

PAPER • OPEN ACCESS

VAEs with structured image covariance applied to compressed sensing MRI

To cite this article: M A G Duff et al 2023 Phys. Med. Biol. 68 165008

View the article online for updates and enhancements.

You may also like

- Equivariant neural networks for inverse problems Elena Celledoni, Matthias J Ehrhardt, Christian Etmann et al.
- <u>An inner-outer iterative method for edge</u> preservation in image restoration and reconstruction Silvia Gazzola, Misha E Kilmer, James G Nagy et al.
- <u>On Krylov methods for large-scale CBCT</u> reconstruction Malena Sabaté Landman, Ander Biguri, Sepideh Hatamikia et al.

JOIN US | ESTRO 2024 In-Booth Talks, Demos, & Lunch Symposium

Browse talk schedule >



This content was downloaded from IP address 18.118.137.243 on 03/05/2024 at 12:11

Physics in Medicine & Biology

PAPER

OPEN ACCESS

CrossMark

RECEIVED 19 January 2023

REVISED 16 June 2023

ACCEPTED FOR PUBLICATION 5 July 2023

PUBLISHED

3 August 2023

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



VAEs with structured image covariance applied to compressed sensing MRI

M A G Duff¹ , I J A Simpson², M J Ehrhardt¹ and N D F Campbell³

¹ Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, United Kingdom

² Predictive Analytics Lab, Department of Informatics, University of Sussex, BN1 9RH, United Kingdom

³ Department of Computer Science, University of Bath, Bath, BA2 7AY, United Kingdom

 $\label{eq:constraint} E-mail: M.A.G.Duff@bath.ac.uk, I.Simpson@sussex.ac.uk, M.Ehrhardt@bath.ac.uk and N.Campbell@bath.ac.uk and N.Campbell@bath.a$

Keywords: imaging, inverse problems, generative models, MRI, variational autoencoders, machine learning

Abstract

Objective. This paper investigates how generative models, trained on ground-truth images, can be used as priors for inverse problems, penalizing reconstructions far from images the generator can produce. The aim is that learned regularization will provide complex data-driven priors to inverse problems while still retaining the control and insight of a variational regularization method. Moreover, unsupervised learning, without paired training data, allows the learned regularizer to remain flexible to changes in the forward problem such as noise level, sampling pattern or coil sensitivities in MRI. *Approach.* We utilize variational autoencoders that generate not only an image but also a covariance uncertainty matrix for each image. The covariance can model changing uncertainty dependencies caused by structure in the image, such as edges or objects, and provides a new distance metric from the manifold of learned images. *Main results.* We evaluate these novel generative regularizers on retrospectively sub-sampled real-valued MRI measurements from the fastMRI dataset. We compare our proposed learned regularization against other unlearned regularization approaches and unsupervised and supervised deep learning methods. *Significance.* Our results show that the proposed method is competitive with other state-of-the-art methods and behaves consistently with changing sampling patterns and noise levels.

1. Introduction

Compressed sensing magnetic resonance imaging (MRI) provides the benefits of MRI with faster acquisition times, providing accurate reconstructions from sparsely sampled k-space data. Sophisticated mathematical reconstruction techniques allow for high-quality images to be produced with just a subset of the full MRI measurements, which are quicker to obtain. This reduces costs but can also improve image quality by reducing motion artifacts. Mathematically, we seek to recover the MRI magnitude image, $x \in \mathcal{X} = \mathbb{R}^d$, from observed measurements, $y \in \mathcal{Y} = \mathbb{C}^m$. The two are related by a linear forward model, $A: \mathcal{X} \to \mathcal{Y}$ giving the equation Ax = y. In MRI, A is composed of a Fourier transform and a subsampling operator that takes just a subset of the Fourier data. This is an ill-posed inverse problem, where measurements are incomplete and so multiple solutions may exist.

This difficulty can be mitigated by incorporating prior information; we consider this to be given in the form of a regularizer, \mathcal{R} , in a variational regularization framework (Lustig *et al* 2007, Knoll *et al* 2011, Xu and Noo 2022)

$$\arg\min_{x} d(Ax, y) + \mathcal{R}(x), \tag{1}$$

where $d: \mathcal{Y} \times \mathcal{Y} \to [0, \infty]$ is a similarity measure, ensuring that the reconstructed image matches the data, and $\mathcal{R}: \mathcal{X} \to [0, \infty]$ is a regularization term which is small if the image satisfies some desired property.

We extend recent work, including (Bora *et al* 2017, Dhar *et al* 2018, Tripathi *et al* 2018, Duff *et al* 2021, Habring and Holler 2022) that consider the use of a generative machine learning model as part of the regularizer,

called *generative regularizers*. A latent generative model is designed to take a sample from a distribution in a lowdimensional latent space and generate data similar to the training distribution. Common deep generative models include variational autoencoders (VAEs) (Kingma and Welling 2014), generative adversarial networks (GANs) (Goodfellow *et al* 2014) and normalizing flows (Rezende and Mohamed 2015). The generative model is trained on high-quality data, obtained from fully sampled measurements, with no knowledge of the forward model, *A*, or any sub-sampled data. The idea of generative regularizers is to penalize inverse problem solutions that are dissimilar to the learned distribution on ground truth, or high-quality reconstructed, images.

The generative part of a VAE consists of a network, such that for any point $z \in \mathcal{Z} = \mathbb{R}^n$, within the learned latent space, the network outputs a distribution on the image space \mathcal{X}

$$p_{G,\Sigma}(x|z) = \mathcal{N}(x; G_{\theta}(z), \Sigma_{\theta}(z))$$
(2)

for parameterized functions $G_{\theta}: \mathbb{Z} \to \mathcal{X}$ and $\Sigma_{\theta}: \mathbb{Z} \to \mathbb{R}^{d \times d}$, $d := \dim(\mathcal{X})$. In a standard VAE model, Σ_{θ} is taken to be a multiple of the identity matrix, i.e. $\Sigma_{\theta}(z) := \rho I$, where the predicted variance, ρ , is either fixed or learned. This assumes that the reconstruction error of the generated image is independent and identically distributed for all pixels. In reality, parts of an image, such as the background, will be easy to model whereas sharp edges and other high frequencies are harder to model accurately. Similarly, as images are commonly piecewise smooth, there are often local correlations in the errors made by the model. A well-known issue with VAEs is their tendency to produce smooth images (see e.g. Ruthotto and Haber (2021)), missing the sharp edges that exist in real data. In this work, we consider the effect of a more expressive covariance matrix, Σ_{θ} , in a generative regularizer.

1.1. Contributions

We propose an adaptive generative regularizer where edge and correlations in the data are modeled with a structured covariance network. We demonstrate the strength of this model by reconstructing knee MRI images from retrospectively sub-sampled real-valued data from the fastMRI dataset (Zbontar *et al* 2018). Our contributions are outlined below.

- An adaptation of the structured covariance model introduced by Dorta *et al* (2018) for magnitude images from the fastMRI dataset of knee MR images.
- An extension of the denoising example from Dorta *et al* (2018) for use in inverse problems with a non-trivial forward model, producing a novel generative regularizer.
- A demonstration of how the prior provided by the VAE with structured covariance can be explicitly visualized. We see that the structured covariance model has learned long-range correlations between pixels, taking into account image structure.
- An ablation study to compare three different options for the decoder covariance, Σ_θ(z): Σ is a fixed constant multiplying the identity matrix, Σ is a diagonal matrix with a learned diagonal and the proposed structured covariance where Σ is dense. We show the most flexible, dense covariance method, produces the best inverse problem results.
- Comparisons of our proposed regularizer with a variety of regularizers: the unlearned total variation (TV) (Rudin *et al* 1992) and least squares with early stopping methods. Also with three unsupervised methods: a deep image prior approach with a pre-trained generator from Narnhofer *et al* (2019), the original Compressed Sensing using Generative Models work by Bora *et al* (2017) and a Plug-and-Play method by Ryu *et al* (2019). Finally, with a state-of-the-art, learned, end-to-end method, variational networks (VNs) (Hammernik *et al* 2018). We demonstrate that our method is competitive with the state-of-the-art methods, yet offers superior generalization to unseen noise statistics and sampling patterns.

2. Related work

Deep learning approaches to image reconstruction in inverse problems is a growing field (Arridge *et al* 2019). There are several supervised deep learning approaches (Wang *et al* 2016, Hammernik *et al* 2018, Hyun *et al* 2018, Oh *et al* 2018, Quan *et al* 2018, Zhu *et al* 2018, Mardani *et al* 2019) that require datasets of subsampled measurements paired with high-quality reconstructions. With any change in the forward model, e.g. a different k-space sampling pattern, or noise level, new data needs to be acquired and models retrained. Furthermore, with these methods, care needs to be taken to ensure the image reconstruction is consistent with the observed measurements and these methods can also be unstable to small perturbations in the measured data (Antun *et al* 2020). In contrast, deep image priors (Narnhofer *et al* 2019, Ulyanov *et al* 2020), have no data requirements and

instead use an untrained convolutional neural network as a prior for the inverse problem. The prior is implicit and comes from the architecture choices, another choice to make, but also requires regularization in the form of early stopping to prevent over-fitting to the potentially noisy data.

We choose to take an unsupervised approach, where we have example ground truth images but no paired data. The image modeling, for training the regulariser, and the forward modeling for the inverse problem reconstruction, are completely decoupled. We model the images using a generative model incorporating it as part of a regulariser for the inverse problem reconstruction. There has been previous work in this area (Bora *et al* 2017, Dhar *et al* 2018, Tripathi *et al* 2018, Duff *et al* 2021, Habring and Holler 2022), and we extend it with the addition of a structured image covariance network. We train our generative model without paired training data and without knowledge of the forward problem and point to the interesting works of Zhang *et al* (2021) and Jalal *et al* (2021) for examples of these settings. Other unsupervised approaches include plug-and-play methods which iteratively call a learned denoiser within a larger optimization or inference algorithm (Ryu *et al* 2019, Ahmad and Bouman 2020).

When choosing and training a generative model for use in inverse problems, the generative model must be able to produce a whole range of possible solutions. A common issue with GANs (Goodfellow *et al* 2014) is, that they do not capture the full distribution of images they were trained on. This failure can be subtle (Arora *et al* 2018) and therefore difficult to identify. GANs also do not have an encoder, and finding a latent space value that corresponds to a particular image is a non-convex optimization problem with multiple local minima. In comparison, a VAE is more suitable for use in generative regularizers because they can reconstruct images across the span of the training distribution although with the consequence of fewer high frequencies. We consider VAEs over other generative models such as normalizing flows or invertible neural networks, both of which have been applied to inverse problems (Kingma and Dhariwal 2018, Jalal *et al* 2021), because of the regularizing effect of a lower dimensional latent space in the VAE. Dorta *et al* (2018) proposed the use of a structured covariance as part of a VAE for denoising, and the novel addition of this work is the application to MRI reconstruction.

The Dorta *et al* (2018) approach parameterizes a dense covariance through the Cholesky decomposition of the corresponding precision matrix, the inverse covariance matrix, Σ_{θ}^{-1} . Other approaches could include a diagonal plus low-rank parameterization of the covariance matrix, as seen for a segmentation task in Monteiro *et al* (2020). While we believe a low-rank approximation may not give the fine detail required for image reconstruction, a combination of the two approaches could be an avenue for future work.

3. Method

We build a probabilistic model for the reconstructed image, *x*, given an observation, *y*. First, we consider the likelihood of the measurement, *y*, given image, *x*, denoted p(y|x), and usually taken to be $\mathcal{N}(y; Ax, \gamma^2 I)$ for additive Gaussian noise over the observations with standard deviation, γ . We choose a prior on the images *x* given by a pre-trained generative model, $p_{G,\Sigma}(x|z)$ as in (2). Finally, let p_Z be prior on the latent space used to train the generator, usually $\mathcal{N}(z; 0, I)$. Combining these parts, we have that

$$p(x, z|y) \propto p(y|x, z) \, p_{G,\Sigma}(x|z) \, p_{Z}(z) \, = \, p(y|x) \, p_{G,\Sigma}(x|z) \, p_{Z}(z), \tag{3}$$

where the equality follows because y is independent of z given x.

Ideally, we would seek to marginalize out the latent vector, *z*, however this integral is intractable, except by expensive sampling over the full distribution of \mathcal{Z} ; instead, we take a maximum *a posteriori* (MAP) estimate. By taking logarithms, maximizing (3) with respect to *x* and *z* is equivalent to variational regularization (1) with

$$d(Ax, y) = \frac{1}{2\gamma^2} ||Ax - y||_2^2$$
(4)

and

$$\mathcal{R}(x) = \min_{z} \left(\log(|\Sigma_{\theta}(z)|) + \frac{\|x - G_{\theta}(z)\|_{\Sigma_{\theta}(z)}^{2}}{2} + \frac{\|z\|_{2}^{2}}{2} \right),$$
(5)

where we denote the weighted norm by $||x||_M^2 := x^T M^{-1}x$ and the determinant of a matrix Σ by $|\Sigma|$. Equation (4) ensures that *x* explains the observation *y*, while the second term in (5) constrains *x* to be close to images in the range of the generator.

3.1. VAEs with structured covariance

The generative model is trained to minimize the distance between the generated, $p_{G,\Sigma}(\cdot; \theta)$, and training, p_{Im} , distribution. A VAE is derived by choosing the distance measure to be a Kullback–Leibler divergence and then maximizing a lower bound approximation to this distance (Kingma and Welling 2014). The intractable distribution over the latent space, $p_{G,\Sigma}(z|x; \theta)$, is approximated by an encoder

$$q(z|x;\psi) = \mathcal{N}(z;\mu_{\psi}(x),\operatorname{diag}(\sigma_{\psi}^{2}(x))) =: \mathcal{N}_{z,\psi}(x)$$
(6)

with neural networks μ_{ψ} , σ_{ψ}^2 , parameterized with weights, ψ . Following the derivation in Kingma and Welling (2014), training a VAE becomes a minimization with respect to ψ and θ of

$$\mathbb{E}_{x \sim p_{\mathrm{Im}}} \mathbb{E}_{z \sim \mathcal{N}_{z,\psi}(x)}(\mathrm{nll}(x, G_{\theta}(z), \Sigma_{\theta}(z)) + d_{\mathrm{KL}}(\mathcal{N}_{z,\psi}|p_{\mathcal{Z}})), \tag{7}$$

where we write the negative log-likelihood as

$$\operatorname{nll}(x, G, \Sigma) = \log(|\Sigma|) + \frac{1}{2} ||x - G||_{\Sigma}^{2} + \operatorname{constants}$$
(8)

and the expectation over x is calculated empirically over the training set via sampling. The expectation over z is also calculated empirically and for more details see the original paper (Kingma and Welling 2014).

3.2. Structure of $\Sigma_{\theta}(z)$

Using a dense matrix Σ_{θ} is computationally infeasible for even moderate-sized images as it is expensive to calculate both the log determinant and the inverse required for the norm. We use the computationally efficient structured uncertainty prediction networks as developed by Dorta *et al* (2018). For each point in the latent space, an additional decoder, parameterized by θ , outputs a sparse lower triangular matrix, $L_{\theta}(z)$, with the diagonal constrained to be positive. This forms the Cholesky decomposition of the precision matrix, Σ_{θ}^{-1} , such that $\Sigma_{\theta} = (L_{\theta}L_{\theta}^T)^{-1}$. The Cholesky decomposition is taken to be sparse: $(L_{\theta})_{i,j} = 0$ if $j \neq N_i$ where N_i is the set of neighbors of the pixel *i*. This leads to a sparse precision matrix, Σ_{θ}^{-1} . A zero value at entry *i* and *j* in Σ_{θ}^{-1} means these pixels are independent, conditioned on all other image pixels. The pixels could still be correlated and thus this still allows for a dense covariance matrix. Indeed we would not choose to make Σ_{θ} sparse because zero values in the covariance indicate two uncorrelated pixels and we wish to allow dependencies across images. The local sparsity structure is also amenable to parallelization on the GPU and for more details see Dorta *et al* (2018). A pictorial view of the full network including an example sparse Cholesky matrix is given in figure 1.

3.3. Objective functions

This sparse Cholesky decomposition of the precision makes predicting a dense covariance matrix computationally feasible. The negative log-likelihood from (7) becomes

$$\operatorname{nll}(x, G, \Sigma) = -2\sum_{i=1}^{d} \log(L_{ii}) + \frac{1}{2} \|L(x-G)\|_{2}^{2}$$
(9)

up to constant terms and with, as above, $\Sigma^{-1} = LL^{T}$. The log determinant is now just a summation over the diagonals of the Cholesky matrix, we have removed the dense matrix inversion and there is no need to build the full covariance matrix. This also simplifies the regularizer from (5), giving

$$\mathcal{R}(x) = \min_{z} \left(nll(x, G_{\theta}(z), \Sigma_{\theta}(z)) + \frac{1}{2} ||z||_{2}^{2} \right).$$
(10)

Moreover, sampling from the extended VAE is possible by solving a sparse system of equations to get a sample $x = G_{\theta}(z) + (L_{\theta}(z)^T)^{-1}u$ where $u \sim \mathcal{N}(0, I)$.

We note that the minimization problems in (7) and (1) with (10) are not well defined as they are not bounded from below. Pixel values that can be determined with high accuracy e.g. those of a consistent black background, can have extremely low variance and the log determinant term may become large and negative. We both bound the size of the diagonals of L_{θ} using a scaled tanh activation and added a very small amount of noise to the black background to deal with this. Investigating other priors on L_{θ} is an interesting area for future work.

4. Experiments

4.1. Dataset

The NYU fastMRI knee dataset contains, amongst other data, 796 fully sampled knee MRI magnitude volumes (Zbontar *et al* 2018, Knoll *et al* 2020), without fat suppression. This data was acquired on one of three clinical 3T systems (Siemens Magnetom Skyra, Prisma and Biograph mMR) or one clinical 1.5T system, using a 15-channel knee coil array and conventional Cartesian 2D TSE protocol employed clinically at NYU School of Medicine. We use the ground truth data from the fastMRI single coil challenge, where the authors used emulated single-coil methodology to simulate single-coil data from the multi-coil acquisition. In addition, in the fastMRI dataset, the images are cropped to a square, centering the knee. We extract 3872 training and 800 test ground truth slices from this dataset, selecting images from near the center of the knee, resize the images, using the Python imaging library (Clark 2015) function with an anti-aliasing filter, to 128×128 pixels and linearly rescale each image to



the range (0, 1). The training and test datasets correspond to the training and validation sets of the fastMRI dataset and images from the same volume are always contained in the same dataset.

4.2. Forward problem

Our forward problem is inspired by the fastMRI single-coil reconstruction task; reconstructing images approximating the ground truth from under-sampled single-coil MR data (Zbontar *et al* 2018). The ground truth images are Fourier transformed and subsampled. To this end, a mask selects the points in k-space corresponding to a sampling pattern. We use both radial and cartesian sampling patterns, selecting radial spokes and horizontal lines in k-space, respectively. We use the operator discretization library (Adler *et al* 2017) in Python and take the same radial subsampling MRI operators as in Bungert and Ehrhardt (2020). Note that this is a relatively simple MRI model, yet a good starting point to test the feasibility of the proposed approach. We discuss its limitations in section 6.

4.3. Model architecture

See figure 1 for a comparison of the VAE and the VAE with structured image covariance. All the networks are built of resnet-style blocks, which consist of a convolutional layer with stride 1, a resizing layer, followed by two more convolutional layers, and then a ReLU activation function. The output of this process is then added to a resized version of the original input to the block. The resizing layer is either a bilinear interpolation for an upsampling layer, increasing width and height by a factor of 2; convolutions with stride 2 for a down-sampling layer or a convolution with stride 1 for a resnet block that maintains image size. We choose a latent space of size 100. The generative network consists of a single dense layer outputting 8 × 8 images with 16 channels before a resnet block without resizing to give 8 × 8 images with 256 channels, then four up-sampling blocks are applied giving image sizes 16×16 , 32×32 , 64×64 and 128×128 with channels 512, 256, 128 and 64 and final another resnet block without resizing to reduce the channels down to one output image. The covariance network is identical but outputs a 128×128 image with 5 channels, from which the sparse Cholesky matrix is formulated, based on

code from Dorta *et al* (2018)⁴. The sparsity pattern is chosen such that the Cholesky matrix $L_{\theta}(z)_{i,j}$ is non-zero only if pixels *i* and *j* lie in the same 5 × 5 patch. The encoder is reversed copy of the generator, with down-sampling layers replacing the upsampling layers. We include drop-out layers during training.

4.4. Model training

Training the VAE with structured covariance is done in a two-stage process. First the generated mean, G_{θ} , and the encoder, $\mathcal{N}_{z,\psi}(x)$, are trained with a standard VAE loss (Kingma and Welling 2014) i.e. $\Sigma_{\theta} = \rho I$ where ρ is a fixed hyperparameter. The weights for the mean decoder and the encoder are then fixed before the covariance model is trained using the full structured noise loss given by (9). The choice of this two-stage training is two-fold: firstly, it forces as much information as possible to be stored in the weights of the mean, and not the covariances, and secondly, it allows us to compare the effect of just changing the covariance models in the ablation study. Models were built and trained in TensorFlow using an NVIDIA RTX 2080-8GB GPU. It took approximately 8 h to train the means and then 24 and 30 h for the diagonal and structured covariance models. The code will be available on publication.

4.5. Ablation study

We consider three variations on the structure of $\Sigma_{\theta}(z)$: Σ_{θ} is a fixed constant multiplying the identity matrix, Σ_{θ} is a diagonal matrix with a learned diagonal and our proposed method where Σ_{θ} is a dense matrix. We call these options: *mean+identity, mean+diagonal* and *mean+covar*^{*}. The *mean+identity* model is just taken to be the output of the first part of the *mean+covar*^{*} training described above. For *mean+diagonal*, we again take the learned generated mean, G_{θ} , and the encoder, $\mathcal{N}_{z,\psi}(x)$ from the *mean+identity* model and then optimize (7) for the covariance network, but choose the off diagonals of the Cholesky matrix, $L_{\theta}(z)$, to be zero, so that the final covariance matrix is diagonal.

4.6. Proposed reconstruction method

For the proposed method *mean+covar*^{*} and the versions *mean+identity* and *mean+diagonal* we choose a variation on the above regulariser (10)

$$\mathcal{R}(x) = \min_{z} \lambda \bigg(\operatorname{nll}(x, \, G_{\theta}(z), \, \Sigma_{\theta}(z)) + \frac{\mu}{2} \|z\|_{2}^{2} \bigg).$$
(11)

The addition of the two regularization parameters λ and μ is in recognition that the optimization problem in (10) is a non-convex problem. We use alternating gradient descent with backtracking line search (see e.g. algorithm 9.2 of Boyd and Vandenberghe (2004)) where gradient descent steps are taken, alternating in the *x* and *z* space, with step size chosen to insure the objective decreases. We initialize at a rough reconstruction, given by the adjoint of the forward operator, and the encoding of the adjoint for *x* and *z*, respectively. Regularization parameters were selected via a grid search to maximize the peak signal-to-noise ratio (PSNR).

4.7. Unlearned method comparisons

We compare against TV regularization (Rudin *et al* 1992) implemented using the primal–dual hybrid gradient method (Chambolle and Pock 2011), with the regularization parameter chosen to maximize PSNR. As a baseline, we also calculate the least squares solution, $\min_x ||Ax - y||_2^2$, optimized using gradient descent with backtracking line search.

4.8. Data driven prior comparisons

For another example of a generative regulariser, we take our trained mean generator G_{θ} and implement the method of Bora *et al* (2017), minimizing with respect to *z* the objective

$$L(z) = \frac{1}{2} \|AG(z) - y\|_2^2 + \mu \|z\|_2^2,$$
(12)

searching for an image in the *range* of the generator that best matches the measurements. The regularization parameter, μ , is chosen to maximize PSNR. We also implement the approach of Narnhofer *et al* (2019) which takes a trained generator but then, after observing data, tweaks the weights of the network, optimizing z^* , $\theta^* \in \operatorname{argmin} ||AG_{\theta}(z) - y||_2^2$, we call this *Narnhofer19*. We use the same mean generator as in our other experiments, gradient descent with backtracking for the *z* optimization and TensorFlow's inbuilt Adam optimizer for the network weight update. We choose 1000 iterations to find an initial value of *z* and then select the number of iterations for the alternating *z* and θ updates via a grid search to maximize the PSNR.

Finally, we take the method of Ryu *et al* (2019), specifically their Plug-and-play alternating direction method of multipliers (PnP-ADMM) algorithm, which takes an iterative ADMM approach but replaces the proximal

function in the image space with a learned denoiser. Specifically, following the derivation in Boyd *et al* (2011) we have

$$x_{k} \in \arg\min_{x} \left\{ \frac{1}{2\sigma^{2}} \mathcal{D}(Ax, y) + \frac{1}{2\eta} \|x - u_{k-1} + v_{k-1}\| \right\},$$
(13)

$$u_{k} \in \arg\min_{x} \left\{ \frac{1}{2\eta} \|x_{k} - v_{k-1} - x\|_{2}^{2} + \mathcal{R}(x) \right\},$$
(14)

$$v_k = v_{k-1} + (x_k - u_k). \tag{15}$$

The first equation (13) can be solved analytically and, in (14), instead of specifying a regularization term, the whole term is replaced by a learned denoiser. The variables σ and η are constants, x_0 is a rough reconstruction given by the adjoint operator and u_0 and v_0 are initialized at zero. We use the code implemented by the authors⁵ for both training the denoiser and for the iterations. We take their RealSN-DnCNN denoiser using the default settings, training on the fastMRI dataset as with the other methods. When running the iterative method on measured data, we choose the number of iterations to maximize PSNR.

4.9. Supervised end-to-end method comparison

Finally, we also demonstrate comparisons with a VN (Hammernik *et al* 2018). This is an end-to-end approach that takes (5) with $d(Ax, y) = \frac{\lambda}{2} ||Ax - y||_2^2$ and treats the regulariser \mathcal{R} as some unknown function. Optimizing (5) by gradient descent will lead to updates of the form

$$x_{t+1} = x_t - \alpha_t \lambda A^T (A x_t - y) + \nabla \mathcal{R}(x_t),$$
(16)

where A^T is the adjoint of A and α_t is a step-size. The authors replace the unknown $\nabla \mathcal{R}(x_t)$ term with a learned component, inspired by a fields of experts model (Roth and Black 2009). Fixing the number of iterations and unrolling leads to an end-to-end method that includes information from the forward and its adjoint. We use the same network design and parameters as the original paper. The iterations are first initialized with a rough reconstruction given by the adjoint. The gradient of the regularizer consists of a convolution with 24 filters, kernel size of 11 and stride 1; a component-wise activation function; and then a convolution transpose with the same dimensions to return an image with the same dimensions as the input. The activation function consists of a weighted sum of 31 radial basis functions with learnable means, variance and weights. All learnable parameters are allowed to vary independently in each layer and the step size α_t and parameter λ are also learned. We train three variations on our data: 25 radial spokes with 0.05 and 0.0125 added noise and 125 radial spokes with 0.05 added noise.

5. Results

5.1. Ablation study

To compare the covariance models *mean+identity*, *mean+diagonal* and *mean+covar**, visualizations of samples from the learned covariances are given in figure 2. We see that the *mean+diagonal* places uncertainty in the right locations, but without structure, the samples are speckled and grainy. The *mean+covar** models can match some of the missing structures from the generated mean images. We note that these are random samples and therefore not expected to match the residual precisely but rather illustrate statistical similarity.

Figure 3 compares the effectiveness of the three covariance models as generative regularisers. The same observed data was used for each reconstruction method for each image. The results for *mean+covar** give a consistently higher PSNR value than *mean+diag* and *mean+identity* and we use this method throughout the rest of the results section. These results support our hypothesis that learning a more flexible distance measure allows us to better fit the data and gives a better regularizer for our inverse problem.

5.2. Prior introspection

We can explicitly examine the learned structured covariance, to visually assess the prior information passed to the generative model. To do this, we take a test image, *x*, and use the encoder to give a latent vector, *z*, that corresponds to the test image. From this we can calculate the structured covariance, $\Sigma_{\theta}(z) \in \mathbb{R}^{d \times d}$, where $d = \dim \mathcal{X}$. Each row of the covariance matrix corresponds to a pixel in the generated image, and the row can be reshaped into an image, showing the correlations between the chosen pixel and all others. Two example images with chosen pixels highlighted with a star are shown in figure 4 and further examples are given in a supplementary video⁶. Positive correlations are given in red, and negative correlations are in blue. We can see

⁵ https://github.com/uclaopt/Provable_Plug_and_Play

⁶ https://youtube.com/watch?v=_bi2D7rJ0OA



Figure 2. Comparing the covariance models *mean+diagonal* and *mean+covar**. The first column gives the ground truth, and the second column a reconstruction in the *range* of the VAE generator. The third provides the residual that can be considered as one sample of a zero mean Gaussian distribution with unknown covariance. The final two columns give single residual samples from our learned covariances *mean+diagonal* and *mean+covar** models calculated using $(L_{\theta}(z)^T)^{-1}u$ where *u* is one sample from $\mathcal{N}(0, I)$ and $L_{\theta}(z)$ is diagonal and lower triangular for the *mean+diagonal* and *mean+covar** models, respectively. Columns are rescaled for better visualization.





that the structure of the edges is present, and that, despite the local structure of the precision matrix, longerrange correlations have been learned.

5.3. Comparison with unlearned methods

We compare the proposed method with TV and least squares reconstructions for a range of noise levels and the number of radial spokes in the sampling pattern. Image examples can be seen in columns 2 and 3 of figures 5 and 6. In figure 7, we see that the results of *mean+covar** track TV but with improvements across the range of radial spokes and noise levels tested. Especially in the examples in figure 6 you can see the piece-wise constant shapes typical of a TV reconstruction, whereas for the same measured data *mean+covar** has managed to reconstruct more of the fine detail. As expected, due to the difficult nature of the inverse problem, the least squares reconstruction does poorly. As a rough indicator, our un-optimized implementation took 1.8 s for reconstructing one image using least squares, and 6.4 s using TV, for one choice of regularization parameter. For one choice of the regularization parameter, *mean+covar** took 78.2 s.

8





Figure 5. Example reconstructions of a test image for different amounts of measured data: 5, 25, 45 and 125 radial spokes, all with additive Gaussian noise with standard deviation 0.05. The columns give different reconstruction methods. The PSNR values are added in white and the red boxes indicate the settings the highlighted variational network has been trained on.

















5.4. Comparison with other data-driven priors

Comparisons with *Narnhofer19*, *PnP-ADMM* and *range* are given in figure 8 and columns 4 and 5 of figures 5 and 6. We see that the results of searching in the *range* of the generator saturate so that even with increased data in the form of radial spokes, or better data in the form of less noise, the reconstructions do not see significant improvement. The example image reconstructions reflect this; they show overly smoothed images that, although have the right structure, do not contain any of the fine detail of the ground truth. This could be because the generative model is not expressive enough to match this fine detail. Alternatively, it could be because the



Figure 10. Example reconstructions of an image using measured data taken with different horizontal random sampling unseen in the training of the variational networks. The first column gives the k-space mask used to sample the data. The remaining columns give reconstructions for the different methods.

non-convex optimization has failed to find a good enough minimum to match the measured data. The results of *Narnhofer19* are much more detailed and very similar in PSNR values to our *mean+covar**. In the images, there is some evidence of smoothing of details e.g. in figure 6. The *PnP-ADMM* method did better for a higher number of radial spokes when we would expect the adjoint to be similar to images in the training set of high-quality reconstructions, however, it is still not competitive with our proposed method or *Narnhofer19*. Our unoptimized implementation took 106.7 s for *Narnhofer19* and 23.9 s for the *range* method. The *PnP-ADMM* method took 0.025 s to run, terminating at a small number of iterations.

5.5. Comparison with supervised end-to-end methods

We now compare *mean+covar*^{*} with VNs (Hammernik *et al* 2018), trained end-to-end, using paired training data obtained using particular noise and sampling settings. Figure 9 shows the PSNR results for 50 test images, with vertical lines highlighting the settings where the networks were trained. Example images are shown in columns 6 and 7 of figures 5 and 6 with the trained for setting highlighted in red. We note that the VNs achieve the best results, in comparison to the other methods, for the settings they were trained on. This is as expected for end-to-end reconstructions. We see that the results are less optimal the further from the trained for setting, with some particularly poor results, e.g. in figure 5, while our unsupervised method remains consistent. Also as expected, the VNs are quick to implement once trained, taking 3.4 s for one image.

5.6. Generalization to a different sampling pattern

Finally, we test the generalization ability of the methods by changing the sampling pattern. We mask horizontal rows in the k-space, taking 16 (out of 128) center rows and a uniformly random selection of other rows with a given probability, p. Fully sampled images correspond to p = 1. Example sampling patterns and one example reconstruction is given in figure 10. We see an improvement over TV for *mean+covar**. *Narnhofer19* gives good PSNR values but the images are potentially over-smoothed. For example, the vertical lines in the top left part of the zoomed-in region are consistently missing from the *Narnhofer19* reconstructions. The VN has not been trained on this horizontal sampling pattern and although gives a reasonable reconstruction, seems to have some artifacts, for example in the top right and bottom left of the image.

6. Discussion

The numerical results show that the generative regularizer *mean+covar** consistently outperforms TV, as shown in figure 7, and is competitive with other unsupervised generative model regularization schemes, figure 8, over a range of noise levels and sampling patterns. This is an important contribution, demonstrating a generative model, trained on an MRI dataset, provides an effective prior for image reconstruction with varying levels of information. As part of the ablation study, figures 2 and 3, we demonstrated that models *mean+diag* and *mean+identity* provided some regularization but *mean+covar** provided the most flexibility to fit the data and thus the best results. Searching in the *range* of the generator we saw evidence that the generator was not expressive enough to fit the data and that the reconstructed images did not continue to improve with more or higher-quality data, this matches with similar results in Bora *et al* (2017), Duff *et al* (2021).

The generative regularizers presented are not end-to-end methods and still require the use of an iterative optimization scheme to reconstruct an image given an observation, furthermore, the optimization objective is non-convex. Fortunately, this optimization is practically straightforward as the gradients can be calculated automatically, e.g. using TensorFlow, and the optimization can be initialized with the adjoint reconstruction. The benefit of the unsupervised approach is that retraining is not required for different forward models. Figure 9 shows that the VN end-to-end method (Hammernik *et al* 2018), although provides the best results on the forward model it was trained on, sees its success drop off as you test it on other forward model settings. An additional benefit of generative regularizers is that you can visually inspect the learned image prior and we demonstrate this in figure 4.

Training the Cholesky weights for the precision matrix proved tricky. As shown in figure 2, the structured covariance model has learned some of the residual structure that we expect but at the expense of applying higher variances across the whole image. This may make the structured covariance model too permissive, allowing it to fit noise or artifacts from the measurements. With an improved model, we expect that the regularization parameters in (11) would not be required, they could be set by the hierarchical Bayesian derivation. Future work could consider what additional priors or alternative modeling schemes could help the learning of this structured covariance model, such as described in Dorta *et al* (2018).

Although this work demonstrates a data-driven approach to inverse problems that is flexible and powerful there are several ways the MRI simulation could be made more realistic. As discussed in section 4, our experiments are inspired by the single coil fastMRI challenge which uses magnitude images and, as the name suggests, a single-coil acquisition. In reality, modern MRI scans use a multi-coil acquisition and usually produce complex images with a non-trivial phase. A multi-coil acquisition could be incorporated into our framework with a change of the forward model, e.g. via a SENSE formulation (Pruessmann *et al* 1999). Modeling the structures and correlations between the real and complex parts of an MR image is an interesting open problem and subject to future work. Finally, the radial masks we used to simulate a radial sampling pattern, were just an approximation of the actual MR physics and future work could consider bringing this closer to real-world applications.

As in most machine learning approaches, we have assumed that the images we wish to reconstruct are similar to those in the training dataset used to train the generative model. This is not obvious in medical imaging, where damage, tumors, and illness can lead to different image presentations that may be far from those of healthy volunteers used to create datasets. Modeling this, e.g. by sparse deviations (Dhar *et al* 2018, Duff *et al* 2021) away from the range of a generative model, is an interesting area of future development.

Another consideration for future work is to consider going beyond the MAP estimate and considering quantifying uncertainty in the reconstructed image. One possible avenue would be to produce a range of reconstructions by sampling from the learned decoder for an optimized value of *z*. Further work would be required to interpret and present these results in a meaningful and theoretically justified way.

7. Conclusion

Generative regularizers provide a bridge between black-box deep learning methods and traditional variational regularization techniques for image reconstruction. We propose an adaptive generative regularizer that models image correlations with a structured noise network. The proposed method is trained unsupervised, using only high-quality reconstructions and is thus adaptable to different noise and k-space sampling levels. Our results show that generative regularizers are most effective when the underlying generative model outputs both an image but also a non-trivial covariance matrix for each point in the latent space. The covariance provides a learned metric that guides where the reconstruction can or cannot vary from the learned generative model. We demonstrated the success of this approach through comparisons with other unsupervised and supervised approaches.

Acknowledgments

MD is supported by a scholarship from the EPSRC Centre for Doctoral Training in Statistical Applied Mathematics at Bath (SAMBa), under the project EP/L015684/1. MJE acknowledges support from the EPSRC (EP/S026045/1, EP/T026693/1, EP/V026259/1) and the Leverhulme Trust (ECF-2019-478). NDFC acknowledges support from the EPSRC CAMERA Research Centre (EP/M023281/1 and EP/T022523/1) and the Royal Society.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://doi.org/ https://fastmri.med.nyu.edu/.

ORCID iDs

M A G Duff [®] https://orcid.org/0000-0003-1014-3147 I J A Simpson [®] https://orcid.org/0000-0001-5605-6626 M J Ehrhardt [®] https://orcid.org/0000-0001-8523-353X N D F Campbell [®] https://orcid.org/0000-0003-2130-4903

References

Adler J, Kohr H and Öktem O 2017 Operator Discretization Library ODL (https://doi.org/10.5281/zenodo.1442734)

Ahmad R, Bouman C A, Buzzard G T, Chan S, Liu S, Reehorst E T and Schniter P 2020 Plug-and-play methods for magnetic resonance imaging: using denoisers for image recovery *IEEE Signal Process Mag.* 37 105–16

Antun V, Renna F, Poon C, Adcock B and Hansen A C 2020 On instabilities of deep learning in image reconstruction and the potential costs of AI *Proc. Natl Acad. Sci.* 117 30088–95

Arora S, Risteski A and Zhang Y 2018 Do GANs learn the distribution? Some theory and empirics ICLR 1–16

Arridge S, Maass P, Öktem O and Schönlieb C B 2019 Solving inverse problems using data-driven models *Acta Numer.* **28** 1–174 Bora A, Jalal A, Price E and Dimakis A G 2017 Compressed sensing using generative models *ICML* pp 822–41

Boyd S, Parikh N, Chu E, Peleato B and Eckstein J 2011 Distributed optimization and statistical learning via the alternating direction method

of multipliers *Found. Trends Mach. Learn.* **3** 1–122 Boyd S and Vandenberghe L 2004 *Convex Optimisation* (Cambridge: Cambridge University Press) 9780521833783 (https://doi.org/ 10.1017/CBO9780511804441)

Bungert L and Ehrhardt M J 2020 Robust image reconstruction with misaligned structural information IEEE Access 8 222944–222955

Chambolle A and Pock T 2011 A first-order primal-dual algorithm for convex problems with applications to imaging *J. Math. Imaging Vis.* 40 120–45

Clark A et al 2015 Pillow (Python Imaging Library Fork) (https://doi.org/10.5281/zenodo.8104287)

Dhar M, Grover A and Ermon S 2018 Modeling sparse deviations for compressed sensing using generative models *ICML* 80 (*PMLR*) 1214–23

Dorta G, Vicente S, Agapito L, Campbell N D F and Simpson I 2018 Training vaes under structured residuals arXiv:1804.01050v3

Dorta G, Vicente S, Agapito L, Campbell N D F and Simpson I 2018 Structured uncertainty prediction networks CVPR (Salt Lake City, UT, USA, 2018) 5477–85

Duff M, Campbell N D F and Ehrhardt M J 2021 Regularising inverse problems with generative machine learning models arXiv:2107.11191 Goodfellow I J, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial nets *NeurIPS* 27

Habring A and Holler M 2022 A generative variational model for inverse problems in imaging SIAM J. Math. Data Sci. 4 306–35

Hammernik K, Klatzer T, Kobler E, Recht M P, Sodickson D K, Pock T and Knoll F 2018 Learning a variational network for reconstruction of accelerated MRI data Magn. Reson. Med. 79 3055–71

Hyun C M, Kim H P, Lee S M, Lee S and Seo J K 2018 Deep learning for undersampled MRI reconstruction *Phys. Med. Biol.* **63** (13) 135007 Jalal A, Arvinte M, Daras G, Price E, Dimakis A G and Tamir J I 2021 Robust compressed sensing MRI with deep generative priors *NeurIPS* **34** 14938–54

Kingma D P and Dhariwal P 2018 Glow: generative flow with invertible 1 x 1 convolutions NeurIPS pp 10215-24

Kingma D P and Welling M 2014 Auto-encoding variational Bayes ICLR

Knoll F *et al* 2020 fastMRI: a publicly available raw k-Space and DICOM dataset of knee images for accelerated mr image reconstruction using machine learning *Radiol.: Artif.* Intell. 21 e190007

Knoll F, Bredies K, Pock T and Stollberger R 2011 Second order total generalized variation (TGV) for MRI *Magn. Reson. Med.* **65** 480–91 Lustig M, Donoho D and Pauly J M 2007 Sparse MRI: the application of compressed sensing for rapid MR imaging *Magn. Reson. Med.* **58** 1182–95

Mardani M, Gong E, Cheng J Y, Vasanawala S S, Zaharchuk G, Xing L and Pauly J M 2019 Deep generative adversarial neural networks for compressive sensing MRI *IEEE Trans. Med. Imaging* 38 167–79

Monteiro M, Le Folgoc L, Coelho de Castro D, Pawlowski N, Marques B, Kamnitsas K, van der Wilk M and Glocker B 2020 Stochastic segmentation networks: modelling spatially correlated aleatoric uncertainty *NeurIPS* 33 ed H Larochelle *et al* 12756–67

Narnhofer D, Hammernik K, Knoll F and Pock T 2019 Inverse GANs for accelerated MRI reconstruction *Wavelets and Sparsity XVIII* 11138 (SPIE) 381–92

Oh C, Kim D, Chung J Y, Han Y and Park H W 2018 Eter-net: end to end mr image reconstruction using recurrent neural network *Machine* Learning for Medical Image Reconstruction 11074 (Granada, Spain, September 16, 2018) (Springer) 12–20

Pruessmann K P, Weiger M, Scheidegger M B and Boesiger P 1999 SENSE: sensitivity encoding for fast MRI *Magn. Reson. Med.* 42 952–62 Quan T M, Nguyen-Duc T and Jeong W K 2018 Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic

loss IEEE Trans. Med. Imaging 37 1488–97

Rezende D J and Mohamed S 2015 Variational inference with normalizing flows ICML 37 (PMLR) 1530–38

Roth S and Black M J 2009 Fields of experts Int. J. Comput. Vision 82 205-29

Rudin LI, Osher S and Fatemi E 1992 Nonlinear total variation based noise removal algorithms *Physica* D 60 259–268

Ruthotto L and Haber E 2021 An introduction to deep generative modeling GAMM-Mitteilungen 44 e202100008

Ryu E, Liu J, Wang S, Chen X, Wang Z and Yin W 2019 Plug-and-play methods provably converge with properly trained denoisers *ICML* 97 (PMLR) 5546–57

Tripathi S, Lipton Z C and Nguyen T Q 2018 Correction by projection: denoising images with generative adversarial networks arXiv:1803.04477

Ulyanov D, Vedaldi A and Lempitsky V 2018 Deep image prior CVPR (Salt Lake City, UT, USA, 2018) (IEEE) 9446–54

Wang S, Su Z, Ying L, Peng X, Zhu S, Liang F, Feng D and Liang D 2016 Accelerating magnetic resonance imaging via deep learning *Proc*-Int. Symp. on Biomedical Imaging (Prague, Czech Republic, 2016) 514–7

Xu J and Noo F 2022 Convex optimization algorithms in medical image reconstruction in the age of AI *Phys. Med. Biol.* 67 07TR01 Zbontar J *et al* 2018 fastMRI: an open dataset and benchmarks for accelerated MRI arXiv:1811.08839

Zhang C, Barbano R and Jin B 2021 Conditional variational autoencoder for learned image reconstruction *Computation* **9** 114 Zhu B, Liu J Z, Cauley S F, Rosen B R and Rosen M S 2018 Image reconstruction by domain-transform manifold learning *Nature* **555** 487–492