

Compressed Sensing MRI Reconstruction Regularized by VAEs with Structured Image Covariance

Margaret Duff, Ivor J. A. Simpson, Matthias J. Ehrhardt and Neill D. F. Campbell

Abstract

Learned regularization for MRI reconstruction can 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. One such approach uses generative models, trained on ground-truth images, as priors for inverse problems, penalizing reconstructions far from images the generator can produce. In this work, we utilize variational autoencoders (VAEs) 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. We demonstrate these novel generative regularizers on radially sub-sampled MRI knee measurements from the fastMRI dataset and compare them to other unlearned, unsupervised and supervised methods. 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.

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.

MD and MJE are with the Department of Mathematical Sciences and NDFC is with the Department of Computer Science at the University of Bath, UK, BA2 7AY. IS is with the Department of Informatics, University of Sussex, UK, BN1 9RH.

e-mail: M.A.G.Duff@bath.ac.uk

I. INTRODUCTION

Compressed sensing Magnetic Resonance Imaging (MRI) provides the benefits of MRI with faster acquisition times. 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 an image, $x \in \mathcal{X}$, from observed measurements, $y \in \mathcal{Y}$. The two are related by a linear forward model, $A : \mathcal{X} \rightarrow \mathcal{Y}$ 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. Measurements are incomplete and so multiple solutions may exist and the problem is ill-posed.

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 [22], [24], [36]

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

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

We extend recent work, including [5], [10], [13], [15], [33] that consider the use of a generative machine learning model as part of the regularizer, called *generative regularizers*. A generative model is designed to take a sample from a distribution in a low-dimensional latent space and generate data similar to some training distribution. Common deep generative models include variational autoencoders (VAEs) [21], generative adversarial networks (GANs) [14] and normalizing flows [29]. 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 $z \in \mathcal{Z}$, within the learned latent space, the network outputs a distribution on the image space \mathcal{X} ,

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

for parameterised functions $G_\theta : \mathcal{Z} \rightarrow \mathcal{X}$ and $\Sigma_\theta : \mathcal{Z} \rightarrow \mathbb{R}^{d \times d}$, $d := \dim(\mathcal{X})$. 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 are harder to model accurately. Similarly, as images are commonly piece-wise 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.* [32]), 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.

a) Contributions: We propose to solve compressed MRI reconstruction with highly adaptive generative regularizers. Edge and correlations are modeled with a structured covariance network. We show the strength of this model on real data obtained from the fastMRI dataset. Contributions include:

- A demonstration of the structured covariance model introduced by Dorta *et al.* [12] on the fastMRI dataset of knee MR images.
- An extension of the denoising example from Dorta *et al.* [12] 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, $\Sigma_\theta(z)$: 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) [31] and least squares methods. Also with two unsupervised methods: a deep image prior approach with a pre-trained generator from Narnhofer *et al.* [26] and the original Compressed Sensing using Generative Models work by Bora *et al.* [5]. Finally, with a state-of-the-art, learned, end-to-end method, variational networks [16]. We demonstrate that our method is competitive with the state-of-the-art neural networks, yet offers superior generalization to other noise statistics and sampling patterns.

II. RELATED WORK

Deep learning approaches to image reconstruction in inverse problems is a growing field [4]. There are several supervised deep learning approaches [16], [18], [25], [27], [28], [35], [38] 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 [2]. In contrast, deep image priors [26], [34], 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 [5], [10], [13], [15], [33], and we extend it with the addition of a structured image covariance network.

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 [14] is, that they do not generate across the range of images they were trained on. This failure can be subtle [3] 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 are able to 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 [19], [20], because of the regularizing effect of a lower dimensional latent space in the VAE. Dorta *et al.* [12] 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.

III. METHOD

We build a probabilistic model for the reconstructed image, x , given an observation, y . First, 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

$$p(x, z|y) \propto p(y|x, z) p_{G,\Sigma}(x|z) p_Z(z) \quad (3)$$

$$= p(y|x) p_{G,\Sigma}(x|z) p_Z(z). \quad (4)$$

Ideally, we would seek to marginalize out the latent vector, z , however this integral is intractable, except by expensive sampling; instead, we take a maximum a posteriori (MAP) estimate. By taking logarithms, maximizing (4) with respect to x and z is equivalent to variational regularization (1) with

$$d(Ax, y) = \frac{1}{2\gamma^2} \|Ax - y\|_2^2 \quad (5)$$

and

$$\mathcal{R}(x) = \min_z \left(\log(|\Sigma_\theta(z)|) + \frac{\|x - G_\theta(z)\|_{\Sigma_\theta(z)}^2}{2} + \frac{\|z\|_2^2}{2} \right), \quad (6)$$

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 (5) ensures that x explains the observation y , while the second term in (6) constrains x to be close to images in the range of the generator.

a) 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 [21]. The intractable distribution over the latent space, $p_{G,\Sigma}(z|x; \theta)$, is approximated by an encoder

$$q(z|x; \psi) = \mathcal{N}(z; \mu_\psi(x), \text{diag}(\sigma_\psi^2(x))) =: \mathcal{N}_{x,\psi}(z) \quad (7)$$

with neural networks μ_ψ, σ_ψ^2 , parameterized with weights, ψ . Following the derivation in [21], training a VAE becomes a minimization with respect to ψ and θ of

$$\mathbb{E}_{x \sim p_{\text{Im}}} \mathbb{E}_{z \sim \mathcal{N}_{x,\psi}} (\text{nll}(x, G_\theta(z), \Sigma_\theta(z)) + d_{KL}(\mathcal{N}_{x,\psi} | p_Z)) \quad (8)$$

where we write the negative log-likelihood as

$$\text{nll}(x, G, \Sigma) = \log(|\Sigma|) + \frac{1}{2} \|x - G\|_{\Sigma}^2 + \text{constants} \quad (9)$$

and the expectation over x is calculated empirically over the training set. The expectation over z is also calculated empirically and for more details see the original paper [21].

b) 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.* [12]. 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 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 pixels i, j are not in a small neighborhood. 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 complex dependencies across images. The local sparsity structure is also amenable to parallelization on the GPU and for more details see [11]. A pictorial view of the full network including an example sparse Cholesky matrix is given in figure 1.

c) Objective functions: This sparse Cholesky decomposition of the precision makes predicting a dense covariance matrix computationally feasible. The negative log-likelihood from (8) becomes

$$\text{nll}(x, G, \Sigma) = -2 \sum_{i=1}^d \log(L_{ii}) + \frac{1}{2} \|L(x - G)\|_2^2 \quad (10)$$

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 (6), giving

$$\mathcal{R}(x) = \min_z \left(\text{nll}(x, G_{\theta}(z), \Sigma_{\theta}(z)) + \frac{1}{2} \|z\|_2^2 \right). \quad (11)$$

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 $u \sim \mathcal{N}(0, I)$.

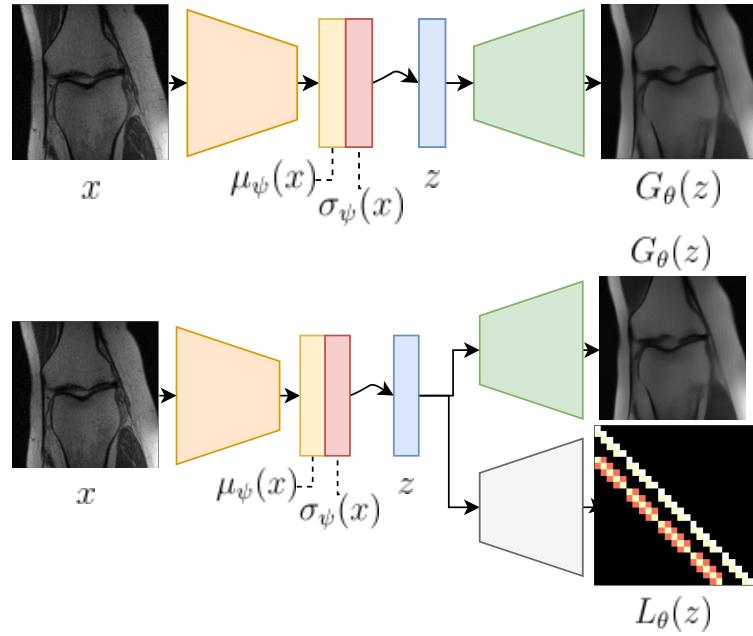


Fig. 1: Comparison of the standard VAE (top) and the Structured Uncertainty Prediction Networks as developed by Dorta *et al.* [12] (bottom). The VAE has an encoder network that outputs a distribution, $\mathcal{N}(z; \mu_\psi(x), \text{diag}(\sigma_\psi^2(x)))$, which is then sampled from to get a latent vector, z , which is passed to the generator network, $G_\theta(z)$. In addition to the usual generator, the proposed network includes a network that takes latent vectors, z , and outputs the weights of a sparse lower triangular matrix, $L_\theta(z)$. This matrix corresponds to the Cholesky decomposition of the inverse covariance for the generated image.

We note that the minimization problems in (8) and (1) with (12) 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 cap the size of the diagonals of L_θ using a scaled tanh activation to deal with this. Investigating other priors on L_θ is an interesting area for future work.

IV. EXPERIMENTS

a) Dataset: The NYU fastMRI knee dataset contains, amongst other data, 796 fully sampled knee MRI magnitude volumes [23], [37], 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. We extract 3,872 training and 800 test ground truth images from this dataset, selecting images from near the center of the knee, resize the images, using the Python imaging library [9] 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.

b) 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 [37]. The magnitude ground truth images are Fourier transformed and subsampled. To this end, a mask selects the points in k-space lying on a set number of radial spokes. The more spokes we measure, the more data is observed. Radial subsampling was chosen, as a common subsampling method, which is illustrative for compressed sensing [17]. We use the operator discretization library (ODL) [1] in Python and take the same radial subsampling MRI operators as in [7].

c) 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 up-sampling 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 8x8 images with 16 channels before a resnet block without resizing to give 8x8 images with 256 channels, then four up-sampling blocks are applied giving image sizes 16x16, 32x32, 64x64 and 128x128 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 128x128 image with 5 channels, from which the sparse Cholesky matrix is formulated, based on code from Dorta *et al.* [12]. 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.

d) Model Training: Training the VAE with structured covariance is done in a two-stage process. First the generated mean, G_θ , and the encoder, $\mathcal{N}_{x,\psi}$, are trained with a standard VAE loss *i.e.* $\Sigma_\theta = \rho I$ where ρ is fixed [21]. The weights for the mean and encoder are then fixed before the covariance model is trained using (8). 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 hours to train the means and then 24hrs and 30hrs for the diagonal and structured covariance models. The code will be available on publication.

e) 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, $\mathcal{N}_{x,\psi}$ from the *mean+identity* model and then optimize (8) 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.

f) 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 (6)

$$\mathcal{R}(x) = \min_z \lambda \left(\text{nll}(x, G_\theta(z), \Sigma_\theta(z)) + \frac{\mu}{2} \|z\|_2^2 \right). \quad (12)$$

The addition of the two regularization parameters λ and μ is in recognition that this is a non-convex problem. We use alternating gradient descent with backtracking line search (see *e.g.* Algorithm 9.2 of [6]) 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).

g) Unlearned method comparisons: We compare against Total Variation (TV) regularization [31] implemented using the Primal-Dual Hybrid Gradient method [8], with 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.

h) 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.* [5], minimizing with respect to z the objective

$$L(z) = \frac{1}{2} \|AG(z) - y\|_2^2 + \mu\|z\|_2^2, \quad (13)$$

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 [26] which takes a trained generator but then, after observing data, tweaks the weights of the network, optimizing $z^*, \theta^* \in \arg \min \|AG_\theta(z) - y\|_2^2$, 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.

i) Supervised end-to-end method comparison: Finally, we also demonstrate comparisons with a variational network (VN) [16]. This is a end-to-end approach which takes (6) with $d(Ax, y) = \frac{\lambda}{2}\|Ax - y\|_2^2$ and treats the regulariser \mathcal{R} as some unknown function. Optimizing (6) by gradient descent will lead to updates of the form

$$x_{t+1} = x_t - \alpha_t \lambda A^T (Ax_t - y) + \nabla \mathcal{R}(x_t), \quad (14)$$

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 [30]. 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.

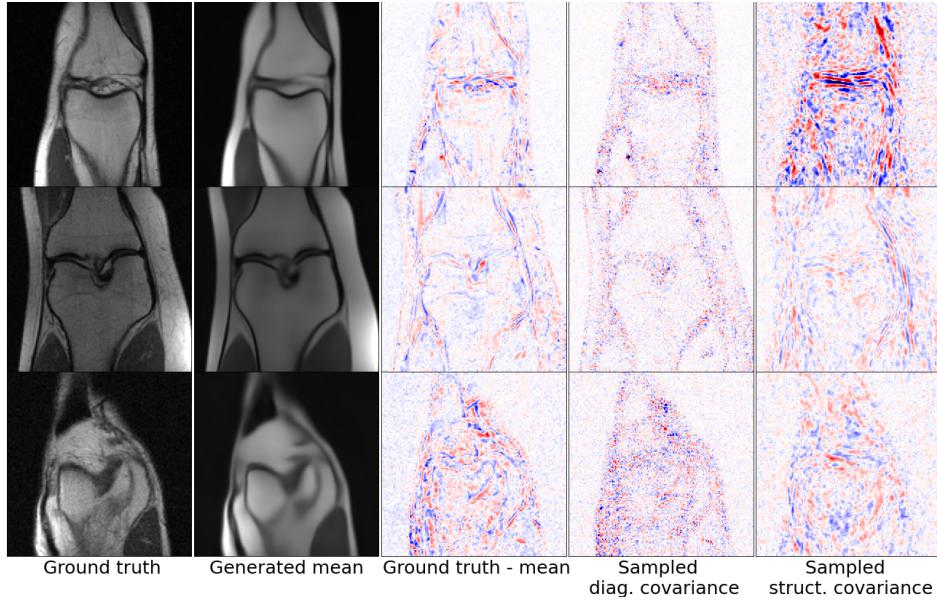


Fig. 2: Comparing the covariance models *mean+diagonal* and *mean+covar**. The first column gives the ground truth, and the second column a reconstruction in the *range* of the VAE generator. The third gives the residual which can be considered as one sample of a zero mean Gaussian distribution with unknown covariance. The final two columns give single residual samples from our learned covariances *mean+diagonal* and *mean+covar** models. Columns are rescaled for better visualization.

V. RESULTS

a) 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. The first three columns compare the original images, the generated mean images and the residual. This residual can be considered as one sample of a zero mean Gaussian distribution with unknown covariance. The final two images show single residual samples from our learned covariances. 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 are able to 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. It is clear

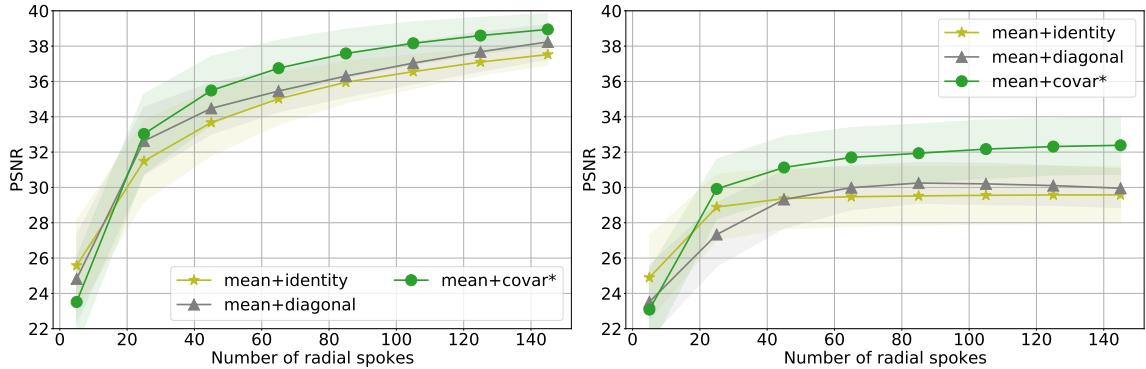


Fig. 3: Ablation study comparing generative regularisers based on VAEs with different noise models *mean+identity*, *mean+diagonal* and *mean+covar**. The plots show PSNR values average over reconstructions of 50 test images. The mean is given by the solid line and standard deviation in the shaded box. The x-axis measures the number of radial spokes in the sampling pattern and the measured data was corrupted with additive Gaussian noise of standard deviation 0.0125 on the left and 0.05 on the right.

that the results for *mean+covar** gives 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.

b) Prior introspection: We are able to 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. Positive correlations are given in red, and negative correlations are in blue. We can see that the structure of the edges is present, and that, despite the local structure of the precision matrix, longer-range correlations have been learned.

c) Comparison with unlearned methods: Figure 7 shows comparisons 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. We see the results

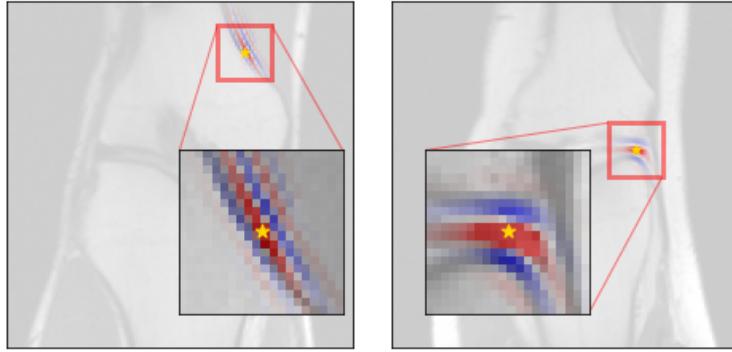


Fig. 4: Visualization of learned covariances between example pixels (yellow stars) and other pixel locations for the *mean+covar* model. Red indicates a high positive correlation, and blue is a strong negative correlation.

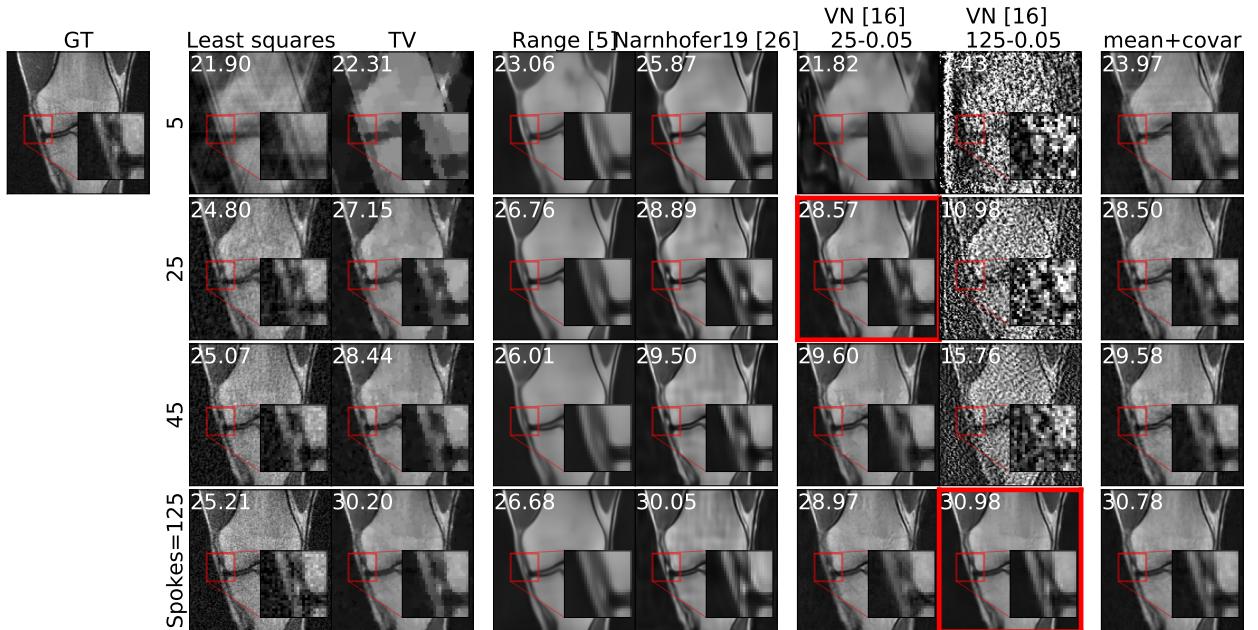


Fig. 5: Example reconstructions of a test image for different amounts of measured data: 5, 25, 45 and 125 radial spokes, all with additive Gaussian noise with standard deviation 0.05. The columns give different reconstruction methods. The PSNR values are added in white and the red boxes indicate the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

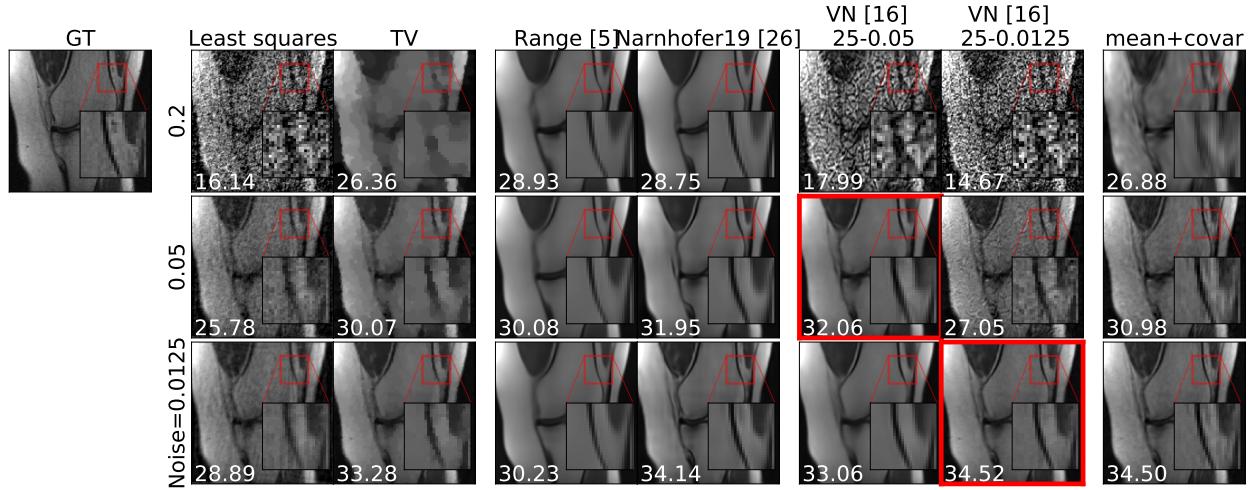


Fig. 6: Example reconstructions of a test image with additive noise of different standard deviations: 0.2, 0.05 and 0.0125, all with a sampling pattern with 25 radial spokes. The layout is as in figure 5.

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 seconds for reconstructing one image using least squares, and 6.4 seconds using TV, for one choice of regularization parameter. For one choice of regularization parameter, *mean+covar** took 78.2 seconds.

d) Comparison with other data-driven priors: Comparisons with *Narnhofer19* and *range* [5] 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 don't see significant improvement. The example image reconstructions reflect this; they show overly smoothed images that, although have the right structure, don't 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 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

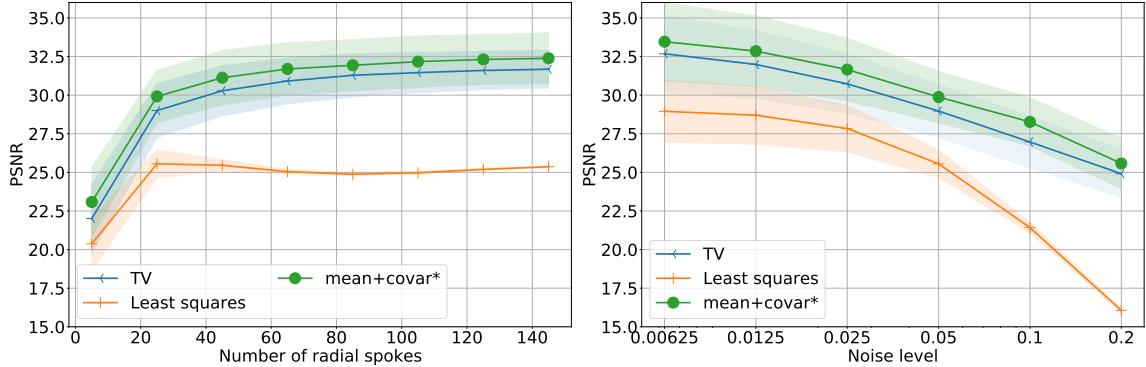


Fig. 7: Comparison of the *mean+covar** method with the training-free methods TV and least squares. The plots show the mean PSNR values for reconstructions averaged over 50 test images, with the standard deviations given by the shaded area. The left plot shows results from measured data with additive Gaussian noise of standard deviation 0.05 and differing numbers of measured radial spokes. The right plot gives results where 25 radial spokes are used for each reconstruction but the standard deviation of the added noise varies.

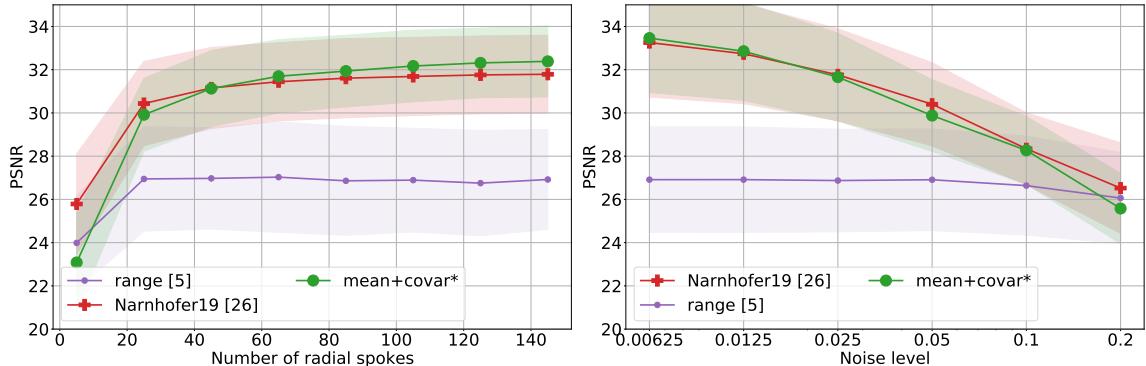


Fig. 8: Comparison of the *mean+covar** method with *range* [5] and *Narnhofer19* [26], generative regularizer approaches. The experimental setup is as figure 7.

in PSNR values to our *mean+covar**. In the images, there is some evidence of smoothing of details *e.g.* in figure 6. Our un-optimized implementation took 106.7 seconds for *Narnhofer19* and 23.9 seconds for the *range* [5] method.

e) Comparison with supervised end-to-end methods: We now compare *mean+covar** with variational networks [16], trained end-to-end on 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

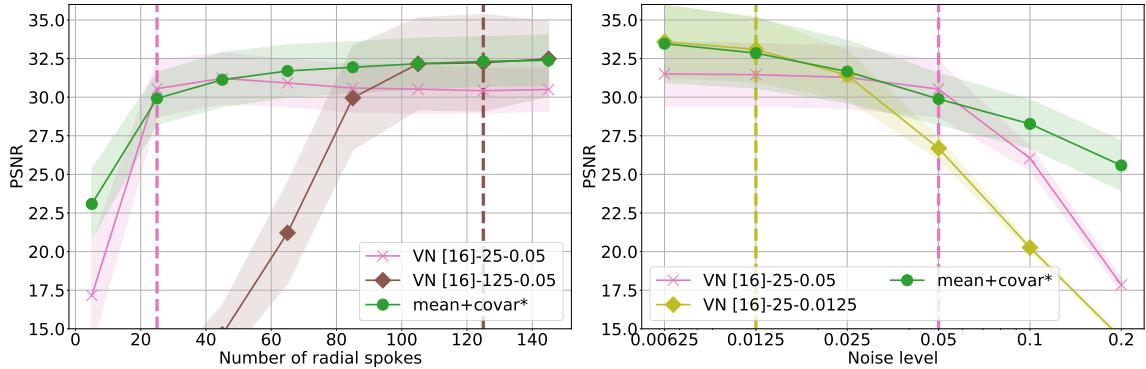


Fig. 9: Comparison of the unsupervised *mean+covar** method with the supervised variational networks [16]. The experimental setup is as figure 7 but in addition, the vertical lines depict the experimental settings the variational networks were trained on.

with the trained for setting highlighted in red. We note that the variational networks 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 variational networks are quick to implement once trained, taking 3.4 seconds for one image

f) 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 variational network 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.

VI. DISCUSSION

The numerical results show that the generative regularizer *mean+covar** consistently outperforms TV, figure 7, and is competitive with other unsupervised generative model regularization schemes,

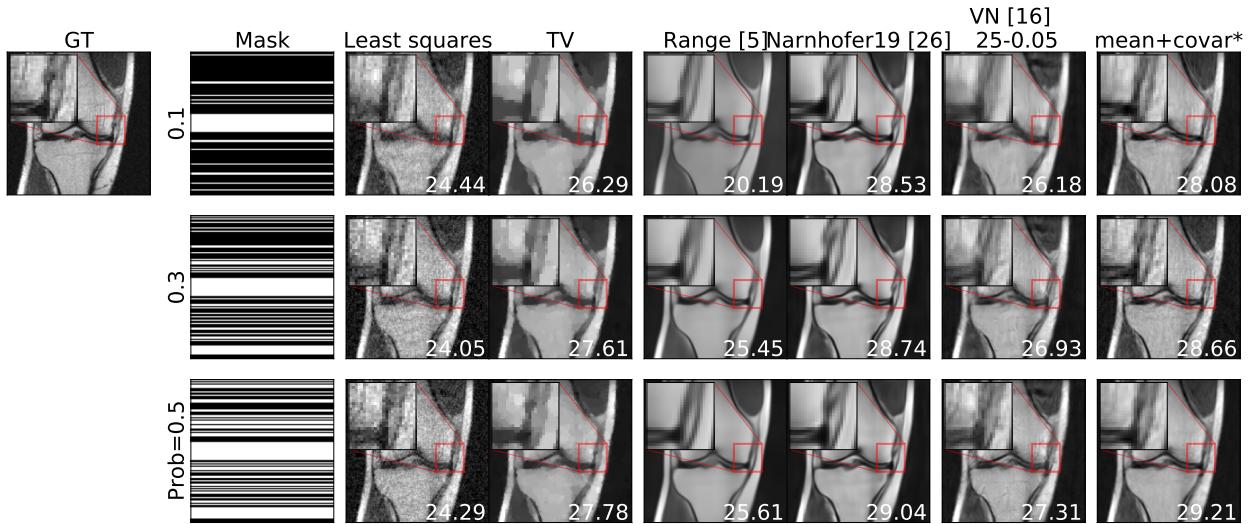


Fig. 10: Example reconstructions of an image using measured data taken with different horizontal random sampling unseen in the training of the variational networks. The first column gives the k-space mask used to sample the data. The remaining columns give reconstructions for the different methods.

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. 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 [5], [13].

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. However, 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 variational network end-to-end method [16], 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 (12) 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 [11]

In this work, 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 [10], [13] away from the range of a generative model, is an interesting area of future development. To bring this work closer to real applications, future work could also consider complex images and parallel MRI.

VII. CONCLUSION

Generative regularizers provide a bridge between black-box deep learning methods and traditional variational regularization techniques for image reconstruction. We propose a highly 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.

REFERENCES

- [1] Adler, J., Kohr, H., Öktem, O.: Operator discretization library (odl) (2017)

- [2] Antun, V., Renna, F., Poon, C., Adcock, B., Hansen, A.C.: On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences* p. 201907377 (2020)
- [3] Arora, S., Risteski, A., Zhang, Y.: Do gans learn the distribution? some theory and empirics. *ICLR* pp. 1–16 (2018)
- [4] Arridge, S., Maass, P., Öktem, O., Schönlieb, C.B.: Solving inverse problems using data-driven models. *Acta Numerica* **28**, 1–174 (2019)
- [5] Bora, A., Jalal, A., Price, E., Dimakis, A.G.: Compressed sensing using generative models. pp. 822–841 (2017)
- [6] Boyd, S., Vandenberghe, L.: *Convex Optimization* (2004)
- [7] Bungert, L., Ehrhardt, M.J.: Robust image reconstruction with misaligned structural information. *IEEE Access* **8**, 222944–222955 (2020)
- [8] Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision* **40**, 120–145 (2011)
- [9] Clark, A.: Pillow (python imaging library fork) (2015)
- [10] Dhar, M., Grover, A., Ermon, S.: Modeling sparse deviations for compressed sensing using generative models. vol. 3, pp. 1990–2005 (2018)
- [11] Dorta, G., Vicente, S., Agapito, L., Campbell, N.D.F., Simpson, I.: Training vaes under structured residuals. *arXiv preprint arXiv:1804.01050* (2018)
- [12] Dorta, G., Vicente, S., Agapito, L., Campbell, N.D., Simpson, I.: Structured uncertainty prediction networks. pp. 5477–5485 (2018)
- [13] Duff, M., Campbell, N.D.F., Ehrhardt, M.J.: Regularising inverse problems with generative machine learning models. *ArXiv Preprint* (2021)
- [14] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *NeurIPS* pp. 2672–2680 (2014)
- [15] Habring, A., Holler, M.: A generative variational model for inverse problems in imaging. *SIAM Journal on Mathematics of Data Science* **4**, 306–335 (2022)
- [16] Hammerik, 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**, 3055–3071 (2018)
- [17] Hollingsworth, K.G.: Reducing acquisition time in clinical mri by data undersampling and compressed sensing reconstruction. *Physics in Medicine and Biology* **60**, R297–R322 (2015)
- [18] Hyun, C.M., Kim, H.P., Lee, S.M., Lee, S., Seo, J.K.: Deep learning for undersampled mri reconstruction. *Physics in Medicine and Biology* **63** (2018)
- [19] Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.I.: Robust compressed sensing mri with deep generative priors. *NeurIPS* (2021)
- [20] Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1×1 convolutions. pp. 10215–10224 (2018)
- [21] Kingma, D.P., Welling, M.: Auto-encoding variational bayes (2014)
- [22] Knoll, F., Bredies, K., Pock, T., Stollberger, R.: Second order total generalized variation (tgv) for mri. *Magnetic Resonance in Medicine* **65**, 480–491 (2011)
- [23] Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdzal, 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**, e190007 (2020)

- [24] Lustig, M., Donoho, D., Pauly, J.M.: Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine* **58**, 1182–1195 (2007)
- [25] Mardani, M., Gong, E., Cheng, J.Y., Vasanawala, S.S., Zaharchuk, G., Xing, L., Pauly, J.M.: Deep generative adversarial neural networks for compressive sensing mri. *IEEE Transactions on Medical Imaging* **38**, 167–179 (2019)
- [26] Narnhofer, D., Hammernik, K., Knoll, F., Pock, T.: Inverse gans for accelerated mri reconstruction. p. 45. SPIE-Intl Soc Optical Eng (2019)
- [27] Oh, C., Kim, D., Chung, J.Y., Han, Y., Park, H.W.: Eter-net: End to end mr image reconstruction using recurrent neural network (2018)
- [28] Quan, T.M., Nguyen-Duc, T., Jeong, W.K.: Compressed sensing mri reconstruction using a generative adversarial network with a cyclic loss. *IEEE Transactions on Medical Imaging* **37**, 1488–1497 (2018)
- [29] Rezende, D.J., Mohamed, S.: Variational inference with normalizing flows. vol. 2, pp. 1530–1538 (2015)
- [30] Roth, S., Black, M.J.: Fields of experts. *International Journal of Computer Vision* **82**, 205–229 (2009)
- [31] Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60**, 259–268 (1992)
- [32] Ruthotto, L., Haber, E.: An introduction to deep generative modeling. *GAMM-Mitteilungen* **44**, e202100008 (3 2021). <https://doi.org/10.1002/gamm.202100008>
- [33] Tripathi, S., Lipton, Z.C., Nguyen, T.Q.: Correction by projection: Denoising images with generative adversarial networks. ArXiv Preprint (2018)
- [34] Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. pp. 1867–1888 (2020)
- [35] Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., Feng, D., Liang, D.: Accelerating magnetic resonance imaging via deep learning. vol. 2016-June, pp. 514–517. IEEE Computer Society (2016)
- [36] Xu, J., Noo, F.: Convex optimization algorithms in medical image reconstruction in the age of ai. *Physics in Medicine and Biology* **67**, 07TR01 (2022)
- [37] Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdza, M., Romero, A., Rabbat, M., Vincent, P., Yakubova, N., Pinkerton, J., Wang, D., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastmri: An open dataset and benchmarks for accelerated mri. ArXiv Preprint (2018)
- [38] Zhu, B., Liu, J.Z., Cauley, S.F., Rosen, B.R., Rosen, M.S.: Image reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492 (2018)