



Motion Correction in Low SNR MRI Using an Approximate Rician Log-Likelihood

Ivor J. A. Simpson¹(✉), Balázs Örszík², Neil Harrison³, Iris Asllani^{2,4},
and Mara Cercignani³

¹ Department of Informatics, University of Sussex, Brighton, UK
i.simpson@sussex.ac.uk

² CISC, Brighton and Sussex Medical School, Brighton, UK

³ CUBRIC, University of Cardiff, Cardiff, UK

⁴ Biomedical Engineering, Rochester Institute of Technology, Rochester, USA

Abstract. Certain MRI acquisitions, such as Sodium imaging, produce data with very low signal-to-noise ratio (SNR). One approach to improve SNR is to acquire several images, each of which takes may take more than a minute, and then average these measurements. A consequence of such a lengthy acquisition procedure is subject motion between each image. This work investigates a solution for retrospective motion correction in this scenario, where the high level of Rician noise renders standard registration tools less effective. We employ a simple generative model for the data based on tissue segmentation maps, and provide a differentiable approximation of the Rician log-likelihood to fit the model to the observations. We find that this approach substantially outperforms a Gaussian log-likelihood baseline on synthetic data that has been corrupted by Rician noise of varying degrees. We also provide results of our approach on real Sodium MRI data, and demonstrate that we can reduce the effects of substantial motion compared to a general purpose registration tool.

Keywords: Motion correction · Rician distribution · Low SNR

1 Introduction

Subject motion is a common issue in long MRI acquisition protocols; in situations where several images have been acquired, motion can be retrospectively corrected using image registration. For brain MRI images with reasonable signal-to-noise-ratio (SNR), general purpose linear image registration tools, e.g. [2, 3, 8, 10, 13], have been shown to be highly effective. However, in low SNR MRI data, such as acquired with Sodium MRI, traditional cost functions may become less effective. One cause is the noise properties of the analysed data, which consists of the magnitude of the complex signal components. The noise in such data is described using a Rician distribution [6]. When the SNR of the acquired complex signal is high, the resulting noise is approximately Gaussian. Conversely, when the SNR




Fig. 1. Illustration of the generative model, which predicts noise free images, \hat{y} , parameterised by: T1 estimated tissue segmentation maps, G , multiplied by estimated tissue intensities, \mathbf{x} . These are transformed by a translation, \mathbf{t} and rotation θ . The error between the observations, \mathbf{y} and predictions is described using a Rician likelihood, which is used to drive the parameter estimation.

is low the Rician distribution is asymmetric and dissimilar from a Gaussian. This distinction is particularly significant for registration approaches considering cost functions derived from a Gaussian, e.g. sum-of-squared differences.

This paper introduces a linear motion correction model using a simple generative model of the data. This is inspired by the seminal ‘‘Unified Segmentation’’ paper [1]. A diagram of our approach is given in Fig. 1. Our model produces noise-free predictions, which are rigidly aligned to each of the observed images. The novel contribution of this work lies in our approximation of the Rician log-likelihood that enables gradient estimates through automatic differentiation [14]. This is in contrast to previous work using Rician likelihoods for motion correction [16], which required a gradient-free optimisation of the transformations.

We demonstrate how our approach can be used to remove substantial motion from Sodium MRI data. Sodium is an emerging imaging modality, with several potential biomedical applications [11, 19]. However, it has poor SNR due to the relatively low concentration and magnetic susceptibility of Sodium, as shown in Fig. 1. Our results illustrate the effectiveness of this approach in removing substantial motion from high noise situations in both real and synthetic datasets.

2 Background: The Rice Distribution

The noise in magnitude MR images is known to follow a Rice distribution [6]:

$$p(\mathbf{y}|\hat{\mathbf{y}}, \sigma) = \text{Rice}(\mathbf{y}; \hat{\mathbf{y}}, \sigma) = \frac{\mathbf{y}}{\sigma^2} \exp\left(\frac{-(\mathbf{y}^2 + \hat{\mathbf{y}}^2)}{2\sigma^2}\right) I_0\left(\frac{\mathbf{y}\hat{\mathbf{y}}}{\sigma^2}\right) \quad (1)$$

where I_0 is a modified Bessel function of the first kind with order zero (described in Sect. 3.2). Unlike the Gaussian, this distribution: is not symmetric with respect to its first parameter, $\hat{\mathbf{y}}$; does not fulfill any of the algebraic conjugacy properties that enable derivation of closed-form parameter updates; it also does not provide an obvious cost function for directly comparing two images, as it requires a parameterisation in terms of the clean signal, $\hat{\mathbf{y}}$. Generative models can be used to provide such a parameterisation [1].

3 Method

We consider a generative model for the image data based on 5 probabilistic tissue segmentation maps, G , derived from a T1 image acquired in the same space. We denote G as a matrix of size $N \times 5$, where N corresponds to the number of voxels. The intensity of any voxel can be predicted by matrix multiplication with \mathbf{x} , a vector containing the intensity for each tissue class. We consider a geometric transformation associated with each observed image:

$$\hat{\mathbf{y}}_i = P(T(G\mathbf{x}, \mathbf{t}_i, \boldsymbol{\theta}_i)) \quad (2)$$

where T provides a rigid transformation of $G\mathbf{x}$, according to translation \mathbf{t} and rotation parameters given by $\boldsymbol{\theta}$. We also include a convolution, P , which corresponds to the point-spread function of the acquisition sequence; this is estimated a-priori from the sequence reconstruction method [18]. The predictions $\hat{\mathbf{y}}_i$ can now be fit to the observed data \mathbf{y}_i using an appropriate likelihood function.

3.1 Priors

In this problem, we are considering the registration of noisy data. Accordingly, the model requires the specification of prior knowledge to enable robust inference. We choose a physiologically based Gaussian prior over the concentration of Sodium, measured in mM, for different tissue types:

$$p(\mathbf{x}) = \mathcal{N}([40, 30, 140, 50, 50], [4, 4, 6, 10, 10]^2) \quad (3)$$

where the means are from [11] and the standard deviations are empirically selected.

The translations have a Normal prior, with a standard deviation specified in mm. The rotations, which are described through an axis-angle representation (in Radians), also employ a Normal prior distribution:

$$\begin{aligned} p(\mathbf{t}_i) &= \mathcal{N}(0, 1.25^2) \\ p(\boldsymbol{\theta}_i) &= \mathcal{N}(0, 0.025^2) \end{aligned}$$

3.2 A Stable Approximation of the Rician Log-Likelihood

Most of the Rician likelihood (Eq. 1) is amenable to efficient calculation in a differentiable manner. However, I_0 , corresponds to a modified Bessel function of the first kind with order zero [20], which is an infinite series:

$$I_0(z) = \sum_{k=0}^{\infty} \frac{(\frac{1}{4}z^2)^k}{(k!)^2} \quad (4)$$

The result can be approximated as a sum of the first N_k terms. However, this necessitates a differentiable form for the factorial in the denominator. By noting both that $k! = \Gamma(k+1)$, where Γ is the Gamma function, and that we only require the log probability, we can write an approximation for $\log I_0(z)$ as:

$$\log I_0(z) \approx \log\text{-sum-exp}(\mathbf{k}(\log(0.25) + 2 * \log(z)) - 2 \ln \Gamma(\mathbf{k} + 1)) \quad (5)$$

where $\ln \Gamma$ refers to the log Gamma function. \mathbf{k} is a vector containing values from 0 to N_k , which is summed over. $\log\text{-sum-exp}(\mathbf{z})$ is a numerically stable and convex function [5] for calculating the logarithm of the sum of exponentiated terms, $\log\text{-sum-exp}(\mathbf{z}) = \log(\sum_i \exp(z_i))$. This implementation is empirically numerically stable, although inefficient in terms of memory as we require multiplying each voxel by N_k values. We found that $N_k = 50$ provided sufficient precision.

3.3 Inference

We perform maximum-a-posteriori (MAP) inference on the model parameters $\Theta = \{\mathbf{x}, \mathbf{t}, \boldsymbol{\theta}, \sigma\}$, with the following cost function:

$$\mathcal{L} = - \sum_i^N [\log p(\mathbf{y}_i | \mathbf{x}, \mathbf{t}_i, \boldsymbol{\theta}_i, \sigma) + \log p(\mathbf{t}_i) + \log p(\boldsymbol{\theta}_i)] + \log p(\mathbf{x}) \quad (6)$$

Updates alternated between two groups of parameters, those that are shared for all images $\Theta_1 = \{\mathbf{x}, \sigma\}$ and those that vary per image $\Theta_2 = \{\mathbf{t}, \boldsymbol{\theta}\}$. The updates for Θ_1 were calculated using batches of 5 images at a time, and Θ_2 were updated per image. To account for the batching in updating Θ_1 , we perform two update steps on these parameters for every step for Θ_2 . The Adam [9] optimiser was used to optimise the model parameters, with a fixed learning rate of $2e^{-2}$ for Θ_1 and $1e^{-3}$ for Θ_2 with $\beta_1 = 0.0$ and $\beta_2 = 0.9$. We stopped the inference after 300 rounds of iterations, at which point the model parameters appeared to have converged. This took approximately 3.5 min for 16 images, or 4.5 min for 32 images on an NVIDIA Quadro RTX 6000 with 24 GB of RAM.

4 Experiments

4.1 Synthetic Data

We generate synthetic data by drawing samples from our generative model with random tissue parameters, drawn from Eq. 3, with additional random voxelwise

variability with standard deviations [4, 4, 6, 10, 10] mM. These synthetic images were then transformed to simulate random motion, with translations sampled from $\mathcal{N}(0, 5 \text{ mm}^2)$ and angles from $\mathcal{N}(0, 0.1^2)$. Each of these images was then corrupted with Rician noise at various levels. We then tried to correct for the simulated motion using our model with either a Gaussian or Rician likelihood.




Fig. 2. Synthetic data experiments where the ground truth translation (mean euclidean distance) and rotation error (mean Frobenius norm of the difference of log matrices) are given in the above plots for varying Rician noise level. The dashed line indicates the average initial error. As can be seen, the error when using a Gaussian likelihood rises very quickly, whereas the Rician likelihood is less affected by noise. In this example, $\sigma = 40$ is roughly equivalent to the Sodium MRI data.

Figure 2 illustrates that using the correct likelihood model has a substantial impact on registration performance, particularly in high noise scenarios.

4.2 Real Sodium MRI

^{23}Na MR images were acquired using a dual-tuned, 2-channel (one channel for sodium and one for proton) birdcage ^{23}Na ^1H coil developed by RAPID Biomedical GmbH on a 3T Siemens Prisma scanner. Sodium images were acquired using the FLORET spiral sequence [15] with parameters TR = 120 ms, TE = 0.2 ms, FOV = $256 \times 256 \times 256$ mm, flip angle = 80° , 3 hubs at 22° , 200 interleaves, pulse duration = 0.5ms and dwell time = 0.01 ms. Each acquisition took 1 min and 10s, and was repeated either 16 or 32 times. The k-space data were transferred offline and image reconstruction was performed in Matlab using 3D re-gridding [15] with density compensation [22]. The data was reconstructed with an isotropic resolution of 4 mm^3 and an image size of $64 \times 64 \times 64$. Examples slices are shown in Fig. 1. A T1-weighted image (2 mm^3 isotropic) was also acquired

using the same coil prior to the Sodium data. This was used for preparing tissue segmentation maps using SPM12.

To enable quantification, a set of 4 Sodium phantoms with known concentrations (30, 50, 70, 120 mM) were attached to the head. We use these to map the tissue specific priors, defined in Eq. 3, to the correct intensity range in each image. This mapping is inferred through linear regression of the median signal for each of these phantoms from the true concentrations.




Fig. 3. Bar chart illustrating the mean and std. dev. voxelwise σ , estimated using `scipy.stats.Rice`, over the motion corrected images for 4 subjects either sleeping (s) or awake (a). These numbers are normalised by the estimated σ in the background.

Using this acquisition protocol, we collected data for 4 subjects either when they are asleep (32 Sodium images) or awake (16 images). The data acquired when sleeping is much more likely to contain motion artefacts due to both the length of the scan and unintentional movements during sleep. Accordingly, we use a more permissive transformation prior (with double the standard deviation for rotation and translations) for these examples.

We experiment with motion correcting the sodium magnitude images using either our proposed approach or “`mcfliirt`” [8], using a cost function of normalized correlation and co-registering to the average image. Nearest neighbour interpolation was used as the final step for both approaches for comparable results without introducing additional smoothness or distortion of noise characteristics.

Validation of the proposed model is complicated by the low SNR exhibited in the motion corrected and averaged images, see Fig. 4 for some examples. Desirable properties of aligned images include: similar values at each voxel over images, and sharp boundaries between regions in the average image. We can measure the first of these by fitting a Rice distribution to each voxel, see Fig. 3.




Fig. 4. Example average images calculated by averaging 32 Sodium MRI acquisition. Norm σ refers to the mean voxelwise Rice σ , normalised by an estimate of σ in the image background. In this example, where a lot of motion was detected, our approach leads to a visibly sharper average image.




Fig. 5. Boxplot illustrating the absolute gradient of the average image in voxels on the boundary of CSF. Larger values indicate the presence of stronger edges.

We observe that for sleeping acquisitions that are corrupted with visible motion, particularly subject 3 and 4, our approach reduces the mean voxelwise noise compared to other methods. However, in some awake acquisitions, the use of either motion correction approach increases σ ; we hypothesise this may be due to interpolation artefacts when correcting sub-voxel motion.

Considering the sharpness of the average image, we can visually observe sharper looking average images in examples with large motion, particularly in subject 3 shown in Fig. 4, where we estimated a mean translation of 7.75 mm (6.49 mm std. dev.) and rotation norm of 0.145 (0.12 std. dev.). To quantify the image sharpness, we examine the distribution of absolute gradient values in voxels that lie on the boundary between CSF and anything else, which should have high contrast. We observe that our motion correction induces stronger edges in most of the acquisitions of sleeping participants.

5 Discussion

The presented approach uses a very simple generative model for the image data, which prevents it overfitting to the high level of noise in the data. However, it also prevents it from making use of strong distinctive features such as the eyes, which contain a high level of Sodium, or the phantoms that are attached to the head. Future work will consider using more complex statistical models and techniques, such as variational inference [7], to build a voxelwise generative model. Amortised inference strategies could also be investigated to improve efficiency [4].

In our experimentation, we observed that in some cases where low motion was observed, our algorithm overestimated the level of movement. We found that this was removed by introducing variable transformation permissiveness based on our prior beliefs on the level of motion. Future work will consider methods for inferring these parameters, and using auto-regressive priors on motion [21].

This work has not investigated preprocessing the data using denoising methods, e.g. [12]; although such approaches may produce cleaner representations for aligning the data, they also manipulate the underlying image statistics being modelled, which may lead to biased results. We also have not compared against the use of robust cost functions [17], although these are generally more suited to heavy tailed rather than asymmetric noise distributions as we have here.

We have published our code on GitHub¹. The data are not currently available for distribution as the initial analysis of a wider dataset is ongoing.

6 Conclusions

This paper has introduced an algorithm for data modelling and motion correction of low SNR MRI data using a differentiable approximation of the Rician log-likelihood. Our synthetic experiments illustrated the importance of choosing the right cost function for generative models for motion correction, as the Gaussian likelihood performs very poorly where the errors take a different form. On real Sodium MRI data, our results provide support for the use of our method in resolving substantial motion artefacts and creating sharper average images.

Acknowledgments. We acknowledge funding from the University of Sussex used in the data acquisition. We thank Guillaume Madelin for providing the Sodium MRI sequence.

References

1. Ashburner, J., Friston, K.J.: Unified segmentation. *Neuroimage* **26**(3), 839–851 (2005)
2. Ashburner, J., Neelin, P., Collins, D., Evans, A., Friston, K.: Incorporating prior knowledge into image registration. *Neuroimage* **6**(4), 344–352 (1997)

¹ <https://github.com/ivorsimpson/sodium-mri-inference>.

3. Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ants). *Insight j* **2**(365), 1–35 (2009)
4. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: VoxelMorph: a learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**(8), 1788–1800 (2019)
5. Boyd, S., Boyd, S.P., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge (2004)
6. Gudbjartsson, H., Patz, S.: The Rician distribution of noisy MRI data. *Magn. Reson. Med.* **34**(6), 910–914 (1995)
7. Hoffman, M.D., Blei, D.M., Wang, C., Paisley, J.: Stochastic variational inference. *J. Mach. Learn. Res.* **14**(40), 1303–1347 (2013). <https://jmlr.org/papers/v14/hoffman13a.bib>
8. Jenkinson, M., Bannister, P., Brady, M., Smith, S.: Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**(2), 825–841 (2002)
9. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
10. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.: Elastix: a toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imaging* **29**(1), 196–205 (2009)
11. Madelin, G., Regatte, R.R.: Biomedical applications of sodium MRI in vivo. *J. Magn. Reson. Imaging* **38**(3), 511–529 (2013)
12. Manjón, J.V., Carbonell-Caballero, J., Lull, J.J., García-Martí, G., Martí-Bonmatí, L., Robles, M.: MRI denoising using non-local means. *Med. Image Anal.* **12**(4), 514–523 (2008)
13. Ourselin, S., Roche, A., Subsol, G., Pennec, X., Ayache, N.: Reconstructing a 3D structure from serial histological sections. *Image Vis. Comput.* **19**(1–2), 25–31 (2001)
14. Paszke, A., et al.: *Automatic differentiation in PyTorch* (2017)
15. Pipe, J.G., Zwart, N.R., Aboussouan, E.A., Robison, R.K., Devaraj, A., Johnson, K.O.: A new design and rationale for 3D orthogonally oversampled k-space trajectories. *Magn. Reson. Med.* **66**(5), 1303–1311 (2011)
16. Ramos-Llordén, G., Arnold, J., Van Steenkiste, G., Van Audekerke, J., Verhoye, M., Sijbers, J.: Simultaneous motion correction and t1 estimation in quantitative t1 mapping: an ml restoration approach. In: 2015 IEEE International Conference on Image Processing (ICIP), pp. 3160–3164. IEEE (2015)
17. Reuter, M., Rosas, H.D., Fischl, B.: Highly accurate inverse consistent registration: a robust approach. *Neuroimage* **53**(4), 1181–1196 (2010)
18. Riemer, F., Solanky, B.S., Stehning, C., Clemence, M., Wheeler-Kingshott, C.A., Golay, X.: Sodium (²³Na) ultra-short echo time imaging in the human brain using a 3D-cones trajectory. *Magn. Reson. Mater. Phys., Biol. Med.* **27**(1), 35–46 (2014)
19. Rose, A.M., Valdes, R., Jr.: Understanding the sodium pump and its relevance to disease. *Clin. Chem.* **40**(9), 1674–1685 (1994)
20. Wolfram Mathworld: Modified Bessel function of the first kind. <https://mathworld.wolfram.com/ModifiedBesselFunctionoftheFirstKind.html>
21. Woolrich, M.W., Jenkinson, M., Brady, J.M., Smith, S.M.: Fully Bayesian spatio-temporal modeling of fMRI data. *IEEE Trans. Med. Imaging* **23**(2), 213–231 (2004)
22. Zwart, N.R., Johnson, K.O., Pipe, J.G.: Efficient sample density estimation by combining gridding and an optimized kernel. *Magn. Reson. Med.* **67**(3), 701–710 (2012)