of each component. We build on the discrete InSpect algorithm previously published by Slator et al. [23], adapting it for the continuous case. This allows us to capture smooth changes in parameters across the image with high resolution values, rather than forcing hard categorisation of pixels into a small set of bins, which discards subtle variation. Unlike standard inversion approaches, the InSpect method introduced here exploits the huge dependence among voxels, dramatically reducing the SNR required for stable inversion, hence enabling spectroscopic imaging in a wide variety of new situations.

2 Methods

Multidimensional Spectrum Estimation. InSpect is based on a continuum model, which assumes that single voxels contain spins with a spectrum of MR properties. For a general n -dimensional MRI experiment the voxel signal is

$$S(\mathbf{t}) = \int \dots \int F(\boldsymbol{\omega}) K(\mathbf{t}, \boldsymbol{\omega}) d\omega_1 \dots d\omega_n \quad (1)$$

where \mathbf{t} is a vector of experimental parameters which are varied to yield contrast in intrinsic MR properties $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)$, via the specific form of the kernel $K(\mathbf{t}, \boldsymbol{\omega})$. F is the n -dimensional spectrum over $\boldsymbol{\omega}$, i.e. the distribution of these values across all spins. For example, in T2*-diffusivity (or T2-diffusivity)

imaging, $\mathbf{t} = (b, T_E)$, the b-value and echo time (TE); $\boldsymbol{\omega} = (ADC, T_2^*)$, the apparent diffusion coefficient and T2*; $K(\mathbf{t}, \boldsymbol{\omega}) = \exp(-T_E/T_2^*) \exp(-bADC)$ is the kernel; and F is the 2D T2*-ADC spectrum.

The standard approach for estimating the spectrum, following [18, 9, 22], proceeds as follows. Equation (1) is first discretised onto an n-dimensional grid, with lengths defined by the user-defined vector $N_{\boldsymbol{\omega}} = (N_{\omega_1}, \dots, N_{\omega_n})$. By choosing an ordering of the grid coordinates, the signal for all MR encodings in the experiment can thus be written in matrix form as

$$\mathbf{S} = \mathbf{K}\mathbf{F} \quad (2)$$

where \mathbf{S} is a column vector, length N_s , of the signals at each MR encoding, K is an N_s by $\prod_{m=1}^n N_{\omega_m}$ matrix of discretised kernel values, and F is an $\prod_{m=1}^n N_{\omega_m}$ length column vector of spectrum values. The spectrum F is then calculated with regularised non-negative least squares

$$\mathbf{F} = \arg \min_{\mathbf{F} \geq 0} \|\mathbf{K}\mathbf{F} - \mathbf{S}\|_2^2 + \alpha \|\mathbf{F}\|_2^2. \quad (3)$$

By solving Equation (3) in each voxel, the spectrum can be estimated across a whole image. Volume fraction maps are then produced by numerically integrating voxelwise spectra over user-defined regions of the spectrum, e.g. [14, 3, 16]. However low SNR can lead to noisy spectrum estimates and hence poor mappings.

Underlying model. Our algorithm automates spectral mapping, and undertakes data-driven regularisation of the Fredholm integral (or Laplace transform) inversion. Rather than naively fitting spectra to each voxel independently, we learn a low-dimensional representation consistent with the whole image.

The first element of the representation is a pre-specified number, M , of canonical spectral components, $\{F_1, F_2, \dots, F_M\}$ (see first column of Figure 1 for an example). Each spectral component has corresponding voxelwise weights across all N image voxels. The weighting of component m in voxel n is denoted z_{nm} , so that the full set of voxelwise weights is

$$\mathbf{z}_n = \{z_{n1}, z_{n2}, \dots, z_{nM}\}_{n=1}^N, \text{ subject to } \sum_{m=1}^M z_{nm} = 1. \quad (4)$$

The second column of Figure 1 shows example voxelwise weights. The signal from each voxel, \mathbf{S}_n , is described by the continuum model of Equation (1) with the effective spectrum in each voxel n modelled as a weighted sum of the canonical spectrum components

$$F(\mathbf{z}_n) = \sum_{m=1}^M z_{nm} \mathbf{F}_m \quad (5)$$

where $\mathbf{z}_n = \{z_{nm}\}_{m=1}^M$ are the component weights for voxel n . The discrete model for a single voxel is therefore given by $\mathbf{S}_n = \mathbf{K}F(\mathbf{z}_n)$. We assume Gaussian noise,

giving the following whole-image log-likelihood

$$\log \pi(\{\mathbf{S}_n\}_{n=1}^N | \{\mathbf{F}_m\}_{m=1}^M, \{\mathbf{z}_n, \sigma_n^2\}_{n=1}^N) = \sum_{n=1}^N \log \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I). \quad (6)$$

Note that we assume all observations in a voxel have the same variance, i.e. the covariance matrix is $\sigma_n^2 I$. For simplicity we denote the log-likelihood $\log \pi(\mathbf{D}|\theta)$, where $\mathbf{D} = \{\mathbf{S}_n\}_{n=1}^N$ is the full multi-contrast MR dataset (either a single image or set of images), and $\theta = \{\{\mathbf{F}_m\}_{m=1}^M, \{\mathbf{z}_n\}_{n=1}^N, \{\sigma_n^2\}_{n=1}^N\}$ are the model parameters.

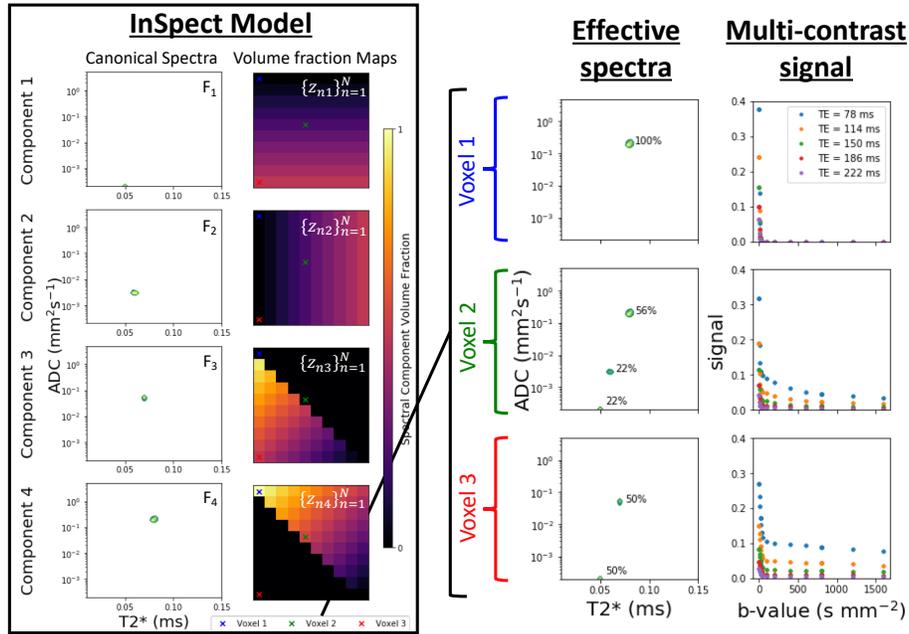


Fig. 1. The model underlying InSpect consists of a set of M components. Each component has an associated canonical spectrum and volume fraction map, which combine to give the effective voxelwise spectra.

InSpect Inference Algorithm We seek the parameters θ that maximise $\log \pi(\mathbf{D}|\theta)$. In practice, we estimate σ_n^2 directly from the data, e.g. for a T2*-diffusion experiment we estimate by calculating the empirical variance of the volumes with $b=0$ and the lowest TE. To derive maximisation steps for the canonical spectral components $\{\mathbf{F}_m\}_{m=1}^M$, and voxelwise weights, $\{z_{nm}\}_{n=1}^N$ we

first rearrange the log-likelihood to give

$$\log \pi(\mathbf{D}|\theta) = \sum_{n=1}^N -\frac{1}{2} \log 2\pi\sigma_n^2 - \frac{1}{2\sigma_n^2} \|\mathbf{S}_n - \mathbf{K}F(\mathbf{z}_n)\|_2^2. \quad (7)$$

We now derive a maximisation step for a canonical spectrum, F_j conditioned on the voxelwise weights $\{\mathbf{z}_n\}_{n=1}^N$, and the other spectra $\{F_m\}_{m \neq j}$. For the canonical spectral components, $\mathbf{F}_1, \dots, \mathbf{F}_M$, we need to solve $F_k = \arg \max_{F_k \geq 0} \log \pi(\mathbf{D}|\theta)$. Taking the derivation of the function with respect to F_k (in numerator layout), setting equal to zero, and rearranging gives

$$\mathbf{F}_k = \arg \min_{\mathbf{F}_k \geq 0} \left\| \left(\sum_{n=1}^N \mathbf{K}z_{nk} \right) \mathbf{F}_k - \sum_{n=1}^N \left(\mathbf{S}_n - \mathbf{K} \left(\sum_{m \neq k} z_{nm} \mathbf{F}_m \right) \right) \right\|_2^2. \quad (8)$$

We can solve Equation (8) with non-negative least squares as described earlier.

To find the maximum likelihood voxelwise weightings for each of the spectral components, we first note the posterior distribution for the model, up to proportionality

$$\pi(\theta|\mathbf{D}) \propto \pi(\mathbf{D}|\theta)\pi(\theta) = \prod_{n=1}^N \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I) \quad (9)$$

where we have assumed a uniform prior on all parameters θ . The posterior distribution for each \mathbf{z}_n - up to proportionality - is therefore

$$\pi(\mathbf{z}_n|\mathbf{S}_n, \theta^{(t-1)}) \propto \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I) \quad (10)$$

We therefore update $\mathbf{z}_n = \{z_{nm}\}_{m=1}^M$ by maximising this (in log-scale), subject to $\sum_{m=1}^M z_{nm} = 1$, i.e.

$$\mathbf{z}_n = \arg \max_{\sum_{m=1}^M z_{nm}=1} \log \mathcal{N}(\mathbf{S}_n; \mathbf{K}F(\mathbf{z}_n), \sigma_n^2 I) \quad (11)$$

which we solve sequentially for voxels $n = 1, \dots, N$ with the interior-point algorithm. The InSpect algorithm is hence the following iterative optimisation:

1. Initialise the canonical spectral components $\{F_1, F_2, \dots, F_M\}$, e.g. by assigning component to distinct elements of the whole-image spectrum.
2. Initialise the spectral weights for all voxels $\{z_{nm}\}_{n=1}^N$, given $\{F_1, F_2, \dots, F_M\}$
3. Update \mathbf{F}_m for some m by solving Equation (8)
4. Update $\{z_{nm}\}_{n=1}^N$ by solving Equation (11) for all voxels
5. Repeat steps 3 and 4 until convergence.

Application to Placenta Diffusion-Relaxometry Data We demonstrate our InSpect algorithm on placental T2*-diffusion data acquired on a 3T clinical MRI scanner using a 32-channel cardiac coil, previously published by Slator et al [24]. The protocol has 66 diffusion-weightings (ranging from $b = 5$ to 1600 s mm^{-2} , including six $b = 0$ volumes) and 5 TEs (78, 114, 150, 186, 222 ms) for a total of 330 contrast-encodings. Other acquisition parameters were FOV = $300 \times 320 \times 84 \text{ mm}$, TR = 7 s, SENSE = 2.5, halfscan = 0.6, resolution = 3 mm^3 . We considered 13 scans from 12 women, of whom 9 were categorised as healthy controls, two had chronic hypertension in pregnancy, and one had pre-eclampsia (PE) with additional fetal growth restriction (FGR). One participant with chronic hypertension was scanned twice, four weeks apart, and developed superimposed PE by the second scan. The algorithm was run on all scans simultaneously, on a manually-segmented ROI comprising the whole placenta and the adjacent section of uterine wall, with $M = 4$ components specified.

Application to Simulations We also tested on simulated diffusion-relaxometry data. Four synthetic canonical spectral components - informed by observed placental spectra [24] - were first defined (Figure 2, first column). We next defined ground truth voxelwise weights on a 50-by-50 image (Figure 2, third column). Given these, we simulated normalised diffusion-relaxometry scans using the same b-values and TEs as the placental data. We added Rician noise with SNR from 25-200 - this is comparable to the placenta data where we calculated SNRs ranging from 100 to 200. We applied InSpect - specifying $M = 4$ components - to these scans. We also fit voxelwise T2*-ADC spectra to all scans by solving Equation (3), with α set to 0.01 using the L-curve method [10], and hence derived volume fraction maps with the standard spectral integration approach [16, 12, 3].

3 Results

Figure 2 and Supplementary Material Figures 4-5 demonstrate that our InSpect algorithm significantly outperforms the voxelwise approach on simulated data. At all noise levels, our maps more accurately recover the ground truth than voxelwise maps. We also accurately, and automatically, recover ground truth spectral components (e.g. Figure 2 first columns) - these have to be manually identified in the standard voxelwise approach.

Figure 3 presents the joint InSpect fit to all participants' placental MR images. The four canonical spectral components (first column in top four rows) have distinct characteristics which suggests they each reflect a different microstructural environment. Although the algorithm imposes no direct anatomical analogue for any of the components, the corresponding maps identify clear anatomical structures which are consistent across control placentas, and show clear differences in dysfunctional placentas. This suggests that the canonical spectral components consistently identify distinct tissue environments, and that those tissue environments are salient to placental dysfunction.

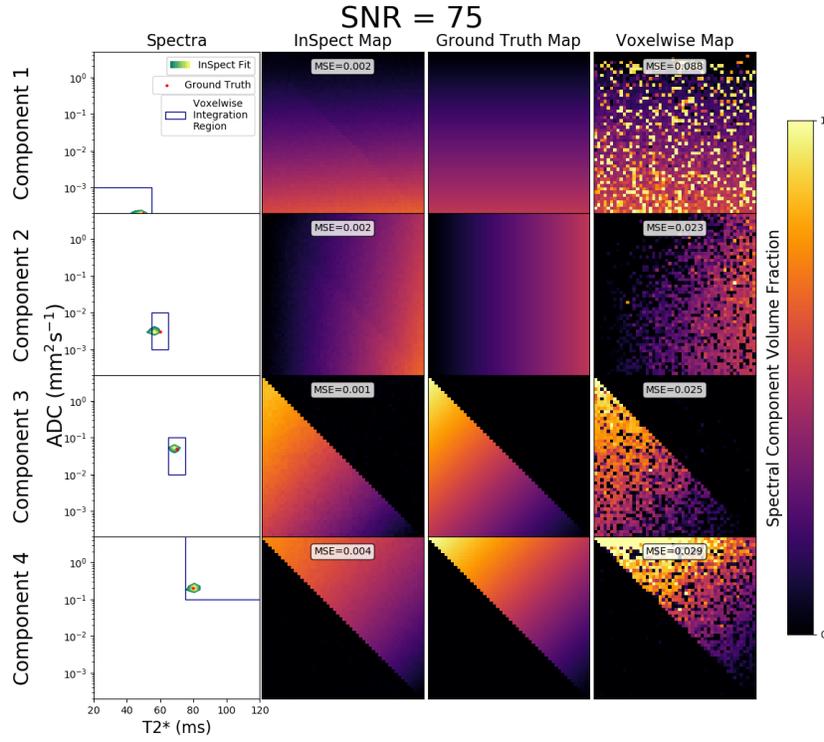


Fig. 2. InSpect algorithm applied to simulated images. The algorithm output is presented in columns 1-2: inferred spectral components (green-yellow peaks, column 1), and corresponding voxelwise weights (column 2). Column 3: ground truth voxelwise weights; ground truth spectral components (red dots in column 1) have fixed $T2^*$ -ADC values of (0.05, 0.0002), (0.06, 0.003), (0.07, 0.05) and (0.08 ms, 0.2 mm²/s). Column 4: maps obtained by numerical integration (within blue regions of column 1) of voxelwise spectral fits. Reported mean square error (MSE) values compare each component’s map and the corresponding ground truth map. See Supplementary Material for fits with other SNR values.

4 Discussion and Conclusion

We introduce and demonstrate an unsupervised machine learning method for spectroscopic imaging. Our algorithm simultaneously estimates a set of canonical spectral components and their mapping across images. This offers potential advantages over typical spectrum estimation methods - such as those utilising marginal distributions [2] or spatial regularisation [12] - which can be unstable with standard MRI noise levels and require manual spectral labelling to obtain parametric maps. Our method also has advantages over blind source separation (BSS) techniques (e.g. [13] and [19]) since we incorporate a well-defined basic MRI model, allowing us to explicitly reconstruct signal components that we

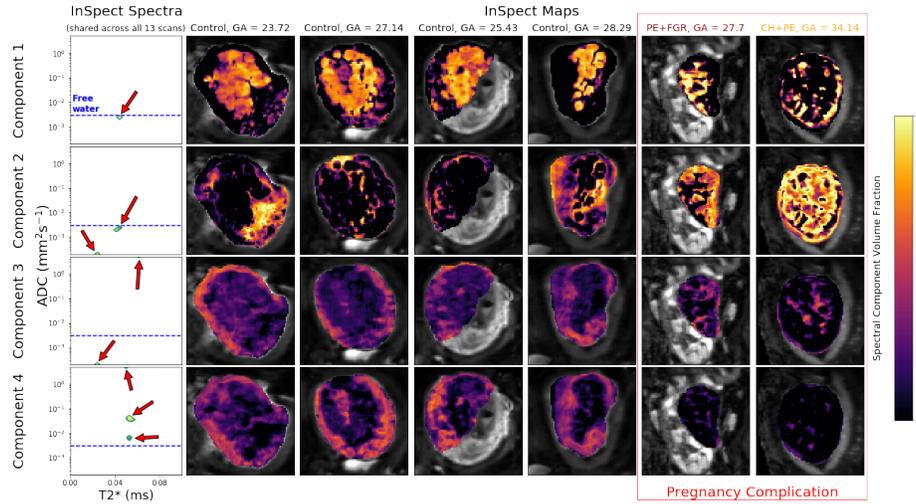


Fig. 3. Four-component InSpect model jointly fit to 13 placenta diffusion-relaxometry scans. The leftmost boxes display the InSpect spectra, which are shared across all 13 participants. The remaining boxes display the corresponding volume fraction maps for 6 of the 13 scans, with rows showing the maps for a single canonical spectral component, and columns showing the maps for a single scan. See Supplementary Material for maps from the remaining 7 scans.

can associate with distinct tissue compartments. However, BSS would be more appropriate when when explicit signal models are unknown or inaccurate.

On simulated data we significantly outperform the standard voxelwise approach, even when the total number of voxels is relatively small compared to a typical clinical scan (Figure 2). On placental diffusion-relaxometry MRI data InSpect maps clearly show anatomical structures (Figure 3). This suggests that these maps provide insight into microstructure and microcirculation across the placenta.

Given the observed spatial patterns and corresponding canonical spectral component characteristics in placental data, we can make initial speculations about the tissue microenvironments associated with each component. Component one maps out lobular structures in the placenta, and consists of a single spectral peak with ADC close to free water. These observations are consistent with the characteristics we would expect of maternal blood pools within the placenta. Component two appears to encircle these lobules, and contains a restricted (i.e. very low ADC) spectral peak. This is consistent with this component representing tissular structures, including the lobule-enclosing septa. In control participants, components three and four are prominent in the uterine wall. Both component-associated spectra contain peaks with higher ADC than free water, suggesting the presence of perfusing blood. This may be maternal blood in uter-

ine wall areas and fetal blood within the placenta. Component 4 is considerably reduced in dysfunctional placentas, likely indicative of pathology.

There are limitations that motivate potential future improvements. The choice of canonical spectral components is user-defined, in future we will explore how to select this automatically, e.g. through model selection statistics, cross validation, and/or prior microstructural knowledge. Whilst the spectral components appear to show anatomical features, we could constrain the spectra to ensure this link, or add a Markov random field to make it more likely. Here we fit our InSpect model to groups of MRI scans simultaneously, which may average over some important within-individual features. In future we will compare to individual scan fits, with focus on which approach best differentiates controls from disease.

Our algorithm is widely applicable across qMRI techniques; to facilitate this our code is freely available at <https://github.com/PaddySlator/inspect>. Applying to single-contrast qMRI is particularly attractive, since there are numerous techniques across tissue types and imaging modalities; our approach could enable significantly improved mappings compared to standard voxelwise model fitting methods. For example, the framework is immediately applicable to spectrum-based analysis of multi-echo T2 relaxometry for myelin water imaging [1]. The method can also be applied to other multidimensional experiments beyond the diffusion-relaxometry example in this paper, such as diffusion exchange spectroscopy [5, 6]. To conclude, our approach opens up spectral techniques to a wide range of situations where they are currently impossible. In particular, it paves the way to turn spectroscopic imaging into a widely used method for clinical research and practice.

References

1. Alonso-Ortiz, E., Levesque, I.R., Pike, G.B.: MRI-based myelin water imaging: A technical review. *Magnetic Resonance in Medicine* **73**(1), 70–81 (2015). <https://doi.org/10.1002/mrm.25198>
2. Benjamini, D., Basser, P.J.: Use of marginal distributions constrained optimization (MADCO) for accelerated 2D MRI relaxometry and diffusometry. *Journal of Magnetic Resonance* **271**, 40–45 (2016). <https://doi.org/10.1016/j.jmr.2016.08.004>
3. Benjamini, D., Basser, P.J.: Magnetic resonance microdynamic imaging reveals distinct tissue microenvironments. *NeuroImage* **163**, 183–196 (2017). <https://doi.org/10.1016/j.neuroimage.2017.09.033>
4. Benjamini, D., Basser, P.J.: Multidimensional correlation MRI. *NMR in Biomedicine* (jan 2020). <https://doi.org/10.1002/nbm.4226>
5. Benjamini, D., Komlosh, M.E., Basser, P.J.: Imaging Local Diffusive Dynamics Using Diffusion Exchange Spectroscopy MRI. *Physical Review Letters* **118**(15), 158003 (apr 2017). <https://doi.org/10.1103/PhysRevLett.118.158003>
6. Breen-Norris, J.O., Siow, B., Walsh, C., Hipwell, B., Hill, I., Roberts, T., Hall, M.G., Lythgoe, M.F., Ianus, A., Alexander, D.C., Walker-Samuel, S.: Measuring diffusion exchange across the cell membrane with DEXSY (Diffusion Exchange Spectroscopy). *Magnetic Resonance in Medicine* (September 2019), 1–9 (2020). <https://doi.org/10.1002/mrm.28207>

7. De Santis, S., Assaf, Y., Jeurissen, B., Jones, D.K., Roebroeck, A.: T1 relaxometry of crossing fibres in the human brain. *NeuroImage* **141**, 133–142 (nov 2016). <https://doi.org/10.1016/J.NEUROIMAGE.2016.07.037>
8. De Santis, S., Barazany, D., Jones, D.K., Assaf, Y.: Resolving relaxometry and diffusion properties within the same voxel in the presence of crossing fibres by combining inversion recovery and diffusion-weighted acquisitions. *Magnetic Resonance in Medicine* **75**(1), 372–380 (jan 2016). <https://doi.org/10.1002/mrm.25644>
9. English, A.E., Whittall, K.P., Joy, M.L., Henkelman, R.M.: Quantitative Two-Dimensional time Correlation Relaxometry. *Magnetic Resonance in Medicine* **22**(2), 425–434 (dec 1991). <https://doi.org/10.1002/mrm.1910220250>
10. Hansen, P.C.: Analysis of Discrete Ill-Posed Problems by Means of the L-Curve. *SIAM Review* **34**(4), 561–580 (1992). <https://doi.org/10.1137/1034115>
11. Hutter, J., Slator, P.J., Christiaens, D., Teixeira, R.P.A.G., Roberts, T., Jackson, L., Price, A.N., Malik, S., Hajnal, J.V.: Integrated and efficient diffusion-relaxometry using ZEBRA. *Scientific Reports* **8**(1), 15138 (dec 2018). <https://doi.org/10.1038/s41598-018-33463-2>
12. Kim, D., Doyle, E.K., Wisnowski, J.L., Kim, J.H., Haldar, J.P.: Diffusion-relaxation correlation spectroscopic imaging: A multidimensional approach for probing microstructure. *Magnetic Resonance in Medicine* **78**(6), 2236–2249 (2017). <https://doi.org/10.1002/mrm.26629>
13. Kim, D., Kim, J.H., Haldar, J.P.: Automatic Tissue Decomposition using Nonnegative Matrix Factorization for Noisy MR Magnitude Images. *ISMRM 2015* (February), 3701 (2015)
14. Kim, D., Wisnowski, J.L., Haldar, J.P.: Improved efficiency for microstructure imaging using high-dimensional MR correlation spectroscopic imaging. In: 2017 51st Asilomar Conference on Signals, Systems, and Computers. pp. 1264–1268. *IEEE* (oct 2017). <https://doi.org/10.1109/ACSSC.2017.8335555>
15. Kim, D., Wisnowski, J.L., Nguyen, C.T., Haldar, J.P.: Multidimensional correlation spectroscopic imaging of exponential decays: From theoretical principles to in vivo human applications. *NMR in Biomedicine* (May 2019), 1–19 (jan 2020). <https://doi.org/10.1002/nbm.4244>
16. Mackay, A., Whittall, K., Adler, J., Li, D., Paty, D., Graeb, D.: In vivo visualization of myelin water in brain by magnetic resonance. *Magnetic Resonance in Medicine* **31**(6), 673–677 (jun 1994). <https://doi.org/10.1002/mrm.1910310614>
17. Melbourne, A., Aughwane, R., Sokolska, M., Owen, D., Kendall, G., Flouri, D., Bainbridge, A., Atkinson, D., Deprest, J., Vercauteren, T., David, A., Ourselin, S.: Separating fetal and maternal placenta circulations using multiparametric MRI (jan 2018). <https://doi.org/10.1002/mrm.27406>
18. Menon, R.S., Allen, P.S.: Application of continuous relaxation time distributions to the fitting of data from model systems and excised tissue. *Magnetic Resonance in Medicine* **20**(2), 214–227 (aug 1991). <https://doi.org/10.1002/mrm.1910200205>
19. Molina-Romero, M., Gómez, P.A., Sperl, J.I., Czisch, M., Sämann, P.G., Jones, D.K., Menzel, M.I., Menze, B.H.: A diffusion model-free framework with echo time dependence for free-water elimination and brain tissue microstructure characterization. *Magnetic Resonance in Medicine* **80**(5), 2155–2172 (nov 2018). <https://doi.org/10.1002/mrm.27181>
20. Ning, L., Gagoski, B., Szczepankiewicz, F., Westin, C.F., Rathi, Y.: Joint RELaxation-Diffusion Imaging Moments to Probe Neurite Microstructure. *IEEE Transactions on Medical Imaging* **39**(3), 668–677 (mar 2020). <https://doi.org/10.1109/TMI.2019.2933982>

21. Pas, K., Komlosh, M.E., Perl, D.P., Basser, P.J., Benjamini, D.: Retaining information from multidimensional correlation MRI using a spectral regions of interest generator. *Scientific Reports* pp. 1–10 (2020). <https://doi.org/10.1038/s41598-020-60092-5>
22. Ronen, I., Moeller, S., Ugurbil, K., Kim, D.S.: Analysis of the distribution of diffusion coefficients in cat brain at 9.4 T using the inverse Laplace transformation. *Magnetic Resonance Imaging* **24**(1), 61–68 (jan 2006). <https://doi.org/10.1016/j.mri.2005.10.023>
23. Slator, P.J., Hutter, J., Marinescu, R.V., Palombo, M., Young, A.L., Jackson, L.H., Ho, A., Chappell, L.C., Rutherford, M., Hajnal, J.V., Alexander, D.C.: InSpect: INtegrated SPECTral Component Estimation and Mapping for Multi-contrast Microstructural MRI. In: *IPMI 2019*, pp. 755–766. Springer, Cham (jun 2019). https://doi.org/10.1007/978-3-030-20351-1_59
24. Slator, P.J., Hutter, J., Palombo, M., Jackson, L.H., Ho, A., Panagiotaki, E., Chappell, L.C., Rutherford, M.A., Hajnal, J.V., Alexander, D.C.: Combined Diffusion-Relaxometry MRI to Identify Dysfunction in the Human Placenta. *Magnetic Resonance in Medicine* (October), 1–22 (mar 2019). <https://doi.org/10.1002/mrm.27733>
25. Veraart, J., Novikov, D.S., Fieremans, E.: TE dependent Diffusion Imaging (TEdDI) distinguishes between compartmental T2 relaxation times. *NeuroImage* **182**, 360–369 (2018). <https://doi.org/10.1016/j.neuroimage.2017.09.030>