# InSpect: INtegrated SPECTral Component Estimation and Mapping for Multi-Contrast Microstructural MRI

Paddy J. Slator[1]*, Jana Hutter[2,3], Razvan V. Marinescu[1], Marco Palombo[1], Alexandra L. Young[1], Laurence H. Jackson[2,3], Alison Ho[4], Lucy C. Chappell[4], Mary Rutherford[2], Joseph V. Hajnal[2,3], and Daniel C. Alexander[1]

[1] Centre for Medical Image Computing, Department of Computer Science, University College London, 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 a novel algorithm for deriving meaningful maps from multi-contrast MRI experiments. Such experiments enable the estimation of multidimensional correlation spectra, in domains such as T1-diffusivity, T2-diffusivity, or T1-T2. These spectra combine information from complementary MR properties, and therefore have the potential for improved quantification of distinct tissue types compared to single-contrast analyses. However, spectral estimation is an ill-conditioned problem which is highly sensitive to noise and requires significant regularisation. We propose an Expectation-Maximisation based method - which we term InSpect - for unified analysis of multi-contrast MR images. The algorithm simultaneously estimates canonical spectra associated with distinct tissue types within an image, and produces maps quantifying the spatial distribution of these spectra. We test the algorithm's capabilities on simulated data, then apply to placental diffusion-relaxometry data. On placental data we identified significant within-organ and across-subject variation in T2*-ADC spectra - showing the potential of InSpect for detailed separation and quantification of distinct microstructural environments.

## 1 Introduction

Multidimensional magnetic resonance (MR) experiments simultaneously measure multiple MR properties, and hence promise more specific characterisation of tissue. Examples of multidimensional MR techniques include various types of correlation spectroscopy, such as relaxometry-relaxometry [6] and diffusion-relaxometry [12]. Several recent papers have leveraged recent advances in scanner hardware to extend these ideas into imaging, in the T1-diffusion [4], T2-diffusion [13,9], and T1-T2-diffusion [2] domains.

An attractive approach to analysing such data is continuum modelling. This assumes that spins have a distribution of values (e.g. relaxivity, diffusivity), which are quantified by a multidimensional spectrum. The exponential dependence on relaxation constants and diffusivity leads to a Laplace transform model on the MR signal; the spectrum can therefore be estimated using an inverse Laplace transform with regularisation [6].

However, fitting the spectrum requires high signal-to-noise ratio (SNR) data, which can make individual voxel fits - and hence the derivation of quantitative maps of spectral variation across an image - particularly problematic. In practice, this means that additional spatial regularisation is often necessary. This can mean averaging the signal over a region of interest (ROI), or also "spectrally integrating" within user-defined regions. This usually involves identifying canonical components in ROI-derived spectra, then integrating voxelwise spectra within the regions corresponding to these components, hence obtaining apparent spectral volume fraction estimates [10,9,2].

Recently, methods have been proposed for increasing robustness of voxelwise spectral fits, utilising marginal distributions [1] or spatial regularisation [9]. These methods can improve subsequent spatial mapping, yet there are inherent limitations to existing approaches which motivate this paper. Specifically, canonical spectral components require manual identification, and therefore may be ill-defined if estimated over inhomogeneous regions, and may not cover the full range of observations over the extent of an image.

In this paper we present a method - named InSpect - which addresses these problems in a data-driven way. Our algorithm automatically segments multidimensional MR images by clustering voxels with similar spectra, and simultaneously infers representative spectra for these clusters. This offers many potential advantages over voxelwise approaches. By averaging over similar voxels we reduce noise in canonical spectrum estimates. Additionally, the method is fully data-driven so is unlikely to miss any important spectral components that appear in the data. In short, the InSpect algorithm seeks a compact representation of the whole image that captures intrinsic variation in the data without overfitting.

The paper proceeds as follows: we first define the InSpect model, then derive an Expectation-Maximisation algorithm for its inference. We finally test the algorithm's utility on simulated and real multidimensional MRI data. We present a general form of InSpect in two dimensions, but emphasise that extension to higher dimensions is simple.

## 2   Methods

### 2.1   Multidimensional Spectrum Estimation

The standard approach for estimating the spectrum from a multidimensional MRI experiment proceeds as follows [6]; for a general 2D multi-contrast MRI experiment the signal, $S$, can be described by a continuum model

$$S(t_1, t_2) = \int \int F(\omega_1, \omega_2) K(t_1, t_2, \omega_1, \omega_2) \, d\omega_1 d\omega_2 \tag{1}$$

where $t_1$ and $t_2$ are experimental parameters which are varied to yield contrast in intrinsic MR properties $\omega_1$ and $\omega_2$, via the specific form of the kernel $K(t_1, t_2, \omega_1, \omega_2)$. $F(\omega_1, \omega_2)$ is the 2D spectrum of $\omega_1$ and $\omega_2$, i.e. the distribution of these values across all spins. Discretising onto a $N_{\omega_1}$ by $N_{\omega_2}$ grid yields

$$S(t_1, t_2) = \sum_{i=1}^{N_{\omega_1}} \sum_{j=1}^{N_{\omega_2}} F(\omega_1^{(i)}, \omega_2^{(j)}) K(t_1, t_2, \omega_1^{(i)}, \omega_2^{(j)}) \tag{2}$$

By choosing a suitable ordering of spectrum coordinates $\omega_1, \omega_2$, the signal for all MR encodings in the experiment can be written in matrix form

$$\mathbf{S} = \mathbf{KF} \tag{3}$$

where $\mathbf{S}$ is a column vector, length $N_s$ of the signals at each encoding, $K$ is an $N_s$ by $N_{\omega_1} N_{\omega_2}$ matrix of discretised kernel values, and $F$ is an $N_{\omega_1} N_{\omega_2}$ length column vector of spectrum values. The spectrum $F$ can then be calculated as follows, including a non-negativity constraint and regularisation term

$$\mathbf{F} = \underset{\mathbf{F} \geq 0}{\arg\min} \, \|\mathbf{KF} - \mathbf{S}\|_2^2 + \alpha \|\mathbf{F}\|_2^2 \,. \tag{4}$$

By solving the above equation with non-negative least squares regression the spectrum can be estimated within a single voxel. However, as mentioned earlier, low SNR can lead to noisy spectrum estimates and hence poor spatial maps. This is the case whether these maps are produced directly from the estimated spectra, or through picking canonical spectral regions and integrating the fitted voxelwise spectra within them. In the following section we describe our novel approach to this problem.

## 2.2   InSpect model

We move from considering the signal from a single voxel, to an image (or volume) consisting of $N$ voxels in total. We assume that the signal from each voxel, $\mathbf{S}_n$, is described by the continuum model of Equation (1) as described in the previous section.

Rather than naively fitting spectra to each $\mathbf{S}_n$, we seek a data-driven lower dimensional representation of the spectral image. Thus, we assume a small number of voxel types each defined by a canonical spectrum. We then seek to simultaneously estimate the set of canonical spectra and the assignment of voxels to spectra. This has the effect of grouping voxels into clusters, so a byproduct is a segmentation of the image into distinct regions based on spectral properties.

We start from the assumption that there are $M$ clusters of voxels, with each cluster having an associated spectrum $\mathbf{F}_m$, which we estimate from the data. In practice, we do not know each voxel's cluster a-priori and also need to estimate it. These cluster membership indices are hence model latent states, which we write as $z_n = m$ if voxel $n$ belongs to cluster $m$. The set of latent states for the whole image is therefore $\mathbf{z} = \{z_n\}_{n=1}^N$, where $z_n \in \{1, ..., M\}$. The cluster membership

probabilities are modelled with a categorical distribution, the parameters are the number of clusters, $M$, and the cluster probabilities $\{p_m\}_{m=1}^{M}$.

We consider a Gaussian model on the observed signal. For a voxel, $n$ belonging to cluster $m$ the expected value for each element of the observed signal vector, $S_n$, is the corresponding $\mathbf{KF}_m$ term. We also assume that all observations in a voxel have the same variance $\sigma_n^2$. We write this as follows

$$\mathbf{S}_n \sim \mathcal{N}(\mathbf{KF}_m, \sigma_n^2). \tag{5}$$

Assuming that voxels are independent, the complete data model likelihood for an image given the parameters $\theta$ is therefore

$$\pi(\mathbf{D}, \mathbf{z}|\theta) = \pi(\mathbf{z}|\theta)\pi(\mathbf{D}|\mathbf{z}, \theta) = \prod_{n=1}^{N} p_{z_n} N(\mathbf{S}_n; \mathbf{KF}_{z_n}, \sigma_n^2) \tag{6}$$

where $N(\mathbf{S}_n; \mathbf{KF}_{z_n}, \sigma_n^2)$ refers to the product over the Normal PDFs of each measured signal value within the voxel, and we have denoted $\mathbf{D} = \{\mathbf{S}_n\}_{n=1}^{N}$ for notational simplicity. By summing over all possible clusterings we get the marginal likelihood

$$\pi(\mathbf{D}|\theta) = \sum_{\text{all } \mathbf{z}} \pi(\mathbf{D}, \mathbf{z}|\theta) = \sum_{m=1}^{M} \prod_{n=1}^{N} p_{z_n} N(\mathbf{S}_n; \mathbf{KF}_{z_n}, \sigma_n^2) \tag{7}$$

The maximum likelihood estimate (MLE) of $\theta$ is the value that maximises this. This calculation is intractable in practice, so we derive an Expectation-Maximisation (EM) algorithm [5]. The full set of model parameters is $\theta = \{p_1, ..., p_M, \mathbf{F}_1, ..., \mathbf{F}_M, \sigma_1^2, ..., \sigma_N^2\}$, but in practice we estimate the $\sigma_n^2$'s by calculating the empirical variance of the observed data at each voxel, so do not consider this in the EM algorithm. We also need to choose the number of clusters $M$ - ideally from the data. We do this by fitting the model for a range of reasonable $M$, then comparing model selection statistics - such as the Bayesian information criterion (BIC) and Akaike information criterion (AIC) - across this range.

## 2.3   Expectation-Maximisation algorithm

To implement an EM algorithm we first calculate an expression for the expected value of $\log \pi(\mathbf{D}, \mathbf{z}|\theta)$ (complete data log-likelihood) with respect to $\pi(\mathbf{z}|\mathbf{D})$ (posterior probability of latent states given the data and current parameters). Using the notation of Bishop [3] we write this function as

$$Q(\theta, \theta^{(t-1)}) = \mathbb{E}_{\mathbf{z}|\mathbf{D}, \theta^{(t-1)}} \left[\log \pi(\mathbf{D}, \mathbf{z}|\theta)\right] = \sum_{\mathbf{z}} \pi(\mathbf{z}|\mathbf{D}, \theta^{(t-1)}) \log \pi(\mathbf{D}, \mathbf{z}|\theta) \tag{8}$$

where $\theta^{(t)}$ denotes the model parameters at step $t$ of the algorithm. Marginalising over the clusters, applying the independence of voxels, substituting the likelihood

(Equation (6)), and taking the log gives the final expression for $Q$

$$Q(\theta, \theta^{(t-1)}) = \sum_{n=1}^{N} \sum_{m=1}^{M} w_{nm} \left[ \log p_m - \frac{N_s}{2} \log 2\pi\sigma_n^2 - \left\| \frac{\mathbf{S}_n - \mathbf{KF}_m}{\sqrt{2\sigma_n^2}} \right\|_2^2 \right] \qquad (9)$$

where $w_{nm} = \pi(z_n = m|\mathbf{S}_n, \theta^{(t-1)})$. These $w_{nm}$ terms are calculated (for all $n$ and $m$) in the E-step, and the M-step maximises $Q(\theta, \theta^{(t-1)})$ with respect to the parameters $\theta$.

**E-step** For the E-step we first note the full posterior distribution for the model

$$\pi(\theta, \mathbf{z}|\mathbf{D}) \propto \pi(\mathbf{D}|\theta, \mathbf{z})\pi(\mathbf{z}, \theta) = \prod_{n=1}^{N} p_{z_n} N(\mathbf{S}_n; \mathbf{KF}_{z_n}, \sigma_n^2) \qquad (10)$$

where we have assumed a uniform prior on both the latent states $\mathbf{z}$, and the parameters $\theta$. The posterior distribution for each $z_n$ - normalised across all clusters - is therefore

$$w_{nm} = \pi(z_n = m|\mathbf{S}_n, \theta^{(t-1)}) = \frac{p_m N(\mathbf{S}_n; \mathbf{KF}_m, \sigma_n^2)}{\sum_{i=1}^{M} p_i N(\mathbf{S}_n; \mathbf{KF}_i, \sigma_n^2)}. \qquad (11)$$

**M-step** In the M-step we optimise the parameters given the current cluster weights $w_{nm}$. In other words we solve

$$\theta = \arg\max_{\theta} \ Q(\theta, \theta^{(t-1)}) \qquad (12)$$

sequentially for each parameter $\theta = \{p_1, ..., p_M, \mathbf{F}_1, ..., \mathbf{F}_M\}$. For $p_1, ...p_M$ maximising this equation, whilst implementing the constraint $\sum_{m=1}^{M} p_m = 1$ using a Lagrange multiplier (e.g. Bishop [3]), gives

$$p_m = \frac{1}{N} \sum_{n=1}^{N} w_{nm} \qquad (13)$$

in other words the current mean cluster weight over all voxels.

For the canonical cluster-associated spectra, $\mathbf{F}_1, ...\mathbf{F}_M$, we have (ignoring constant terms, and swapping $\arg\max_{\mathbf{F}_m}$ for $\arg\min_{\mathbf{F}_m}$ due to the sign)

$$\mathbf{F}_m = \arg\min_{\mathbf{F}_m} \sum_{n=1}^{N} w_{nm} \left\| \frac{\mathbf{S}_n - \mathbf{KF}_m}{\sqrt{2}\sigma_n} \right\|_2^2 \qquad (14)$$

Taking the derivative with respect to $\mathbf{F}_m$, setting equal to zero and rearranging gives

$$\mathbf{KF}_m = \frac{1}{W_m} \sum_{n=1}^{N} w_{nm} \mathbf{S}_n \qquad (15)$$

where $W_m = \sum_{n=1}^{N} w_{nm}$. In other words, we need to minimise the difference between $\mathbf{KF}_m$ and the mean signal over all voxels normalised by cluster weights. This implies the following modification to Equation (4) to calculate the spectrum associated with each cluster

$$\mathbf{F}_m = \underset{\mathbf{F}_m > 0}{\arg\min} \; \left\| \mathbf{KF}_m - \frac{1}{W_m} \sum_{n=1}^{N} w_{nm} \mathbf{S}_n \right\|_2^2 + \alpha \|\mathbf{F}_m\|_2^2 \qquad (16)$$

which we can solve with non-negative least squares regression as described earlier. By iterating E and M steps we therefore calculate the MLE of model parameters (cluster probabilities and cluster-associated spectra for each cluster), and posterior distribution of latent states (cluster weights for each voxel).

## 2.4   Application to Placenta Diffusion-Relaxometry Data

We applied our algorithm to data from combined diffusion-relaxometry experiments previously published by Slator *et al.* [11]. In this work the authors varied the b-value and echo time (TE) using a sequence called ZEBRA [8], yielding simultaneous diffusivity and T2* contrast. The sequence consists of 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 = 300×320×84 mm, TR = 7 s, SENSE = 2.5, halfscan = 0.6, resolution = 3mm$^3$.

   We fit InSpect to placental scans from 12 participants, 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 two times, four weeks apart, and developed superimposed pre-eclampsia by the second scan. A placenta and uterine wall region of interest (ROI) was manually segmented on all images.

   There are multiple approaches possible when applying InSpect to a dataset with multiple participants: fit to each image independently and find separate canonical spectra and clusters for each; fit to all images simultaneously and estimate a common set of spectra; or any number of hybrid approaches. The best approach will depend on the specific application. One important consideration is the extent to which one wants to probe within-image heterogeneity, as opposed to across-image differences. In this paper our aim is to gain an initial idea of typical placental T2*-ADC spectra, and their spatial distributions, across healthy and unhealthy tissue. We therefore fit InSpect in two ways: first to each participant's scan individually for various values of $M$, to determine a parsimonious number of clusters from the data. Then, given this information, we fix the number of clusters and fit InSpect to the data from all patients simultaneously.

   The number of clusters which best explained the data in individual image fits, as measured by the BIC and AIC, varied across participants. For the placentas of participants diagnosed with a pregnancy complication the typical number of clusters was three. For healthy placentas we tended to see BIC and AIC

values levelling off at around eight clusters. Given this information, we fixed the number of clusters at eight when fitting InSpect to the data from all participants simultaneously.

We also also naively fit voxelwise T2*-ADC spectra to all scans by solving Equation (4), and derived spectral volume fraction maps from these by integrating in six regions (defined later) of T2*-ADC space, mirroring the approach in references [10,9,2].

### 2.5   Application to Simulated Data

We performed simulations, using the same contrast encodings (i.e. b-value, TE pairs) as the placental data, to test InSpect. We constructed a minimal example by first creating a synthetic image of 1000 voxels, with the voxels split evenly between two clusters. Each cluster was next associated a simple one-component spectrum with a fixed T2* and ADC value. Given a voxel's T2* and ADC we simulated the signal with a simple joint model as follows

$$S(b, T_E) = \exp\left(-b\mathrm{ADC}\right)\exp\left(-T_E/T_2^*\right) \tag{17}$$

adding Gaussian measurement noise with a realistic SNR level of 20 (e.g. [11]) in all voxels. We ran a total of 72 simulations by varying the properties of cluster-associated spectra (details in Figure 1 caption). This allowed us to investigate how different cluster-associated spectra need to be for InSpect to distinguish them, and if this varies across regions of T2*-ADC space.

We first fit spectra in the standard voxelwise manner (using Equation (4)) for all simulated images. We subsequently derived spectral volume fraction maps from the voxelwise fits, following the aforementioned integration approach (e.g. [10,9,2]). We defined the areas in which to spectrally integrate as two regions separated by the midpoint of the two cluster-associated spectra. We fit InSpect to all synthetic images separately with the number of clusters set to two, thereby automatically obtaining segmentations and canonical spectrum estimates.

### 2.6   Algorithm Implementation Details

For the voxelwise fits we set the value of $\alpha$ in Equation (4) at 0.01 using the L-curve method [7]. We used the same $\alpha$ value in Equation (16) for the InSpect fits. For all InSpect fits we initialised the cluster weights by sampling from a uniform random distribution. We estimate $\sigma^2$ voxelwise by calculating the empirical variance of the b=0 volumes with the lowest TE. We determined convergence of the EM algorithm by manually checking for stationarity of parameters, 25 EM steps were typically sufficient for convergence. The EM algorithm output is robust over multiple runs from different random initialisations.

## 3   Results

Figure 1 shows that InSpect performs significantly better than voxelwise mapping by spectral integration; recovering the ground truth clusters in the majority

of simulations. These results are not a complete characterisation of the ability of InSpect, but they do show that it works for various parameter combinations across the T2*-ADC domain under the same contrast encodings as the placental data. The limits of InSpect - under this acquisition sequence - were reached when T2* was low, especially for components with close ADC values. This reflects the fact that most signal has attenuated for these T2* values at the acquired TEs.

Figure 2 shows the InSpect fit to the placental data - with all 13 scans fit at once. The InSpect maps show clear structure - both within organ and across participants. The corresponding eight cluster-associated spectra show varying number of peaks, T2*, and ADC values.

We also compare InSpect mapping to spectral integration of voxelwise fits (Figure 3). The voxelwise maps show similar features to InSpect maps, although they average over some small-scale features. Visualisation of these substructures may be possible from voxelwise fits given a sufficiently delicate partitioning of the spectral domain.

## 4   Discussion

We demonstrate InSpect, a data-driven method for mapping spectral components in multi-contrast microstructural MRI experiments. This is an alternative to mapping by spectral integration of voxelwise fits, which is inherently unstable because of high noise, and fails to exploit the likely low variation in underlying spectra across individual (or groups of) images. InSpect automatically maps the spectral components, whereas voxelwise integration requires user-defined choices of regions in the spectral domain. Another advantage of InSpect over voxelwise approaches is speed; the method avoids computationally heavy non-negative least squares fitting in every voxel.

On simulated data we show that InSpect significantly improves mapping based on spectral properties, even when the total number of voxels is relatively small compared to a typical clinical scan (Figure 1). On placental diffusion-relaxometry MRI data InSpect segmented clear anatomically-linked structures (Figure 2A). This suggests that InSpect maps provide insight into microstructure and microcirculation across the placenta. Clusters 8 and 6 likely represent the centre and periphery of placental lobules respectively. The spectra associated with these clusters both have very high T2*, potentially reflecting the high oxygen levels in healthy placentas. We also observed very clear differences in cluster assignments between scans of participants who were diagnosed with a pregnancy disorder compared to controls (Figure 2C).

In fact, we observed a canonical spectrum type which appears solely in pathological cases (cluster 1 in Figure 2). This demonstrates a significant advantage when fitting the model to the full data set - as we did for the placental data - as opposed to individual images. By sharing data across subjects we automatically quantify spectral differences between control and dysfunctional placentas. It also ensures consistent interpretation across all images in a data set, avoiding problems such as the matching of clusters across images. However, this comes at
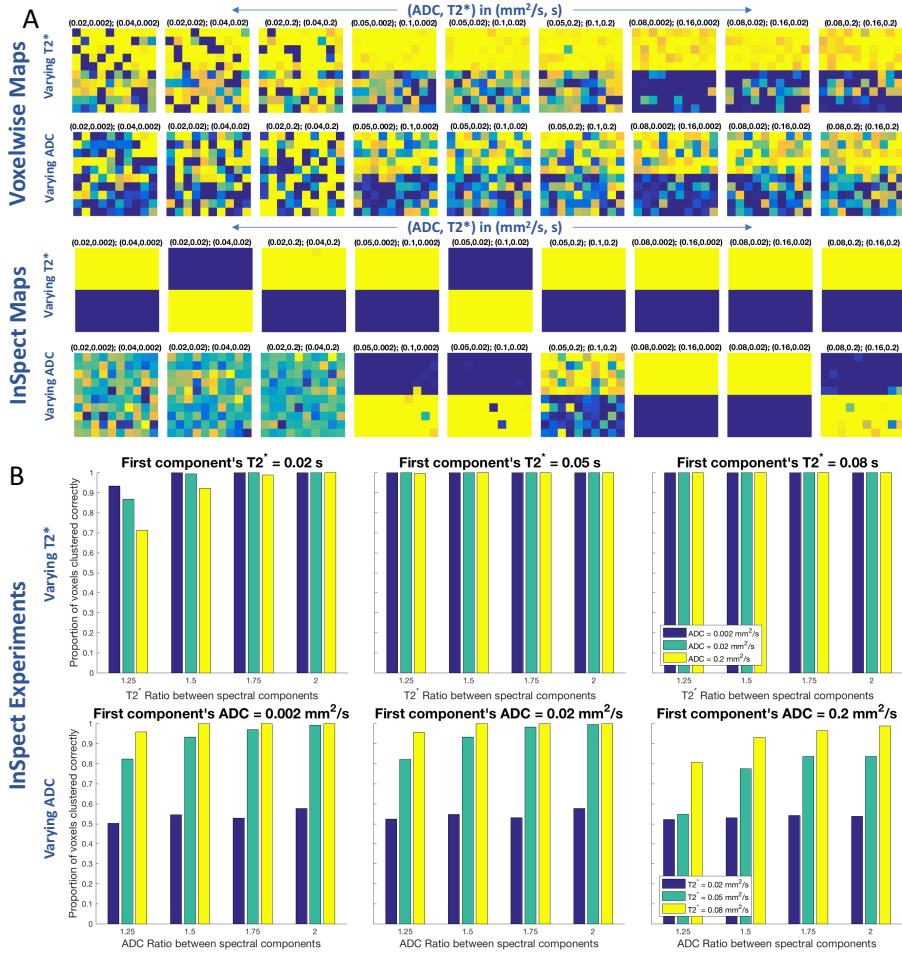
**Fig. 1.** InSpect applied to synthetic diffusion-relaxometry images. Each image comprises two clusters - located in the top and bottom halves of the image. Each cluster is associated with a one component T2*-ADC spectrum. Three quantities were varied across simulations: distance between the two cluster-associated spectral components (either in T2* space with ADC fixed, or vice versa, we used 4 distances between 1.25 and 4), value of the fixed parameter (3 in total), and value of the varying parameter (3 in total). There were hence a total of 72 simulations. SNR was set to 20 for all images. A) Volume fraction maps obtained from spectral integration of voxelwise fits (e.g. [10,9,2]) in two domains delineated by the midpoint of the simulation components, and InSpect segmentation maps (posterior weights on clusters in each voxel). Note that in some cases where InSpect exactly infers the simulated ground truth the cluster labels are swapped. B) Bar plots showing the proportion of voxels where InSpect successfully recovered the ground truth for various combinations of T2* and ADC characterising the two cluster-associated spectra.
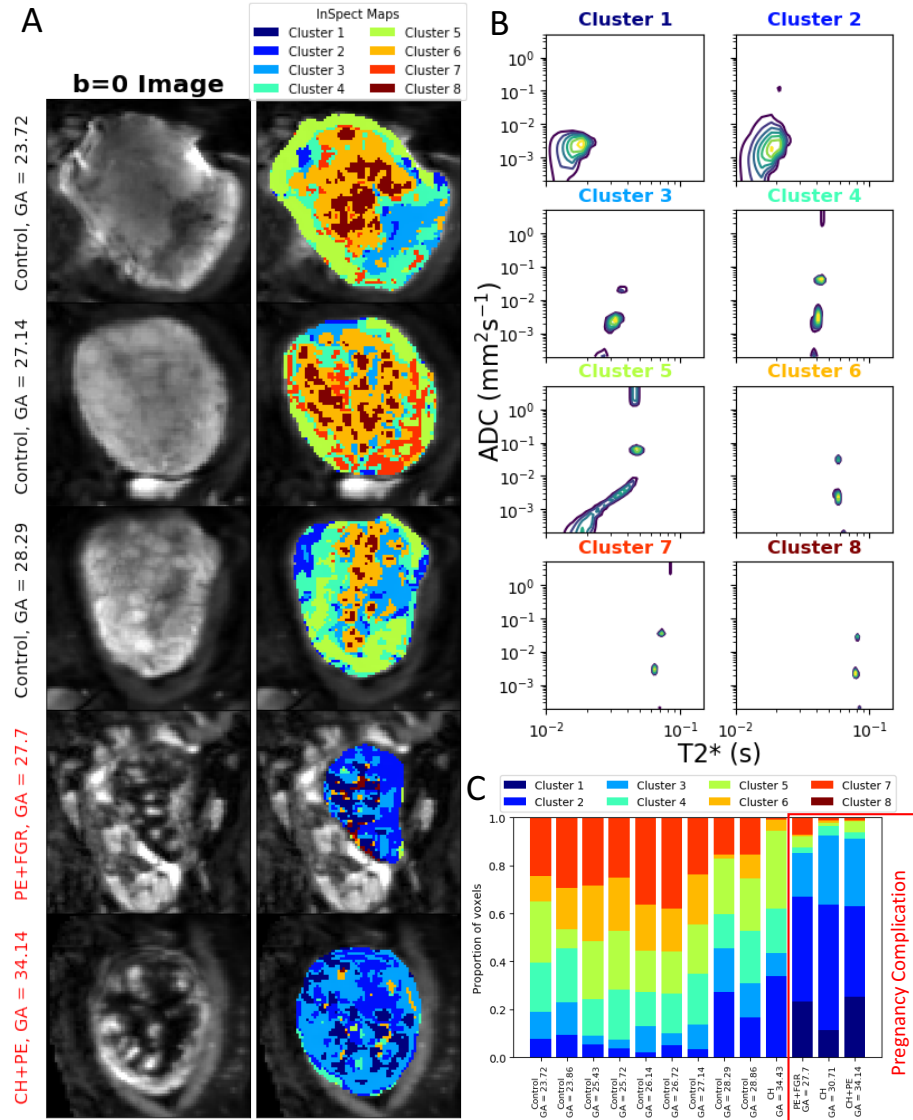
**Fig. 2.** InSpect simultaneously applied to 13 placental diffusion-relaxometry images. A) InSpect maps and b=0 images for five of the 13 participants, including two who had been diagnosed with a pregnancy complication at the time of the scan. B) Corresponding T2*-ADC spectra for the eight clusters in A). C) Proportion of voxels assigned to each cluster for all participants.
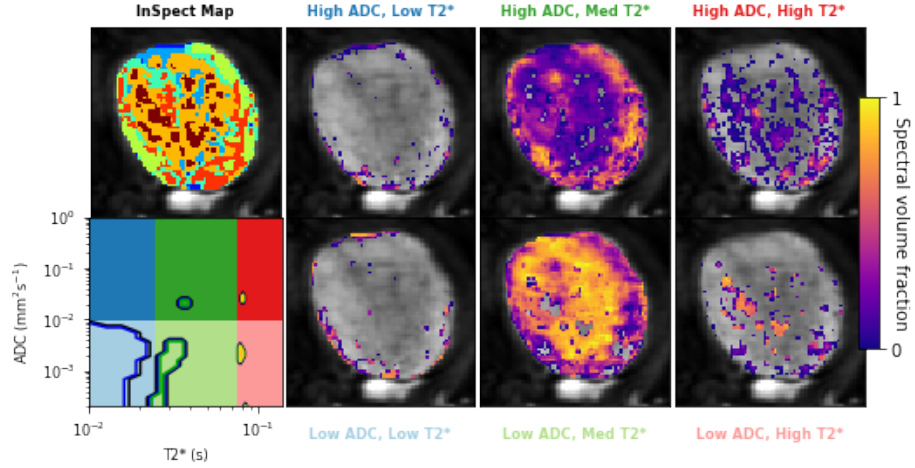
**Fig. 3.** Comparison between InSpect mapping (top left panel) and and spectral integration of voxelwise fits for a single participant. Voxelwise maps (all but the top left panel) were obtained by the standard spectral integration approach (e.g. [10,9,2]) within the six regions displayed in the bottom left panel. The bottom left panel also outlines three inferred InSpect spectra (clusters 1, 5 and 8 in Figure 2), each in a different color, for comparison.

the expense of averaging inter-organ heterogeneity - within-participant fits may make sense in other contexts.

An advantage of the voxelwise spectral integration approach over InSpect is that it implicitly considers partial volume effects. This observation implies a clear direction for future work - an extension of the discrete clustering approach we propose here to continuous spectral component mapping. This would have continuous latent states, rather than discrete, and therefore would better quantify partial voluming. These developments should work towards the ultimate aim of using multi-contrast MRI experiments to improve tissue microstructure measurements; enabling the quantification of tissue properties - both healthy and diseased - which cannot usually be separated by a single MR contrast. These methods have wide-ranging potential applications: spectroscopic imaging techniques are prominent in the brain and other organs; there many parameters other than T2* and ADC that could be exploited; and extending to higher dimensional spectra is straightforward.

## 5   Conclusion

We introduce the InSpect algorithm, which simultaneously undertakes two complementary tasks - multidimensional spectrum estimation and mapping - by sharing information across voxels. Although we applied InSpect to diffusion-

relaxometry data, the method is immediately applicable to any multidimensional MR imaging experiment.

# References

1. 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)
2. Benjamini, D., Basser, P.J.: Magnetic resonance microdynamic imaging reveals distinct tissue microenvironments. NeuroImage **163**, 183–196 (2017)
3. Bishop, C.M.: Pattern Recognition and Machine Learning (2006)
4. 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 (2016)
5. Dempster, A.P., Laird, N., D.B. Rubin, Rubin, D.: Maximum Likelihood from Incomplete Data via the EM Algorithm. Journal of the Royal Statistical Society. Series B (Methodological) **39**(1), 1–38 (1977)
6. 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 (1991)
7. Hansen, P.C.: Analysis of Discrete Ill-Posed Problems by Means of the L-Curve. SIAM Review **34**(4), 561–580 (1992)
8. 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 (2018)
9. 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)
10. 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 (1994)
11. 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 (2019)
12. Van Dusschoten, D., Moonen, C.T., Adrie De Jager, P., Van As, H.: Unraveling diffusion constants in biological tissue by combining Carr- Purcell-Meiboom-Gill imaging and pulsed field gradient NMR. Magnetic Resonance in Medicine **36**(6), 907–913 (1996)
13. Veraart, J., Novikov, D.S., Fieremans, E.: TE dependent Diffusion Imaging (TEdDI) distinguishes between compartmental T2 relaxation times. NeuroImage **182**, 360–369 (2018)