Compressed Sensing Recovery of Medical Images using Deep Image Prior

David Van Veen
vanveen@utexas.edu

Ajil Jalal
ajiljalal@utexas.edu

Eric Price
ecprice@cs.utexas.edu

Sriram Vishwanath
sriram@austin.utexas.edu

Alexandros G. Dimakis
dimakis@austin.utexas.edu

Abstract

We propose a novel method for compressed sensing recovery of medical images using untrained deep generative models. Our method is based on the recently proposed Deep Image Prior (DIP), wherein the convolutional weights of the network are optimized to match the observed measurements. We extend this approach to solve any differentiable inverse problem. Our algorithm requires approximately 4 - 6$\times$ fewer measurements compared to classical compressed sensing approaches such as Lasso. More recent approaches based on *pre-trained* generative models are severely limited, as they can reconstruct only simple signals and also require vast amounts of training data. Our approach circumvents these limitations by using an *untrained* network. As such, we can apply our method to medical imaging problems for which signals are complex and data acquisition is expensive.

1 Introduction

We consider the well-studied compressed sensing problem of recovering an unknown signal $x^* \in \mathbb{R}^n$ by observing a set of noisy measurements $y \in \mathbb{R}^m$ of the form

$$y = Ax^* + \eta.$$

Here $A \in \mathbb{R}^{m \times n}$ is a known measurement matrix, typically generated with random independent Gaussian entries. Since the number of measurements m is smaller than the dimension n of the unknown vector x^*, this is an under-determined system of noisy linear equations and hence ill-posed. There are many solutions, and some structure must be assumed on x^* to have any hope of recovery.

Pioneering research [8, 5, 6] established that if x^* is assumed to be sparse in a known basis, a small number of measurements will be provably sufficient to recover the unknown vector in polynomial time using methods such as Lasso [28, 31]. While sparsity in a known basis has proven successful, more complex models with additional structure have been recently proposed such as model-based compressive sensing [3] and manifold models [15, 14, 9]. Recently Bora et al. [4] showed that deep generative models [11, 19, 25] can be used as excellent priors for images and were able to reconstruct images with up to 10$\times$ fewer measurements compared to Lasso for a given reconstruction error. Compressed sensing using deep generative models was further improved in very recent work [30, 12, 18, 27, 10, 2, 13, 21].

Inspired by these impressive benefits of deep generative models, we chose to investigate the application of such methods for medical imaging, a canonical application of compressive sensing. A significant problem, however, is that all these previous methods require the existence of *pre-trained* models. While generative models for medical images exist [22, 26, 33], they are still far from perfect. Instead of addressing this important problem in generative models, we use the Deep Image Prior (DIP) method proposed by Ulyanov et al. [32].

32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada.

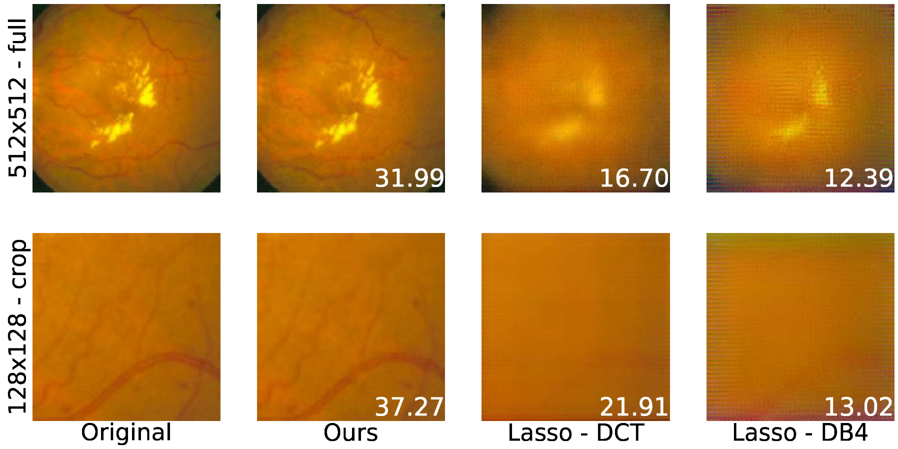

Figure 1: Comparison of our algorithm (CS-DIP) vs. baselines on RGB retinopathy medical images. We show a full 512x512 image (first row) and a 128x128 cropped portion of the same image (second row) to demonstrate finer details. For each row we show the original image (first column), our reconstruction (second column), and baseline reconstructions by Lasso in both the DCT basis and DB4 wavelet basis (third, fourth columns). Values in lower right of reconstructed images denote PSNR (dB). We recommend viewing in color.

Our Contributions: In this paper we propose DIP for compressed sensing (CS-DIP). Our basic method is simple: we initialize a DCGAN generator with random weights and optimize them using gradient descent to make the network produce an output which *agrees with the observed measurements* as much as possible.

Our results show that we require approximately $4 - 6\times$ fewer measurements to obtain similar reconstruction error compared to Lasso. This benefit is not as significant as the $8 - 10\times$ obtained by Bora et al. [4], but we have the advantage of not requiring a pre-trained generative model. We can therefore apply our method to various medical imaging datasets without depending on a high quality generative model.

2 Methods

Let $x^* \in \mathbb{R}^n$ be the signal that we are trying to reconstruct, $A \in \mathbb{R}^{m \times n}$ be the measurement matrix, and $\eta \in \mathbb{R}^m$ be independent noise. Given the measurement matrix A and the observations $y = Ax^* + \eta$, we wish to reconstruct an $\hat{x}$ that is close to x^*. A generative model is a deterministic function $G(z; w)$: $\mathbb{R}^k \to \mathbb{R}^n$ which takes as input a seed $z \in \mathbb{R}^k$ and a set of parameters (or "weights") $w \in \mathbb{R}^d$, producing an output $G(z; w) \in \mathbb{R}^n$. In this paper we apply a DCGAN [25] model and restrict the signals to be images.

Our approach is to find a set of weights for the convolutional network such that the measurement matrix applied to the network output, i.e. $AG(z; w)$, matches the measurements y we are given. Hence we initialize an *untrained* network $G(z; w)$ with some fixed z and solve the following optimization problem:

$$w^* = \arg\min_{w} \|y - AG(z; w)\|^2.$$

This problem is non-convex, as $G(z; w)$ is a complex feed-forward neural network. Still we can use gradient-based optimizers for any model and measurement process that is differentiable. Ulyanov et al. observed that convolutional networks such as DCGAN are biased toward smooth, natural images; hence the network structure alone provides a good prior for reconstructing images in problems such as inpainting and denoising [32]. Our finding is that this applies to general linear measurement processes. We restrict our solution to lie in the span of a convolutional network; if a sufficient number of measurements m is given, we obtain an output such that $x^* \approx G(z; w^*)$.

Note that this method uses an untrained generative model and optimizes over the network weights w. In contrast previous methods such as that of Bora et al. [4] use a trained model and optimize in the

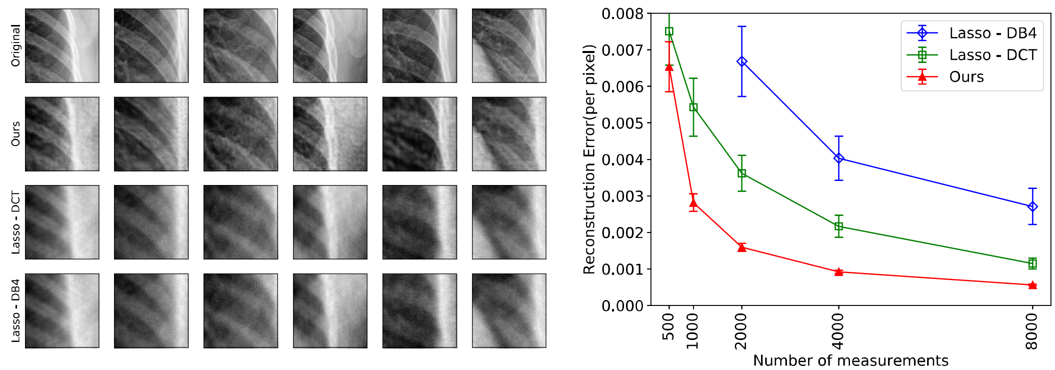

(a) Reconstructions (m=2000, n=65536 pixels)

(b) MSE for Chest X-Rays

Figure 2: In Figure 2a we show reconstruction results on Chest X-rays for $m = 2000$ measurements (out of $n = 65536$ pixels). We show original images (top row), reconstructions by our algorithm (second row), and baseline reconstructions by Lasso in both the DCT basis and DB4 wavelet basis (bottom two rows). In Figure 2b we compare the performance of our algorithm with baselines, plotting per-pixel reconstruction error (MSE) vs. number of measurements, where vertical bars indicate 95% confidence intervals.

latent z-space, solving $z^* = \arg\min_z \|y - AG(z; w)\|^2$. We instead initialize a random z and keep this fixed throughout the optimization process.

3 Experiments and Results

As is standard in compressed sensing, the measurement matrices $A \in \mathbb{R}^{m \times n}$ are created by sampling Gaussian I.I.D. random variables for each matrix entry, such that $A_{i,j} \sim \mathcal{N}(0, \frac{1}{m})$. Recall m is the number of measurements, and n is the number of pixels in the ground truth image. To quantitatively evaluate the performance of our algorithm, we use per-pixel mean-squared error (MSE) between the reconstruction $\hat{x}$ and true image x^*, i.e. $\frac{\|\hat{x} - x^*\|^2}{n}$. We compare against Lasso in the DCT [1] and DB4 basis [7]. For further details on details on datasets and optimizers used, please see Appendix 5.1. Also included in the appendix are additional experiments on the MNIST dataset [20].

3.1 Results - Chest X-Rays

In Figure 2b we plot the reconstruction error with varying number of measurements m of n = 65536. On this dataset we also significantly outperform Lasso, which requires roughly $4\times$ more measurements to achieve the same error. Figure 2a demonstrates reconstructions using our algorithm compared to Lasso. We observe that our algorithm produces reasonable reconstructions with as few as $m = 2000$ measurements, while Lasso's output is quite blurry. The reconstructed images for $m = 1000, 4000, 8000$ are in the appendix (Fig. 5, 6, 7).

3.2 Results - Retinopathy

In Figure 1 we show reconstructions using our algorithm (CS-DIP) and Lasso. We see that our algorithm can reconstruct fine details in the retina scans, such as blood vessels and bright spots. Meanwhile for the same number of measurements, Lasso fails to produce a comprehensible image.

4 Conclusion

We demonstrate how to perform compressed sensing on medical images using untrained, randomly initialized convolutional neural networks. Our approach requires 4 - $6\times$ fewer measurements compared to classical Lasso methods. Compared to more recent methods with generative models, ours is unique in that we don't require the network to be pre-trained. Thus our method can reconstruct complex signals without training data and can hence be applied to medical imaging.

References

[1] Nasir Ahmed, T_ Natarajan, and Kamisetty R Rao. Discrete cosine transform. *IEEE transactions on Computers*, 100(1):90–93, 1974.

[2] Muhammad Asim, Fahad Shamshad, and Ali Ahmed. Solving bilinear inverse problems using deep generative priors. *CoRR*, abs/1802.04073, 2018.

[3] Richard G Baraniuk, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. Model-based compressive sensing. *IEEE Transactions on Information Theory*, 56(4):1982–2001, 2010.

[4] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. *arXiv preprint arXiv:1703.03208*, 2017.

[5] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.

[6] Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.

[7] Ingrid Daubechies. Orthonormal bases of compactly supported wavelets. *Communications on pure and applied mathematics*, 41(7):909–996, 1988.

[8] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.

[9] Armin Eftekhari and Michael B Wakin. New analysis of manifold embeddings and signal recovery from compressive measurements. *Applied and Computational Harmonic Analysis*, 39(1):67–109, 2015.

[10] Alyson K Fletcher and Sundeep Rangan. Inference in deep networks in high dimensions. *arXiv preprint arXiv:1706.06549*, 2017.

[11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[12] Aditya Grover and Stefano Ermon. Amortized variational compressive sensing. 2018.

[13] Paul Hand and Vladislav Voroninski. Global guarantees for enforcing deep generative priors by empirical risk. *arXiv preprint arXiv:1705.07576*, 2017.

[14] Chinmay Hegde and Richard G Baraniuk. Signal recovery on incoherent manifolds. *IEEE Transactions on Information Theory*, 58(12):7204–7214, 2012.

[15] Chinmay Hegde, Michael Wakin, and Richard Baraniuk. Random projections for manifold learning. In *Advances in neural information processing systems*, pages 641–648, 2008.

[16] AD Hoover, Valentina Kouznetsova, and Michael Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical imaging*, 19(3):203–210, 2000.

[17] Stefan Jaeger, Sema Candemir, Sameer Antani, Yì-Xiáng J Wáng, Pu-Xuan Lu, and George Thoma. Two public chest x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6):475, 2014.

[18] Maya Kabkab, Pouya Samangouei, and Rama Chellappa. Task-aware compressed sensing with generative adversarial networks. *arXiv preprint arXiv:1802.01284*, 2018.

[19] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[21] Morteza Mardani, Hatef Monajemi, Vardan Papyan, Shreyas Vasanawala, David Donoho, and John Pauly. Recurrent generative adversarial networks for proximal learning and automated compressive image recovery. *arXiv preprint arXiv:1711.10046*, 2017.

[22] Dong Nie, Roger Trullo, Jun Lian, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with context-aware generative adversarial networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 417–425. Springer, 2017.

[23] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.

[24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[25] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[26] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, pages 146–157. Springer, 2017.

[27] Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. *arXiv preprint arXiv:1802.08406*, 2018.

[28] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[29] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.

[30] Subarna Tripathi, Zachary C Lipton, and Truong Q Nguyen. Correction by projection: Denoising images with generative adversarial networks. *arXiv preprint arXiv:1803.04477*, 2018.

[31] Joel A Tropp. Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE transactions on information theory*, 52(3):1030–1051, 2006.

[32] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *arXiv preprint arXiv:1711.10925*, 2017.

[33] Jelmer M Wolterink, Tim Leiner, Max A Viergever, and Ivana Išgum. Generative adversarial networks for noise reduction in low-dose ct. *IEEE transactions on medical imaging*, 36(12):2536–2545, 2017.