

SUREMAP: PREDICTING UNCERTAINTY IN CNN-BASED IMAGE RECONSTRUCTIONS USING STEIN’S UNBIASED RISK ESTIMATE

Ruangrawee Kitichotkul, Christopher A. Metzler, Frank Ong, Gordon Wetzstein

Department of Electrical Engineering at Stanford University
`{rk22, cmetzler}@stanford.edu`

ABSTRACT

Convolutional neural networks (CNN) have emerged as a powerful tool for solving computational imaging reconstruction problems. However, CNNs are generally difficult-to-understand black-boxes. Accordingly, it is challenging to know when they will work and, more importantly, when they will fail. This limitation is a major barrier to their use in safety-critical applications like medical imaging: Is that blob in the reconstruction an artifact or a tumor?

In this work we use Stein’s unbiased risk estimate (SURE) to develop per-pixel confidence intervals, in the form of heatmaps, for compressive sensing reconstruction using the approximate message passing (AMP) framework with CNN-based denoisers. These heatmaps tell end-users how much to trust an image formed by a CNN, which could greatly improve the utility of CNNs in various computational imaging applications.

Index Terms— Compressive Sensing, Approximate Message Passing, CNN, MRI

1. INTRODUCTION

Computational imaging (CI) systems, like magnetic resonance imaging (MRI), can generally be described by the equation

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}, \quad (1)$$

where $\mathbf{y} \in \mathbb{C}^m$ denotes the measurements, $\mathbf{A} \in \mathbb{C}^{m \times n}$ models the linear measurement operator/matrix, $\mathbf{x} \in \mathbb{C}^n$ is the vectorized latent image, and $\boldsymbol{\eta} \in \mathbb{C}^m$ is additive noise. The goal of a computational imaging reconstruction algorithm is to reconstruct \mathbf{y} from \mathbf{x} .

When $m \ll n$ the reconstruction problem is underdetermined, and is known as compressive sensing (CS). CS reconstruction algorithms impose a prior, implicitly or explicitly, to form a reconstruction, $\hat{\mathbf{x}}$, of \mathbf{x} from \mathbf{y} . While historically this

prior was sparsity in some basis [1], the sparsity model has largely been superseded: Modern “hand-designed” methods achieve far better performance by imposing more elaborate priors, such as non-local self-similarity [2]. Meanwhile, learning-based methods, which impose priors with convolutional neural networks (CNNs), offer better performance still [3].

CNNs learn priors from vast quantities of training data, which they use to tune thousands to millions of parameters. In general, it is unclear how each parameter contributes to the performance of the algorithm and it is difficult to know if and when a CNN-based method will successfully reconstruct an image.

Expected mean squared error (MSE), i.e. risk, is the gold standard for evaluating a CS reconstruction algorithm. However, in general computing the risk requires access to the ground truth image – which defeats the point of reconstruction in the first place.

In this work, we demonstrate that when used in conjunction with the approximate message passing (AMP) framework [4], which decouples the CS reconstruction problem into a series of additive white Gaussian noise (AWGN) denoising problems, one can accurately calculate the expected per-pixel MSE associated with CS reconstruction using Stein’s unbiased risk estimate (SURE) [5]. Consequently, we can generate heatmaps of low-pass filtered per-pixel MSE estimates *without requiring access to the latent image*. We also apply this framework to the Variable Density AMP (VDAMP) algorithm [6], an MRI reconstruction algorithm which decouples the problem into a series of additive colored Gaussian noise denoising problems. These uncertainty heatmaps could inform end-users about the reliability of image reconstructions and could also serve as supplementary information for an artifact-removal algorithm [7] or to guide an adaptive sampling strategy [8].

2. RELATED WORK

Researchers have long sought to qualify the uncertainty associated CNN-based reconstructions. The importance of this problem was recently highlighted in [9, 10], where the authors showed how slight perturbations to a compressively sampled MRI signal can lead to vastly different, but still plausible

R.K. was supported by the Stanford Research Experience for Undergraduates (REU) program. C.M. was supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at Stanford University administered by Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence (ODN). G.W. was supported by an NSF CAREER Award (IIS 1553333), a Sloan Fellowship, and a PECASE by the ARL.

looking, reconstructions.

If one assumes the latent image lies in the range of a generative network, one can use RIP-like conditions to guarantee recovery when the network is sufficiently expansive [11, 12] or invertible [13]. By looking at the distribution of an invertible network’s latent variables, one can then estimate the uncertainty associated with a reconstruction [14].

Alternatively, when dealing with probabilistic neural networks, as exemplified by variational autoencoders [15], one can sample from $p(\hat{\mathbf{x}}|\mathbf{y})$, and thereby reason about the variance, but not the bias, associated with the reconstruction $\hat{\mathbf{x}}$ [16]. Similarly, bootstrap and jackknife resampling methods [17] as well as a combination of variational dropout and input-dependent noise models [18] can be used to estimate the variance of a reconstruction. One can even train a CNN to identify motion artifacts [19].

The majority of these method however can only characterize the variance associated with the reconstruction. They do not accurately predict the mean squared error, which is effected by bias as well.

Recently, Edupuganti et al. predicted the per-pixel mean-squared error associated with reconstructed MRI images using SURE [20]. However, in order to apply SURE, their method assumes that the difference between the true signal \mathbf{x} and an initial estimate, formed with density compensated least squares (DCLS), follows a distribution that is both Gaussian and *white*. As demonstrated in Figure 1, the latter assumption does not hold in practice: The “effective noise”, i.e., the difference between the estimate and the truth, demonstrates obvious structure when represented in the wavelet domain. These correlations invalidate the standard SURE approach, which applies only to i.i.d. Gaussian noise.

3. BACKGROUND

3.1. Stein’s Unbiased Risk Estimate

SURE was first developed by its namesake several decades ago [5]. Given a noisy signal $\mathbf{y} = \mathbf{x} + \eta$, where η follows a Gaussian distribution with known covariance $\sigma^2 \mathbf{I}$, SURE states that one can form an unbiased estimate of the mean squared error (MSE), $\frac{1}{n} \|\mathbf{x} - f(\mathbf{y})\|^2$, via the expression

$$S(\mathbf{y}, f(\mathbf{y})) = \frac{1}{n} \|\mathbf{y} - f(\mathbf{y})\|^2 - \sigma^2 + \frac{2\sigma^2}{n} \text{div}_{\mathbf{y}}(f(\mathbf{y})), \quad (2)$$

where $\text{div}_{\mathbf{y}}(f(\mathbf{y}))$ denotes its divergence, defined as

$$\text{div}_{\mathbf{y}}(f(\mathbf{y})) = \sum_{n=1}^N \frac{\partial f_n(\mathbf{y})}{\partial y_n}. \quad (3)$$

Since its introduction, SURE has been used extensively to tune algorithms. It is at the heart of the well-known SURE-shrink denoising algorithm [21] and has been used extensively for tuning the parameters within various iterative reconstruction algorithms as well [22, 23, 6, 24]. SURE has also been combined with deep learning to train CNNs without ground truth data [25, 26, 27] and been used to predict the error asso-




Fig. 1: DCLS effective noise. An illustration of the effective noise in the wavelet domain following a density compensated least squares reconstruction of a compressively sampled MRI image: The noise does not follow an i.i.d. Gaussian distribution.




Fig. 2: VDAMP effective noise. An illustration of the effective noise in the wavelet domain within an iteration of VDAMP while reconstructing a compressively sampled MRI image: The effective noise is approximately i.i.d. within each wavelet subband.

ciated with a denoising algorithm’s reconstruction [28].

3.2. Approximate Message Passing

Approximate message passing (AMP), presented in Algorithm 1, is a simple iterative algorithm for reconstructing a signal from i.i.d. Gaussian measurements [4], i.e., $\mathbf{A}_{i,j} \sim \mathcal{N}(0, 1)$ for all i, j . AMP resembles a projected gradient descent algorithm but comes with an additional term, $\frac{1}{m} \text{div}_{\mathbf{r}_t}(\mathbf{x}_{t+1}) \mathbf{z}_t$, known as the Onsager correction. The Onsager correction ensures that at every iteration the effective noise, that is the difference between \mathbf{r}_t and the ground truth signal \mathbf{x} , follows a white Gaussian distribution with variance $\hat{\sigma}^2$.

Algorithm 1: AMP

Input : Observation $\mathbf{y} \in \mathbb{R}^m$, Denoiser $f(\cdot)$,
Measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$

Output: Reconstructed image $\hat{\mathbf{x}}$

Initialize $\mathbf{x}_0 = \mathbf{0}_n$, $\mathbf{z}_0 = \mathbf{y}$;

for $t = 0, \dots, T - 1$ **do**

$\mathbf{r}_t = \mathbf{x}_t + \mathbf{A}^T \mathbf{z}_t$
 $\hat{\sigma}_t = \|\mathbf{z}_t\|_2 / \sqrt{m}$
 $\mathbf{x}_{t+1} = f(\mathbf{r}_t; \hat{\sigma}_t)$
 $\mathbf{z}_{t+1} = \mathbf{y} - \mathbf{A} \mathbf{x}_{t+1} + \frac{1}{m} \text{div}_{\mathbf{r}_t}(\mathbf{x}_{t+1}) \mathbf{z}_t$

end

return \mathbf{x}_T

Variable Density AMP (VDAMP) is a recent extension

to AMP designed to solve the CS reconstruction problem when dealing with variable density sampled Fourier measurements [6]. Through multiscale updates in the wavelet domain, it ensures that the effective noise follows a colored Gaussian distribution with a *known covariance matrix*. This covariance matrix is diagonal when represented in the wavelet domain. Figure 2 illustrates the empirical distribution of the effective noise associated with VDAMP.

While originally designed with simple, soft-thresholding based denoisers $f(\cdot)$, both AMP and VDAMP can be extended to incorporate more advanced denoisers, such as CNNs. The resulting Denoising-based AMP (D-AMP) and VDAMP (D-VDAMP) algorithms offer state-of-the-art performance when dealing with i.i.d. Gaussian and variable density sampled Fourier measurements, respectively [29, 30, 31].

4. METHOD

In this work, we combined SURE with the denoising-based version of AMP and VDAMP to generate per-pixel mean-squared error estimates associated with reconstructions of compressively sampled images.

4.1. Uncertainty Quantification for D-AMP

At each iteration, D-AMP solves a denoising problem described by

$$\mathbf{r}_t = \mathbf{x} + \eta_t; \mathbf{x}, \mathbf{r}_t \in \mathbb{R}^n; \eta_t \sim \mathcal{N}(0, \sigma_\eta^2 \mathbf{I}_n), \quad (4)$$

where η_t is the noise at the t -th iteration, and \mathbf{r}_t is the corresponding noisy image. The final estimate formed by T iterations of AMP is $\hat{\mathbf{x}} = f(\mathbf{r}_T)$. Because this is the output of a simple AWGN denoising problem, SURE can be used to estimate the mean squared error associated with this reconstruction.

When a closed form expression for $\text{div}_{\mathbf{y}}(f_{\theta}(\mathbf{r}_T))$ is not available, it can be estimated with the following Monte-Carlo estimate [32]

$$\text{div}_{\mathbf{r}_T}(\hat{\mathbf{x}}) \approx \frac{1}{K} \sum_{k=1}^K \frac{1}{\epsilon} \mathbf{b}_k^T (f(\mathbf{r}_T + \epsilon \mathbf{b}_k) - f(\mathbf{r}_T)), \quad (5)$$

where $\mathbf{b}_k \sim \mathcal{N}(0, \mathbf{I}_n)$ and $\epsilon \in \mathbb{R}$ is a small number, chosen to be $\frac{\max(\mathbf{r}_T)}{1000}$ in this work, and K is the number of Monte-Carlo samples used in the approximation. To generate a per-pixel SURE heatmap, we compute SURE for overlapping patches of the reconstruction and average the result.

4.2. Uncertainty Quantification for D-VDAMP

At each iteration, D-VDAMP solves a denoising problem described by

$$\mathbf{r}_t = \mathbf{x} + \eta_t; \mathbf{x}, \mathbf{r}_t \in \mathbb{R}^n; \eta_t \sim \mathcal{CN}(0, \Psi^t \text{diag}(\tau_t) \Psi), \quad (6)$$

where $\mathcal{CN}(0, \Psi^t \text{diag}(\tau_t) \Psi)$ denotes a circular Gaussian distribution with independent real and imaginary parts, each of which has mean 0 and the covariance matrix $\frac{1}{2} \Psi^t \text{diag}(\tau_t) \Psi$, and Ψ denotes the wavelet transform matrix. (We use a four-level 2-D Haar transform throughout this paper.) As before,

the final estimate associated with the denoising-based version of VDAMP is $\hat{\mathbf{x}} = f(\mathbf{r}_T)$.¹

As demonstrated in the Generalized SURE work [33], an unbiased risk estimate for removing colored Gaussian noise $\eta_T \sim \mathcal{N}(0, \Sigma)$ is

$$S(\hat{\mathbf{x}}, \mathbf{r}_T) = \frac{1}{n} \|\hat{\mathbf{x}} - \mathbf{r}_T\|^2 + \frac{2}{n} \text{div}_{\mathbf{u}}(\hat{\mathbf{x}}) - \text{tr}(\Sigma), \quad (7)$$

where $\mathbf{u} = \Sigma^{-1} \mathbf{r}_T$.

We can extend this estimate to the complex case by noting that $\|\mathbf{x} - \hat{\mathbf{x}}\|^2 = \|\mathcal{R}(\mathbf{x} - \hat{\mathbf{x}})\|^2 + \|\mathcal{I}(\mathbf{x} - \hat{\mathbf{x}})\|^2$. We next note that with $\Sigma = \Psi^t \text{diag}(\tau_t) \Psi$, the similarity invariance of the trace function implies $\text{tr}(\Sigma) = \sum_{i=1}^n \tau_T^{(i)}$. We are then left with

$$S(\hat{\mathbf{x}}, \mathbf{r}_T) = \|\hat{\mathbf{x}} - \mathbf{r}_T\|^2 - \sum_{i=1}^n \tau_T^{(i)} + \frac{2}{n} (\text{div}_{\mathcal{R}(\mathbf{u})}(\mathcal{R}(\hat{\mathbf{x}})) + \text{div}_{\mathcal{I}(\mathbf{u})}(\mathcal{I}(\hat{\mathbf{x}}))), \quad (8)$$

where $\mathbf{u} = \Psi \text{diag}(\frac{1}{2} \tau_t)^{-1} \Psi^t \mathbf{r}_T$.

To estimate the divergence, we let $\tilde{f}(\mathbf{u}) = f(\Sigma \mathbf{u})$. Now we have

$$\text{div}_{\mathbf{u}}(f(\mathbf{r}_T)) = \text{div}_{\mathbf{u}}(\tilde{f}(\mathbf{u})),$$

and can use the Monte-Carlo approximation (5) to obtain

$$\text{div}_{\mathbf{u}}(f(\mathbf{r}_T)) \approx \frac{1}{K} \sum_{k=1}^K \frac{1}{\epsilon} \mathbf{b}_k^T (f(\mathbf{r}_T + \epsilon \Sigma \mathbf{b}_k) - f(\mathbf{r}_T)), \quad (9)$$

which we apply independently to both the real and imaginary parts of $f(\mathbf{r}_T)$.

As before, we generate per-pixel SURE heatmaps by averaging the overlapping estimated risks of square patches.

5. EXPERIMENT

5.1. Setting

We test our SURE heatmap generation method with CS reconstructions using D-AMP (Gaussian measurement matrices) and D-VDAMP (subsampling Fourier measurement matrices). For D-AMP, the sampling rate, m/n , is 5% and the SNRs are 23dB and 18dB for the natural image and the MR image, respectively. For D-VDAMP, the sampling rate is 25% and the SNR is 20dB. The Fourier coefficients were selected using polynomial variable density sampling [34]. Both D-AMP and D-VDAMP used a collection of DnCNN [35] denoisers trained for multiple noise levels from $\sigma = 0$ to $\sigma = 300/255$. The natural images were 512×512 while the MR images were 320×320 .

5.2. Accuracy-resolution trade-off

We first investigate the accuracy of the MSE estimate as a function of patchsize. Figure 4 compares the average difference squared between the SURE estimate and the true

¹The original VDAMP work, which was based on soft wavelet thresholding, included an additional gradient step after denoising \mathbf{r}_T [6]. In [31], the authors found this term hurts the algorithm's performance when dealing with more advanced denoising algorithms, and so we do not adopt it here.

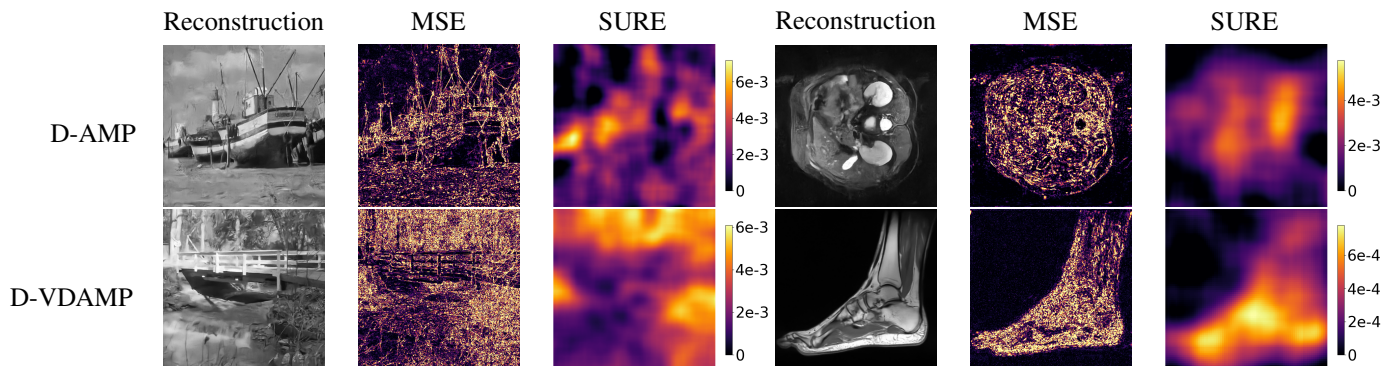


Fig. 3: SURE heatmaps and MSE heatmaps of CS reconstructions with D-AMP and D-VDAMP along with the reconstructed images. For all heatmaps in this figure, the patch size is 48×48 pixels, and the number of Monte-Carlo samples for divergence estimation is 2. MRI images are from MRIdata.org.




Fig. 4: Normalized absolute difference between the SURE heatmap and the patch-average (effectively low-pass filtered) MSE heatmap, which is generated by averaging overlapping patches of MSEs in the same fashion as the SURE heatmap generation. Data is for a CS reconstruction using D-AMP.

MSE of the image as one increases the patch sizes used in the SURE estimates. We observe that, due primarily to the reduced variance of the data fidelity term ($\|\hat{\mathbf{x}} - \mathbf{r}_T\|^2$), the SURE heatmaps become more accurate as the patch size increases. Increasing the number of Monte-Carlo samples, K , has only a slight effect on the accuracy of the estimate. Figure 5 compares the heatmaps formed with various patch sizes. While smaller patch-sizes are higher resolution, larger patch sizes result in more accurate MSE estimates. We found 48×48 patches provided a nice trade-off between resolution and accuracy.

5.3. Results

Figure 3 generates the SURE heatmaps for D-AMP and D-VDAMP reconstructions using a patch size of 48×48 pixels. While somewhat low resolution, the shapes and magnitudes of the heatmaps closely follow the true pixelwise MSEs. These heatmaps, which *do not require the ground truth*, could prove valuable for medical diagnosis and other safety-critical applications.




Fig. 5: SURE heatmaps with small patch sizes. The left heatmap is the MSE. The middle and the right heatmaps are SURE heatmaps of a CS reconstruction with D-AMP generated by using patch widths of 1 pixel and 16 pixels respectively. The number of Monte-Carlo samples, K , is 3 for both heatmaps.

6. REFERENCES

- [1] Emmanuel J Candès, Justin Romberg, and Terence Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on information theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [2] Weisheng Dong, Guangming Shi, Xin Li, Yi Ma, and Feng Huang, “Compressive sensing via nonlocal low-rank regularization,” *IEEE Transactions on Image Processing*, vol. 23, no. 8, pp. 3618–3632, 2014.
- [3] Jian Sun, Huibin Li, Zongben Xu, et al., “Deep admm-net for compressive sensing mri,” in *Advances in neural information processing systems*, 2016, pp. 10–18.
- [4] David L Donoho, Arian Maleki, and Andrea Montanari, “Message-passing algorithms for compressed sensing,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18914–18919, 2009.
- [5] Charles M Stein, “Estimation of the mean of a multivariate normal distribution,” *The Annals of Statistics*, vol. 9, no. 6, pp. 1135–1151, 1981.
- [6] Charles Millard, Aaron T Hess, Boris Mailhé, and Jared Tanner, “Approximate message passing with a colored aliasing model for variable density fourier sampled images,” *arXiv preprint arXiv:2003.02701*, 2020.
- [7] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang, “Toward convolutional blind denoising of real photographs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1712–1722.

- [8] Shihao Ji, Ya Xue, and Lawrence Carin, "Bayesian compressive sensing," *IEEE Transactions on signal processing*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [9] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen, "On instabilities of deep learning in image reconstruction-does ai come at a cost?," *arXiv preprint arXiv:1902.05300*, 2019.
- [10] Nina M Gottschling, Vegard Antun, Ben Adcock, and Anders C Hansen, "The troublesome kernel: why deep learning for inverse problems is typically unstable," *arXiv preprint arXiv:2001.01258*, 2020.
- [11] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis, "Compressed sensing using generative models," in *International Conference on Machine Learning*, 2017, pp. 537–546.
- [12] Paul Hand and Vladislav Voroninski, "Global guarantees for enforcing deep generative priors by empirical risk," in *Conference On Learning Theory*. PMLR, 2018, pp. 970–978.
- [13] Anna C Gilbert, Yi Zhang, Kibok Lee, Yuting Zhang, and Honglak Lee, "Towards understanding the invertibility of convolutional neural networks," *arXiv preprint arXiv:1705.08664*, 2017.
- [14] Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W Pellegrini, Ralf S Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe, "Analyzing inverse problems with invertible neural networks," *arXiv preprint arXiv:1808.04730*, 2018.
- [15] Diederik P Kingma and Max Welling, "An introduction to variational autoencoders," *arXiv preprint arXiv:1906.02691*, 2019.
- [16] Jonas Adler and Ozan Öktem, "Deep posterior sampling: Uncertainty quantification for large scale inverse problems," in *International Conference on Medical Imaging with Deep Learning-Extended Abstract Track*, 2019.
- [17] Mark Tygert, Rachel Ward, and Jure Zbontar, "Compressed sensing with a jackknife and a bootstrap," *arXiv preprint arXiv:1809.06959*, 2018.
- [18] Ryutaro Tanno, Daniel Worrall, Enrico Kaden, Aurobrata Ghosh, Francesco Grussu, Alberto Bizzi, Stamatios N Sotiropoulos, Antonio Criminisi, and Daniel C Alexander, "Uncertainty quantification in deep learning for safer neuroimage enhancement," *arXiv preprint arXiv:1907.13418*, 2019.
- [19] Thomas Küstner, Annika Liebgott, Lukas Mauch, Petros Martirosian, Fabian Bamberg, Konstantin Nikolaou, Bin Yang, Fritz Schick, and Sergios Gatidis, "Automated reference-free detection of motion artifacts in magnetic resonance images," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 31, no. 2, pp. 243–256, 2018.
- [20] Vineet Edupuganti, Morteza Mardani, Shreyas Vasanawala, and John Pauly, "Uncertainty quantification in deep mri reconstruction," *IEEE Transactions on Medical Imaging*, 2020.
- [21] David L Donoho and Iain M Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *Journal of the american statistical association*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [22] Ali Mousavi, Arian Maleki, and Richard G Baraniuk, "Parameterless optimal approximate message passing," *arXiv preprint arXiv:1311.0035*, 2013.
- [23] Chunli Guo and Mike E Davies, "Near optimal compressed sensing without priors: Parametric sure approximate message passing," *IEEE Transactions on Signal Processing*, vol. 63, no. 8, pp. 2130–2141, 2015.
- [24] Charles Millard, Aaron T Hess, Boris Mailhe, and Jared Tanner, "An approximate message passing algorithm for rapid parameter-free compressed sensing mri," in *2020 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2020, pp. 91–95.
- [25] Christopher A Metzler, Ali Mousavi, Reinhard Heckel, and Richard G Baraniuk, "Unsupervised learning with stein's unbiased risk estimator," 2018.
- [26] Shakarim Soltanayev and Se Young Chun, "Training deep learning based denoisers without ground truth data," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., pp. 3257–3267. Curran Associates, Inc., 2018.
- [27] Magauyiya Zhussip, Shakarim Soltanayev, and Se Young Chun, "Training deep learning based image denoisers from under-sampled measurements without ground truth and without image prior," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10255–10264.
- [28] Charles-Alban Deledalle, Vincent Duval, and Joseph Salmon, "Non-local methods with shape-adaptive patches (nlm-sap)," *Journal of Mathematical Imaging and Vision*, vol. 43, no. 2, pp. 103–120, 2012.
- [29] Christopher A Metzler, Arian Maleki, and Richard G Baraniuk, "From denoising to compressed sensing," *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5117–5144, 2016.
- [30] Christopher A Metzler, Ali Mousavi, and Richard Baraniuk, "Learned d-amp: Principled neural network based compressive image recovery," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., pp. 1772–1783. Curran Associates, Inc., 2017.
- [31] Christopher A. Metzler and Gordon Wetzstein, "D-vdamp: Denoising-based approximate message passing for compressive mri," *Under Review*.
- [32] Sathish Ramani, Thierry Blu, and Michael Unser, "Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [33] Yonina C Eldar, "Generalized sure for exponential families: Applications to regularization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 471–481, 2009.
- [34] Michael Lustig, David Donoho, and John M Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [35] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.