

Bayesian MRI Reconstruction with Structured Uncertainty Distributions

Teo Deveney¹[0000–0002–6367–952X], Ivor Simpson²[0000–0001–5605–6626], and
Neill Campbell^{1,3}[0000–0003–2130–4903]

¹ University of Bath, Bath, UK, BA2 7AY
`t.d.deveney@bath.ac.uk`

² University of Sussex, Brighton, UK, BN1 9RH
`i.simpson@sussex.ac.uk`

³ University College London, London, UK, NW1 2AE
`neill.campbell@ucl.ac.uk`

Abstract. We present a Bayesian hierarchical approach to magnetic resonance imaging (MRI) reconstruction using learned structured uncertainty distributions. Our method allows for reconstruction of complex-valued MRI images in a probabilistic manner that goes beyond standard pixelwise uncertainty. We use a variational autoencoder architecture (VAE) prior with an expressive correlated Gaussian decoding distribution obtained via a sparse parameterisation of the precision matrix, and model the posterior uncertainty in the latent and image space using a similarly correlated variational approximation. The resulting posterior is fully marginalisable over the VAE latent, and provides interpretable insights into the spatial structure of the reconstruction distribution that are not seen in existing methods. Diagnostic posterior pixelwise correlations and residual structure show a principled decay of prior correlation influence with increasing data, and we demonstrate that these modelled posterior statistics are representative of the true reconstruction error. This allows us to answer questions like "how much data is required to resolve a local region to a specific spatial accuracy". We also provide numerical experiments demonstrating that our method maintains excellent pixelwise reconstruction performance and well-calibrated posterior coverage even in extremely sparse data scenarios.

Keywords: Bayesian uncertainty quantification · Generative regularisation · MRI.

1 Introduction

Magnetic resonance imaging (MRI) is a widely used technique that allows the reconstruction of images given Fourier space (k-space) measurements. Compressed sensing (CS) MRI exploits the fact that once transformed into the Fourier domain, images can typically be well approximated by sparse representations. In practice, this allows for significantly shorter scan times with limited effect on the image quality. Image reconstruction in the CS MRI setting can be formulated as

an inverse problem where we seek to recover an image close to the original that best explains the sparse k-space measurements. Traditionally, this is posed as an optimisation problem, solving for an image that optimally balances data fidelity with appropriate regularisation (the so-called variational regularisation framework [14, 11]). Deep learning-based methods have been developed that yield efficient inference and improved image fidelity [7, 1], as well as approaches to learn regularisation [13, 15, 5] that improve reconstruction quality whilst maintaining the robustness and interpretability of the variational framework.

In this work, we propose a Bayesian generalisation to learned regularisation in MRI reconstruction by producing a variational approximation to the posterior distribution over reconstructions using learned image priors. Our results demonstrate competitive reconstruction accuracy, whilst our variational Bayesian construction allows us to extract visually interpretable insights directly from the resulting approximations that surpass the typical variational formulation. For both prior and posterior distributions, we can identify factors such as where our reconstruction is uncertain, how the uncertainty in certain pixels affects our confidence of nearby pixels and the implications of this on localised spatial accuracy, and what alternate instances of these uncertain regions might look like.

1.1 Problem Specification

Mathematically, the CS MRI reconstruction problem can be stated as follows; assume given k-space measurements y , a known forward operator A , and the statistical relation

$$y = Ax + \epsilon, \quad (1)$$

we seek to identify an estimate of x such that the probability of the observed y is high in some sense. We assume that ϵ is a zero-mean Gaussian vector with diagonal covariance $\Lambda = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2)$, thus (1) equivalently implies the likelihood $p(y|x, \Lambda) = \mathcal{N}(y|Ax, \Lambda)$. A direct maximum-likelihood approach could be applied, which would be equivalent to a weighted least-squares reconstruction; however, such approaches are known to perform poorly when data is sparse and so some sort of regularisation is desirable [2]. In the Bayesian setting, regularisation is enforced by defining a prior distribution $p(x)$ over the unknown x . Additionally, in this work we consider the amplitude of ϵ to be uncertain and define a prior $p(\Lambda)$ (independently of $p(x)$) over this, which circumvents the need for manual tuning of the strength of the prior regularisation. From here we can write down the posterior distribution

$$p(x, \Lambda|y) = \frac{p(y|x, \Lambda)p(x)p(\Lambda)}{p(y)}. \quad (2)$$

If Λ is predetermined, then the maximisation of this posterior distribution over x (a MAP estimate) is equivalent to the widely used variational framework [19], and can yield significantly better reconstructions with carefully chosen priors. The problem of approximating the distribution (2), rather than a point estimate, proves decidedly more difficult, largely due to computational challenges in representing distributions over image space.

1.2 Existing Approaches

In addition to deterministic methods based on deep learning and the variational framework, some work has explored deep learning-based uncertainty quantification in inverse problems. At present these approaches to representing uncertainty are overwhelmingly sample-based, typically using denoising-based approaches to construct a prior implicitly through an approximation of the gradient of the logged data distribution (the so-called score function). For example [9] applies this approximate score to unadjusted Langevin sampling to approximate the posterior distribution, and more recently ‘score-based’ [18] or ‘denoising’ [8] diffusion models have been applied to inverse problems [17, 3]. These techniques excel at generating high fidelity samples due their denoising-based mechanism, and it is this denoising-based approach that facilitates the often-praised feature of out-of-distribution generalisation. It is worth noting, however, that the true data distribution that these models target is not supported on out-of-distribution data by definition, thus their ability to generalise to these settings implies significant room for hallucination in cases where data is uncertain compared to priors with more concentrated support over the underlying data.

VAE-based approaches were developed in [20], though the standard pixelwise i.i.d decoder limits prior expressivity and sample-based inference restricts interpretability. Conversely, more expressive correlated decoders have been applied using ‘structured uncertainty prediction networks’ (SUPN) [4, 5] and used to regularise CS MRI, giving some prior interpretability but no posterior uncertainty modelling. This work extends upon that approach through the development of the multichannel SUPN prior models, the extension to a hierarchical Bayesian setting for automatic regularisation strength, and by introducing suitable techniques for interpretable and probabilistic posterior inference.

1.3 Contributions

A summary of our contributions is outlined below:

- We extend the SUPN model for multichannel data and apply this to reconstruct real and complex components in MRI images
- We formulate a Bayesian hierarchical model, and propose a variational inference approach to MRI reconstruction which avoids the need for pre-specification of the measurement error standard deviation
- We demonstrate visually interpretable prior and posterior insights beyond pixel-wise variance estimates. Such as:
 - A principled and representative decay of the correlation structure in the posterior residual as data increases
 - A reduction in the spatial distance that posterior correlations occur between pixels as data increases, implying increasing spatial accuracy
- We demonstrate that even with extremely sparse data, this approach produces posterior distributions with high accuracy and well-calibrated posterior coverage.

2 Methodology

2.1 Structured Uncertainty Priors

For our prior distribution $p(x)$ over reconstructed images, we implement a structured uncertainty prediction network (SUPN), for a detailed description of this model we refer the reader to [4] and additional details in the supplemental material. Here we provide a brief summary in order to motivate its application here. The basis of the SUPN approach is a variational autoencoder (VAE) [10] with a spatially correlated decoding distribution. As such the distribution over x is modelled as

$$p(x) = \int p(x|z)dp(z). \quad (3)$$

Here z is a latent variable with a standard Gaussian prior $p(z) = \mathcal{N}(z|0, I)$, and $p(x|z) = \mathcal{N}(x|\mu_\theta(z), \Sigma_\theta(z))$ is a decoding distribution. In a SUPN model, the decoder outputs a dense covariance matrix Σ_θ through a sparse ($\mathcal{O}(N_{pixels})$) parameterisation of its precision matrix Σ_θ^{-1} . It is known from Gaussian Markov random field (GMRF) theory [16] that precision matrices encode conditional dependence between connected components, so this parameterisation amounts to modelling the image distribution as a GMRF where each pixel is connected to all other pixels within a surrounding connectivity window, then learning the precision terms associated with these connections. The corresponding covariance captures longer range correlations through the Markovian structure of the interconnected pixels as can be seen in Fig. 1. This focus on the finer correlation structures makes this a particularly useful prior for CS MRI since it supplements the higher frequency details that are often lost during subsampling.

Since our data is complex valued, we model real and complex components together as two channels with pixelwise connectivity to allow the real and complex channels to directly depend on one another. When conditioned on an appropriate latent z , the resulting models, like standard VAEs, produces over-smoothed images as their mean output, but unlike a regular VAE, returns a covariance matrix capturing information about the distribution of associated finer scale details. The effect of this can be seen in Fig. 1, which shows the smooth overall shape is captured by the mean, but the bone texture is encoded in the correlation structure. Using this as our prior encourages reconstructions that exhibit realistic fine structures. This prior does not support out-of-distribution data, thus the knee model in Fig. 1 could not be used reconstruct a brain scan without an overwhelming amount of data. Conversely, this prior does produce very reasonable knee reconstructions with as low as 400 k-space measurements (2.5% of the full sample), as shown in Fig. 2. Moreover, since this prior uses a Gaussian decoding distribution we obtain the interpretability of Gaussians, allowing us to do things like read off pixel-wise standard deviation, or as shown in Fig. 1 examine how pixels within the image are correlated with one-another.

Note that the summation of columns two and three in Fig. 1 gives a sample from the prior, and generally produces a detailed and plausible knee reconstruction (an example of this can be seen in the top row of Fig. 2); however, we would




Fig. 1. Visualisation of the structure in our learned prior distribution for both real (top) and complex (bottom) channels. From left to right we show an example image; the prior means; a prior residual sample; and the pixelwise correlation (overlaid on the prior mean), which shows how each pixel is correlated (red positive, blue negative) to the pixel indicated by a white star in the real channel.

not expect the finer details to be representative of the underlying image at this stage, since this is a prior sample before conditioning on data. We see here that there are positive correlations along the contours of the bone’s edge and negative correlation for pixels located orthogonally into the bone as you would expect. Moreover this pattern is reversed for the complex channel since this channel is generally negatively correlated with the real channel.

2.2 Posterior Reconstruction

We seek to approximate the posterior (2) such that we retain the interpretability of our prior distribution. Our representation of $p(x)$ by (3) means that we do not have explicit access to its density, and so we model the joint posterior distribution

$$p(x, z, \Lambda|y) = \frac{p(y|x, \Lambda)p(x|z)p(z)p(\Lambda)}{p(y)}, \quad (4)$$

then recover the reconstruction posterior through the marginalisation $p(x|y) = \int p(x, z, \Lambda|y) dz d\Lambda$. We take a variational Bayes approach to estimating the posterior over x and z , and an expectation maximisation approach to estimate Λ . More explicitly, we introduce a variational posterior approximation $q_\lambda(x, z)$ and perform inference by maximising the evidence lower bound

$$\mathcal{L}(\lambda, \Lambda) = \mathbb{E}_{q_\lambda(x, z)}[\log p(x, z, \Lambda|y) - \log q_\lambda(x, z)]. \quad (5)$$

In practice (5) is maximised by stochastic gradient ascent; that is each iteration we approximate the expectation in (5) with a Monte Carlo approximation

obtained by sampling $(\hat{x}_i, \hat{z}_i \sim q_\lambda)_{i=1:N}$ and computing the estimator

$$\frac{1}{N} \sum_{i=1}^N \log p(y|\hat{x}_i, \Lambda) + \log p(\hat{x}_i|\hat{z}_i) + \log p(\hat{z}_i) + \log p(\Lambda) - \log q_\lambda(\hat{x}_i, \hat{z}_i). \quad (6)$$

Based on this estimate, the parameters (λ, Λ) are updated by gradient ascent. The end result is a distribution $q_\lambda(x, z)$ representing our approximate posterior over (x, z) and a point estimate of the measurement noise Λ . In our implementation we choose $q_\lambda(x, z) = q_{\lambda_x}(x) q_{\lambda_z}(z)$, where $q_{\lambda_z}(z)$ is modelled by a multivariate Gaussian with mean and covariance defined by λ_z , and $q_{\lambda_x}(x)$ is a Gaussian distribution parameterised by λ_x using the same sparse GMRF structure used in the SUPN prior, which allows the interpretability observed in the prior. A discussion on the impact of our chosen structure is included in the supplement. We also define truncated Gaussian priors $p(\Lambda)$ over the components of Λ .

3 Numerical Experiments

We demonstrate our approach on the single coil fastMRI knee dataset [12, 22] at 128×128 resolution, which provides complex k-space data. We reconstruct 16 fully-sampled slices from the central region of 484 scans, giving 7,744 complex-valued training images. We choose a connectivity window for the underlying GMRF that extends 4 pixels in each direction, and train a SUPN prior over the training data. We then estimate the posterior by maximising the evidence lower bound (5) with centre-masked k-space data at several mask sizes. Our code is available at https://github.com/teojd/supn_variational_mri.

The most notable, and unique, benefit of our approach is the interpretability it provides over the posterior spatial structure. For example Fig. 2 displays the residuals and pixelwise correlations for posterior estimates attained at different masking ratios, as well as the prior. This shows, visually, the substantial prior influence on the spatial structure of the reconstruction posterior for small masks, and the decay of this influence as the likelihood dominates with increasing data. This is evidenced in the degrading structure of residual samples and the reducing range of the pixelwise correlations. Such information could be helpful, for example, in determining how much data is required to reconstruct a particular region to a prescribed accuracy. This effect is also shown by the posterior samples themselves, which invariably achieve good agreement with the k-space data, but exhibit greater variability with lower quantities of data, as shown in Fig. 3.

Figure 3 also compares reconstructed samples from our method to the ‘diffusion posterior sampler’ (DPS) [3], which is the current state-of-the-art in generative modelling-based approaches to inverse problems. Samples from our method and DPS are also available as supplementary videos. DPS excels at generating high fidelity reconstructions in this comparison, however, when data are sparse this fidelity comes at the cost of spurious details being added into the reconstruction. In contrast, our model’s more structured prior distribution results in samples that have better agreement with the ground truth as shown by our superior PSNR scores in Fig. 3 in all cases.



Fig. 2. Visualisation of reconstructions of a test-set image and the associated posterior correlations. The first 4 columns show various aspects of our model. In the top row left to right we show a prior residual sample; the prior pixelwise correlation (which illustrates the row of the covariance matrices corresponding to the pixel highlighted with a red star); the prior mean; and a prior sample respectively. Rows 2-6 show the corresponding visualisations for the posteriors after conditioning with 2.5%, 5%, 10%, 20%, 40% mask sizes respectively. These explicitly demonstrate the effect of Bayesian conditioning, as increasing data results in a decay of the prior spatial structure in a way that is consistent with the true residual (as shown in the supplemental material). The final column shows the ground truth in the top row, and the reconstructions resulting from a naive reconstruction (zero-filled inverse Fourier transform) below. Due to space constraints we only show results for the real part, the complex part shows analogous behaviour. Note the colour-maps in columns 1 and 2 have been scaled for interpretability to ensure spatial structure remains visible (actual amplitude decays with increasing data).




Fig. 3. Posterior samples and their PSNR values from our method (left) and DPS (right) using 2.5% mask size (rows 1 and 2), and 40% mask size (rows 3 and 4). Our method scores higher and shows more consistent and realistic samples, particularly at lower data quantities where DPS produces spurious artifacts, such as bone tissue merging or appearing in areas not present in the ground truth.

Posterior statistics such as credible regions can also be analytically extracted from our model, and we demonstrate that these are well calibrated in Table 1 by reporting the actual vs expected posterior coverage at the 1, 2, and 3 standard deviation levels [6]. Conversely, diffusion-based methods require many expensive samples to even estimate credible regions; and attaining analogous correlation statistics to Fig. 2 from these samples would additionally require estimating the empirical covariance, at a further $\mathcal{O}(N_{\text{samples}}N_{\text{pixels}}^2)$ cost that is computationally implausible to estimate for required sample sizes, e.g. $N_{\text{samples}} \sim \mathcal{O}(N_{\text{pixels}})$ [21].

	2.5%	5%	10%	20%	40%
1 Std.Dev. (32%)	36.6%	35.1%	33.9%	33.7%	33.2%
2 Std.Dev. (5%)	8.8%	7.6%	7.2%	7.5%	7.2%
3 Std.Dev. (0.3%)	2.3%	1.8%	1.6%	1.6%	1.4%

Table 1. Proportion of pixels lying outside the posterior credible interval for different mask sizes. Row labels indicate expected proportions for a perfectly calibrated model.

4 Conclusion

This work presents a novel alternative direction to uncertainty quantification in MRI using generative models that does not rely solely on sampling. Our method results in posteriors that are given explicitly as probability distributions, which allows us to provide uniquely insightful visualisations describing the structure and accuracy of the reconstruction, whilst retaining the ability to sample plausible reconstructions from this posterior distribution as with other generative modelling approaches. We present examples of such visualisations, and demonstrate that our posteriors are well calibrated, thus validating our approach.

Acknowledgments. We are grateful to support from the EPSRC CAMERA Research Centre (EP/M023281/1 and EP/T022523/1), the UKRI Strength in Places Fund MyWorld Project (SIPF00006/1), and the EPSRC programme grant in the Mathematics of Deep Learning (EP/L015684/1).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adler, J., Öktem, O.: Learned primal-dual reconstruction. *IEEE Transactions on Medical Imaging* **37**, 1322–1332 (2017), <https://api.semanticscholar.org/CorpusID:26897002>
2. Arridge, S.R., Maass, P., Öktem, O., Schönlieb, C.B.: Solving inverse problems using data-driven models. *Acta Numerica* **28**, 1 – 174 (2019), <https://api.semanticscholar.org/CorpusID:197480023>
3. Chung, H., Kim, J., McCann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: *International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=OnD9zGAGT0k>
4. Dorta, G., Vicente, S., de Agapito, L., Campbell, N.D.F., Simpson, I.J.A.: Structured uncertainty prediction networks. 2018 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 5477–5485 (2018), <https://api.semanticscholar.org/CorpusID:4656813>
5. Duff, M., Simpson, I.J.A., Ehrhardt, M.J., Campbell, N.D.F.: Vaes with structured image covariance applied to compressed sensing mri. *Physics in Medicine & Biology* **68** (2022), <https://api.semanticscholar.org/CorpusID:253116567>
6. Gelman, A.: *Bayesian data analysis*. Chapman Hall/CRC Texts in Statistical Science, CRC Press, Boca Raton, Florida, third edition. edn. (2013 - 2013)
7. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated mri data. *Magnetic Resonance in Medicine* **79**(6), 3055–3071 (Jun 2018). <https://doi.org/10.1002/mrm.26977>, epub 2017 Nov 8
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *ArXiv abs/2006.11239* (2020), <https://api.semanticscholar.org/CorpusID:219955663>
9. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.I.: Robust compressed sensing mri with deep generative priors. In: *Neural Information Processing Systems* (2021), <https://api.semanticscholar.org/CorpusID:236881431>

10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. CoRR **abs/1312.6114** (2013), <https://api.semanticscholar.org/CorpusID:216078090>
11. Knoll, F., Bredies, K., Pock, T., Stollberger, R.: Second order total generalized variation (tgv) for mri. *Magnetic Resonance in Medicine* **65**(2), 480–491 (Feb 2011). <https://doi.org/10.1002/mrm.22595>, epub 2010 Dec 8
12. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence* **2**(1), e190007 (2020). <https://doi.org/10.1148/ryai.2020190007>, <https://doi.org/10.1148/ryai.2020190007>, PMID: 32076662
13. Lunz, S., Öktem, O., Schönlieb, C.B.: Adversarial regularizers in inverse problems. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018)
14. Lustig, M., Donoho, D.L., Pauly, J.M.: Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine* **58** (2007), <https://api.semanticscholar.org/CorpusID:15370510>
15. Mukherjee, S., Dittmer, S., Shumaylov, Z., Lunz, S., Öktem, O., Schönlieb, C.B.: Learned convex regularizers for inverse problems. *ArXiv abs/2008.02839* (2020), <https://api.semanticscholar.org/CorpusID:221083179>
16. Rue, H., Held, L.: *Gaussian markov random fields: Theory and applications* (2005), <https://api.semanticscholar.org/CorpusID:60632241>
17. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. *ArXiv abs/2111.08005* (2021), <https://api.semanticscholar.org/CorpusID:244130146>
18. Song, Y., Sohl-Dickstein, J.N., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *ArXiv abs/2011.13456* (2020), <https://api.semanticscholar.org/CorpusID:227209335>
19. Stuart, A.M.: Inverse problems: A bayesian perspective. *Acta Numerica* **19**, 451–559 (2010). <https://doi.org/10.1017/S0962492910000061>
20. Tezcan, K.C., Baumgartner, C.F., Luechinger, R., Pruessmann, K.P., Konukoglu, E.: Mr image reconstruction using deep density priors. *IEEE Transactions on Medical Imaging* **38**(7), 1633–1642 (2019). <https://doi.org/10.1109/TMI.2018.2887072>
21. Vershynin, R.: How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability* **25** (04 2010). <https://doi.org/10.1007/s10959-010-0338-z>
22. Zbontar, J., Knoll, F., Sriram, A., Muckley, M., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M.G., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastmri: An open dataset and benchmarks for accelerated mri. *ArXiv abs/1811.08839* (2018), <https://api.semanticscholar.org/CorpusID:53759905>