

# Solving Linear Inverse Problems Using the Prior Implicit in a Denoiser

**Zahra Kadkhodaie**  
Center for Data Science, and  
Howard Hughes Medical Institute  
New York University  
zk388@nyu.edu

**Eero P. Simoncelli**  
Center for Neural Science,  
Courant Inst. of Mathematical Sciences, and  
Howard Hughes Medical Institute  
New York University  
eero.simoncelli@nyu.edu

Many problems in image processing and computer vision rely, explicitly or implicitly, on prior probability models. Describing the full density of natural images is a daunting problem, given the high dimensionality of the signal space. Traditionally, models have been developed by combining assumed symmetry properties (e.g., translation-invariance, dilation-invariance), with simple parametric forms (e.g., Gaussian, exponential), often within pre-specified transformed coordinate systems (e.g., Fourier transform, wavelets). While these models have led to steady advances in problems such as denoising (e.g., [1–7]), they are too simplistic to generate complex features that occur in our visual world, or to solve more demanding statistical inference problems.

Nearly all problems in image processing and computer vision have been revolutionized by the use of deep Convolutional Neural Networks (CNNs). These networks are generally trained to perform tasks in supervised fashion, and their embedding of prior information, arising from a combination of the distribution of the training data, the architecture of the network [8], and regularization terms included during optimization, is intertwined with the task for which they are optimized. Generative Adversarial Networks (GANs) [9] have been shown capable of synthesizing novel high-quality images, that may be viewed as samples of an implicit prior density model, and recent methods have been developed to use these priors in solving inverse problems [10, 11]. Although GANs can produce visually impressive synthetic samples, a number of results suggest that these samples are not representative of the full image density (a problem sometimes referred to as “mode collapse”) [12]. Variational autoencoder (VAE) networks [13] have been used in conjunction with Score Matching methodologies to draw samples from an implicit prior [14–18]. Related work has developed algorithms for using MAP denoisers to regularize inverse problems, obtaining high quality results on deblurring and superresolution using conventional denoisers [19].

Here, we develop a simple but general algorithm for solving linear inverse problems using the implicit image prior from a CNN trained (supervised) for denoising. We start with a little-known result from classical statistics [20] that states that a denoiser that aims to minimize squared error of images corrupted by additive Gaussian noise may be interpreted as computing the gradient of the log of the density of noisy images. We develop a stochastic coarse-to-fine gradient ascent procedure based on this denoiser-estimated gradient, and use it to draw high-probability samples from the implicit prior embedded within a CNN trained to perform blind (i.e., unknown noise level) least-squares denoising. The gradient step sizes and the amplitude of injected noise are jointly and adaptively controlled by the denoiser, allowing rapid and reliable convergence. More generally, we combine this procedure with constraints arising from any linear measurement of an image to draw samples from the prior conditioned on this measurement, thus providing a stochastic solution to the inverse problem. We demonstrate that our method, using the prior implicit in a state-of-the-art CNN denoiser, produces high-quality results on image synthesis, inpainting, super-resolution, deblurring and recovery of missing pixels. We also apply our method to recovering images from projections onto a random low-dimensional basis, demonstrating results that greatly improve on those obtained using sparse union-of-subspace priors typically assumed in the compressive sensing literature.

## Image priors, manifolds, and noisy observations

Digital photographic images lie in a high-dimensional space ( $\mathbb{R}^N$ , where  $N$  is the number of pixels), and simple thought experiments suggest that they are concentrated on or near low-dimensional manifolds. For a given photograph, applying any of a variety of local continuous deformations (e.g., translations, rotations, dilations, intensity changes) yields a low-dimensional family of natural-appearing images. These deformations follow complex curved trajectories in the space of pixels, and thus lie on a manifold. In contrast, images generated with random pixels are almost always feature and content free, and thus not considered to be part of this manifold. We can associate with this a prior probability model,  $p(x)$ , by assuming that images within the manifold have constant or slowly-varying probability, while unnatural or distorted images (which lie off the manifold) have low or zero probability.

Suppose we make a noisy observation of an image,  $y = x + z$ , where  $x \in R^N$  is the original image drawn from  $p(x)$ , and  $z \sim \mathcal{N}(0, \sigma^2 I_N)$  is a sample of Gaussian white noise. The observation density  $p(y)$  (also known as the *prior predictive density*) is related to the prior  $p(x)$  via marginalization:

$$p(y) = \int p(y|x)p(x)dx = \int g(y-x)p(x)dx, \quad (1)$$

where the noise distribution is

$$g(z) = \frac{1}{(2\pi\sigma^2)^{N/2}} e^{-\|z\|^2/2\sigma^2}.$$

Equation (1) is in the form of a convolution, and thus  $p(y)$  is a Gaussian-blurred version of the signal prior,  $p(x)$ . Moreover, the family of observation densities over different noise variances,  $p_\sigma(y)$ , forms a Gaussian scale-space representation of the prior [21, 22], analogous to the temporal evolution of a diffusion process.

### Least squares denoising and CNNs

Given a noisy observation,  $y$ , the least squares estimate (also called "minimum mean squared error", MMSE) of the true signal is well known to be the conditional mean of the posterior:

$$\hat{x}(y) = \int xp(x|y)dx = \int x \frac{p(y|x)p(x)}{p(y)} dx \quad (2)$$

Traditionally, one obtains such estimators by choosing a prior probability model,  $p(x)$  (often with parameters fit to sets of images), combining it with a likelihood function describing the noise,  $p(y|x)$ , and solving. For example, the Wiener filter is derived by assuming a Gaussian prior in which variance falls inversely with spatial frequency [23]. Modern denoising solutions, on the other hand, are often based on discriminative training. One expresses the estimation function (as opposed to the prior) in parametric form, and sets the parameters by minimizing the denoising MSE over a large training set of example signals and their noise-corrupted counterparts [24–27]. The size of the training set is virtually unlimited, since it can be constructed automatically from a set of photographic images, and does not rely on human labelling.

Current state-of-the-art denoising results using CNNs are far superior, both numerically and visually, to results of previous methods [28–30]. Recent work [31] demonstrates that these architectures can be simplified by removing all additive bias terms, with no loss of performance. The resulting *bias-free* networks offer two important advantages. First, they automatically generalize to all noise levels: a network trained on images with barely noticeable levels of noise can produce high quality results when applied to images corrupted by noise of any amplitude. Second, they may be analyzed as adaptive linear systems, which reveals that they perform an approximate projection onto a low-dimensional subspace. In our context, we interpret this subspace as a tangent hyperplane of the image manifold at a specific location. Moreover, the dimensionality of these subspaces falls inversely with  $\sigma$ , and for a given noise sample, the subspaces associated with different noise amplitude are nested, with high-noise subspaces lying within their lower-noise counterparts. In the limit as the noise variance goes to zero, the subspace dimensionality grows to match that of the manifold at that particular point.

### Exposing the implicit prior through Empirical Bayes estimation

The trained CNN denoisers mentioned above embed detailed prior knowledge of image structure. Given such a denoiser, how can we obtain access to this implicit prior? Recent results have derived relationships between Score matching density estimates and denoising [14, 32, 17, 33], and have used these relationships to make use of implicit prior information. Here, we exploit a much more direct but little-known result from the literature on Empirical Bayesian estimation. The idea was introduced in [34], extended to the case of Gaussian additive noise in [20], and generalized to many other measurement models in [35]. For the Gaussian noise case, the least-squares estimate of Eq. (2) may be rewritten as:

$$\hat{x}(y) = y + \sigma^2 \nabla_y \log p(y). \quad (3)$$

The proof of this result is relatively straightforward. First note that the gradient of the observation density expressed in Eq. (1) is:

$$\nabla_y p(y) = \frac{1}{\sigma^2} \int (x-y)g(y-x)p(x)dx = \frac{1}{\sigma^2} \int (x-y)p(y,x)dx.$$

Multiplying both sides by  $\sigma^2/p(y)$  and separating the right side into two terms gives:

$$\sigma^2 \frac{\nabla_y p(y)}{p(y)} = \int xp(x|y)dx - \int yp(x|y)dx = \hat{x}(y) - y. \quad (4)$$

Rearranging terms and using the chain rule to compute the gradient of the log gives Miyasawa’s result, as expressed in Eq. (3).

Intuitively, the Empirical Bayesian form in Eq. (3) suggests that denoisers use a form of gradient ascent, removing noise from an observed signal by moving up a probability gradient. But note that: 1) the relevant density is not the prior,  $p(x)$ , but the noisy *observation density*,  $p(y)$ ; 2) the gradient is computed on the *log* density (the associated “energy function”); and 3) the adjustment is not iterative - the optimal solution is achieved in a single step, and holds for any noise level,  $\sigma$ .

## 1 Drawing high-probability samples from the implicit prior

We wish to draw a sample from the prior implicit in a denoiser. Equation (4) allows us to generate an image proportional to the gradient of  $\log p(y)$  by computing the denoiser residual,  $f(y) = \hat{x}(y) - y$ . Previous work [15, 17] developed a related computation in a Markov chain Monte Carlo (MCMC) scheme, combining gradient steps derived from Score-matching and injected noise in a Langevin sampling algorithm to draw samples from a sequence of densities  $p_\sigma(y)$ , while reducing  $\sigma$  in discrete steps, each associated with an appropriately trained denoiser. In contrast, starting from a random initialization,  $y_0$ , we aim to find a high-probability image (i.e., an image from the manifold) using a simpler and more efficient stochastic gradient ascent procedure.

We compute gradients using the residual of a bias-free universal CNN denoiser, which automatically adapts to each noise level. On each iteration, we take a small step in the direction specified by the denoiser, which moves closer to the image manifold, thereby reducing the amplitude of the effective noise. The reduction of noise is achieved by decreasing the amplitude in the directions orthogonal to the observable manifold while retaining the amplitude of image in the directions parallel to manifold which results in the emergence of image content. As the effective noise decreases, the observable dimensionality of the image manifold increases [31], allowing the synthesis of detailed structures. Since the family of observation densities,  $p_\sigma(y)$  forms a scale-space representation of  $p(x)$ , the algorithm may be viewed as an adaptive form of coarse-to-fine optimization [36–39]. Assuming the step sizes are adequately-controlled, the procedure converges to a point on the manifold. Figure 2 illustrates this process in two dimensions.

Each iteration operates by taking a deterministic step in the direction of the gradient (as obtained from the denoising function) and injecting some additional noise:

$$y_t = y_{t-1} + h_t f(y_{t-1}) + \gamma_t z_t, \quad (5)$$

where  $f(y) = \hat{x}(y) - y$  is the residual of the denoising function, which is proportional to the gradient of  $p(y)$ , from Eq. (4). The parameter  $h_t \in [0, 1]$  controls the fraction of the denoising correction that is taken, and  $z_t \sim \mathcal{N}(0, I)$  is a sample of white Gaussian noise, scaled by parameter  $\gamma_t$ , whose purpose is to avoid getting stuck in local maxima. The effective noise variance of image  $y_t$  is:

$$\sigma_t^2 = (1 - h_t)^2 \sigma_{t-1}^2 + \gamma_t^2, \quad (6)$$

where the first term assumes that the denoiser successfully reduces the variance of the noise in  $y_{t-1}$  by a factor of  $(1 - h_t)$ , and the second term is the variance arising from the injected noise. To ensure convergence, we aim to reduce the effective noise variance on each time step, which we express in terms of a parameter  $\beta \in [0, 1]$  as:

$$\sigma_t^2 = (1 - \beta h_t)^2 \sigma_{t-1}^2.$$

Combining this with Eq. (6) allows us to solve for  $\gamma_t$ :

$$\gamma_t^2 = [(1 - \beta h_t)^2 - (1 - h_t)^2] \sigma_{t-1}^2 = [(1 - \beta h_t)^2 - (1 - h_t)^2] \|f(y_{t-1})\|^2 / N, \quad (7)$$

where the second line assumes that the magnitude of the denoising residual provides a good estimate of the effective noise standard deviation, as was found in [31]. This allows the denoiser to adaptively control the gradient ascent step sizes, reducing them as the result approaches the manifold (see Figure 2 for a visualization in 2D). This automatic adjustment of the gradient step size is an important feature of the procedure, which results in a fast and smooth convergence as illustrated in Figure 4. We found that initial implementations with a small constant fractional step size  $h_t$  produced high quality results, but required many iterations - a form of Zeno’s paradox. To improve convergence speed, we introduced a schedule for increasing the step size according to  $h_t = \frac{h_0 t}{1 + h_0(t-1)}$ , starting from  $h_0 \in [0, 1]$ . The sampling algorithm is summarized in Algorithm 1 in section B, and is laid out in a block diagram in Figure 1 in the appendix. Example convergence behavior is shown in Figure 3. Sampling examples are shown in section F.

## 2 Solving linear inverse problems using the implicit prior

Many applications in signal processing can be expressed as linear inverse problems - deblurring, super-resolution, estimating missing pixels (e.g., inpainting), and compressive sensing are all examples. Given a set of linear measurements

of an image,  $x^c = M^T x$ , where  $M$  is a low-rank measurement matrix, one attempts to recover the original image. In Section 1, we developed a stochastic gradient-ascent algorithm for obtaining a high-probability sample from  $p(x)$ . Here, we modify this algorithm to solve for a high-probability sample from the conditional density  $p(x|M^T x = x^c)$ .

Consider the distribution of a noisy image,  $y$ , conditioned on the linear deterministic measurements,  $x^c = M^T x$ ,

$$p(y|x^c) = p(y|y^c)p(y^c|x^c)$$

where  $y^c = M^T y$ . Without loss of generality<sup>1</sup>, we assume the measurement matrix has singular values that are equal to one (i.e., columns of  $M$  are orthogonal unit vectors, and thus  $M^T M = I$ ). It follows that  $M$  is the pseudo-inverse of  $M^T$ , and that matrix  $MM^T$  can be used to project an image onto the measurement subspace. As with the algorithm of Section 1, we wish to obtain a local maximum of this function using stochastic coarse-to-fine gradient ascent. Applying the operator  $\sigma^2 \nabla \log$  yields

$$\sigma^2 \nabla_y \log p(y|x^c) = \sigma^2 \nabla_y \log p(y|y^c) + \sigma^2 \nabla_y \log p(y^c|x^c).$$

The second term is the gradient of the observation noise distribution, projected into the measurement space. If this is Gaussian with variance  $\sigma^2$ , it reduces to  $M(y^c - x^c)$ . The first term is the gradient of a function defined only within the subspace orthogonal to the measurements, and thus can be computed by projecting the measurement subspace out of the full gradient. Combining these gives:

$$\begin{aligned} \sigma^2 \nabla_y \log p(y|x^c) &= (I - MM^T) \sigma^2 \nabla_y \log p(y) + M(x^c - y^c) \\ &= (I - MM^T) f(y) + (Mx^c - MM^T y). \end{aligned} \quad (8)$$

Thus, we see that the gradient of the conditional density is partitioned into two orthogonal components, capturing the gradient of the (log) noisy density, and the deviation from the constraints, respectively. To draw a high-probability sample from  $p(x|x^c)$ , we use the same algorithm described in Section 1, substituting Eq. (8) for the deterministic update vector,  $f(y)$  (see Algorithm 2, and Figure 1). Examples of our algorithm solution to five linear inverse problems are shown in Section G.

### 3 Discussion and related work

We've described a general method of transferring the prior embedded in a denoiser to solve inverse problems. Specifically, we developed a stochastic coarse-to-fine gradient ascent algorithm that uses the denoiser to draw high-probability samples from its implicit prior, and a constrained variant that can be used to solve *any* linear inverse problem. The derivation relies on the denoiser being optimized for mean squared error in removing additive Gaussian noise of unknown amplitude. Denoisers can be trained using discriminative learning (nonlinear regression) on virtually unlimited amounts of unlabeled data, and thus, our method extends the power of supervised learning to a much broader set of problems, with no additional training.

Our method is similar to recent work that uses Score Matching to draw samples from an implicit prior [14–18], but differs in several important ways: (1) our derivation is direct and significantly simpler, exploiting a little-known result from the classical statistical literature on Empirical Bayes estimation [20]; (2) our method assumes a single (universal) blind denoiser, rather than a family of denoisers trained for different noise levels; (3) our algorithm is efficient - we use stochastic gradient ascent to maximize probability, rather than MCMC methods (such as Langevin dynamics) to draw samples from each of a sequence of densities, and our step sizes are automatically adapted by the denoiser to control convergence; and (4) we demonstrate the generality of our algorithm by applying it to five different linear inverse problems. Another related line of research uses a denoiser to regularize an objective function for solving linear inverse problems [19, 40]. Our method differs in that (1) the term arising from our denoiser represents the exact gradient of the (noisy) prior (it is not derived as an approximation); (2) our method requires only that the denoiser must be trained on Gaussian-noise contaminated images to minimize squared error and must operate "blind" (without knowledge of noise level), whereas RED relies on a set of three additional assumed properties; (3) RED has been used to solve MAP estimation problems, whereas we have only applied our method to deterministic linear inverse problems (although we believe it can be generalized to MAP); and (4) our algorithm has two primary hyper-parameters ( $h_0$  and  $\beta$ ) to control step sizes, and is robust to choices of these (see Appendix), whereas RED includes multiple hyper-parameters, including the gradient step size whose adjustment is important for achieving convergence (the issue is resolved through use of an ADMM method).

<sup>1</sup>The assumption that  $M$  has orthogonal columns does not restrict our solution. Consider arbitrary linear constraint,  $W^T x = x_w$ , and the singular value decomposition  $W = USV^T$ , with  $S$  square and  $U$  containing only the basis of the range (column) space, define  $M = U$ , and  $x_c = SV^T x_w$ . Then we can write an equivalent constraint as  $M^T x = x_c$ , where the columns of  $M$  are orthogonal unit vectors,  $M^T M = I$ , and  $MM^T$  is a projection matrix, as desired.

The performance of our method on linear inverse problems should ultimately be quantified by experiments with human observers, but we also plan to measure it using a no-reference perceptual quality metric (e.g., [41]). Handling of nonlinear inverse problems with convex measurements (e.g. recovery from quantized representation, such as JPEG) is a natural extension of the method, in which the algorithm must be modified to incorporate projection onto convex sets. Finally, our method for image generation offers a means of visualizing and interpreting implicit prior of a denoiser, which arises from the combination of architecture, optimization, regularization, and training set. As such, it offers a means of experimentally isolating and elucidating the effects of these components.

## References

- [1] D Donoho. Denoising by soft-thresholding. *IEEE Trans Info Theory*, 43:613–627, 1995.
- [2] E P Simoncelli and E H Adelson. Noise removal via Bayesian wavelet coring. In *Proc 3rd IEEE Int'l. Conf. on Image Processing (ICIP)*, pages 379–382, 1996.
- [3] P Moulin and J Liu. Analysis of multiresolution image denoising schemes using a generalized Gaussian and complexity priors. *IEEE Trans Info Theory*, 45:909–919, 1999.
- [4] Hyvärinen. Sparse code shrinkage: Denoising of nonGaussian data by maximum likelihood estimation. *Neural Computation*, 11(7):1739–1768, 1999.
- [5] J. Romberg, H. Choi, and R. Baraniuk. Bayesian tree-structured image modeling using wavelet-domain hidden Markov models. *IEEE Trans Image Proc*, 10(7), July 2001.
- [6] L Şendur and I W Selesnick. Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. *IEEE Trans Signal Proc*, 50(11):2744–2756, November 2002.
- [7] J Portilla, V Strela, M J Wainwright, and E P Simoncelli. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Trans Image Proc*, 12(11):1338–1351, Nov 2003. Recipient, IEEE Signal Processing Society Best Paper Award, 2008.
- [8] D Ulyanov, A Vedaldi, and V Lempitsky. Deep image prior. *Proc Int'l J Computer Vision*, pages 1867–1888, April 2020.
- [9] 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.
- [10] Viraj Shah and Chinmay Hegde. Solving linear inverse problems using GAN priors: An algorithm with provable guarantees. *2018 IEEE Int'l. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Apr 2018.
- [11] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models, 2017.
- [12] Eitan Richardson and Yair Weiss. On GANs and GMMs. In *Advances in Neural Information Processing Systems*, pages 5847–5858, 2018.
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *Second Int'l. Conf. on Learning Representations (ICLR)*, volume 19, 2014.
- [14] Y Bengio, L Yao, G Alain, and P Vincent. Generalized denoising auto-encoders as generative models. In *Adv. Neural Information Processing Systems (NIPS\*13)*, pages 899–907. MIT Press, 2013.
- [15] Saeed Saremi, Bernhard Schölkopf, Arash Mehrjou, and Aapo Hyvärinen. Deep energy estimator networks. *ArXiv e-prints (arXiv.org)*, 1805.08306, May 2018.
- [16] Zengyi Li, Yubei Chen, and Friedrich T. Sommer. Learning energy-based models in high-dimensional spaces with multi-scale denoising score matching, 2019.
- [17] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems 32*, pages 11918–11930. Curran Associates, Inc., 2019.
- [18] Siavash A Bigdeli, Geng Lin, Tiziano Portenier, L Andrea Dunbar, and Matthias Zwicker. Learning generative models using denoising density estimators. *arXiv preprint arXiv:2001.02728*, 2020.

- [19] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (RED). *CoRR*, abs/1611.02862, 2017.
- [20] K Miyasawa. An empirical Bayes estimator of the mean of a normal population. *Bull. Inst. Internat. Statist.*, 38:181–188, 1961.
- [21] J J Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
- [22] T Lindeberg. Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, 21(2):224–270, 1994.
- [23] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. Technology Press, 1950.
- [24] Y. Hel-Or and D. Shaked. A discriminative approach for wavelet shrinkage denoising. *IEEE Trans. Image Processing*, 17(4), April 2008.
- [25] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. on Image processing*, 15(12):3736–3745, 2006.
- [26] Viren Jain and Sebastian Seung. Natural image denoising with convolutional networks. In *Advances in neural information processing systems*, pages 769–776, 2009.
- [27] Harold C Burger, Christian J Schuler, and Stefan Harmeling. Image denoising: Can plain neural networks compete with BM3D? In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 2392–2399. IEEE, 2012.
- [28] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Trans. Image Processing*, 26(7):3142–3155, 2017.
- [29] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017.
- [30] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE Trans. Patt. Analysis and Machine Intelligence*, 39(6):1256–1272, 2017.
- [31] S Mohan\*, Z Kadkhodaie\*, E P Simoncelli, and C Fernandez-Granda. Robust and interpretable blind image denoising via bias-free convolutional neural networks. In *Int'l. Conf. on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
- [32] Saeed Saremi and Aapo Hyvarinen. Neural empirical Bayes. *Journal of Machine Learning Research*, 20:1–23, 2019.
- [33] Zengyi Li, Yubei Chen, and Friedrich T Sommer. Annealed denoising score matching: Learning energy-based models in high-dimensional spaces. *arXiv preprint arXiv:1910.07762*, 2019.
- [34] H Robbins. An empirical Bayes approach to statistics. *Proc. Third Berkley Symposium on Mathematical Statistics*, 1:157–163, 1956.
- [35] M Raphan and E P Simoncelli. Least squares estimation without priors or supervision. *Neural Computation*, 23(2):374–420, Feb 2011.
- [36] B D Lucas and T Kanade. An iterative image registration technique with an application to stereo vision. In *Proc. 7th Int'l Joint Conf. on Artificial Intelligence*, pages 674–679, Vancouver, 1981.
- [37] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [38] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
- [39] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987.
- [40] Gary Mataev, Peyman Milanfar, and Michael Elad. Deepred: Deep image prior powered by red. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.

- [41] K. Ma, W. Liu, K. Zhang, Z. Duanmu, Z. Wang, and W. Zuo. End-to-end blind image quality assessment using deep neural networks. *IEEE Transactions on Image Processing*, 27(3):1202–1213, 2018.
- [42] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
- [43] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [44] Emmanuel J Candès, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [45] D. L. Donoho. Compressed sensing. *IEEE Trans Info Theory*, 52(4):1289–1306, 2006.
- [46] T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Transactions on Information Theory*, 55(4):1872–1882, 2009.
- [47] Z Wang, A C Bovik, H R Sheikh, and E P Simoncelli. Perceptual image quality assessment: From error visibility to structural similarity. *IEEE Trans Image Processing*, 13(4):600–612, Apr 2004.

## A Description of BF-CNN denoiser

**Architecture.** Throughout the paper, we use BF-CNN, described in [31], constructed from 20 bias-free convolutional layers, each consisting of  $3 \times 3$  filters and 64 channels, bias-free batch normalization, and a ReLU nonlinearity. Note that to construct a bias-free network, we remove all sources of additive bias, including the mean parameter of the batch-normalization in every layer.

**Training Scheme.** We follow the training procedure described in [31]. The network is trained to denoise images corrupted by i.i.d. Gaussian noise with standard deviations drawn from the range  $[0, 0.4]$  (relative to image intensity range  $[0, 1]$ ). The training set consists of overlapping patches of size  $40 \times 40$  cropped from the Berkeley Segmentation Dataset [42]. Each original natural image is of size  $180 \times 180$ . Training is carried out on batches of size 128, for 70 epochs.

## B Algorithms

---

**Algorithm 1: Sampling from prior.** Coarse-to-fine stochastic ascent method for sampling from the implicit prior of a denoiser, using denoiser residual  $f(y) = \hat{x}(y) - y$ .

---

parameters:  $\sigma_0, \sigma_L, h_0, \beta$   
 initialization:  $t = 1$ , draw  $y_0 \sim \mathcal{N}(0.5, \sigma_0^2 I)$   
**while**  $\sigma_{t-1} \leq \sigma_L$  **do**  
      $h_t = \frac{h_0 t}{1 + h_0(t-1)}$ ;  
      $d_t = f(y_{t-1})$ ;  
      $\sigma_t^2 = \frac{\|d_t\|^2}{N}$ ;  
      $\gamma_t^2 = ((1 - \beta h_t)^2 - (1 - h_t)^2) \sigma_t^2$ ;  
     Draw  $z_t \sim \mathcal{N}(0, I)$ ;  
      $y_t \leftarrow y_{t-1} + h_t d_t + \gamma_t z_t$ ;  
      $t \leftarrow t + 1$   
**end**

---



---

**Algorithm 2: Linear inverse problem.** Coarse-to-fine stochastic ascent method for sampling from  $p(x | M^T x = x^c)$ , based on the residual of a denoiser,  $f(y) = \hat{x}(y) - y$ . Note:  $e$  is an image of ones.

---

parameters:  $\sigma_0, \sigma_L, h_0, \beta, M, x^c$   
 initialization:  $t=1$ ; draw  $y_0 \sim \mathcal{N}(0.5(I - MM^T)e + Mx^c, \sigma_0^2 I)$   
**while**  $\sigma_{t-1} \leq \sigma_L$  **do**  
      $h_t = \frac{h_0 t}{1 + h_0(t-1)}$ ;  
      $d_t = (I - MM^T)f(y_{t-1}) + (Mx^c - MM^T y_{t-1})$ ;  
      $\sigma_t^2 = \frac{\|d_t\|^2}{N}$ ;  
      $\gamma_t^2 = ((1 - \beta h_t)^2 - (1 - h_t)^2) \sigma_t^2$ ;  
     Draw  $z_t \sim \mathcal{N}(0, I)$ ;  
      $y_t \leftarrow y_{t-1} + h_t d_t + \gamma_t z_t$ ;  
      $t \leftarrow t + 1$   
**end**

---



## D Visualization of Universal Inverse Sampler on a 2D manifold prior




Figure 2: Two-dimensional simulation/visualization of the Universal Inverse Sampler. Fifty example signals  $x$  are sampled from a uniform prior on a manifold (green curve). First three panels show, for three different levels of noise, the noise-corrupted measurements of the signals (red points), the associated noisy signal distribution  $p(y)$  (indicated with underlying grayscale intensities), and the least-squares optimal denoising solution  $\hat{x}(y)$  for each (end of red line segments), as defined by Eq. (2), or equivalently, Eq. (3). Right panel shows trajectory of our iterative coarse-to-fine inverse algorithm (Algorithm 2, depicted in Figure 1), starting from the same initial values  $y$  (red points) of the first panel. Algorithm parameters were  $h_0 = 0.05$  and  $\beta = 1$  (i.e., no injected noise). Note that, unlike the least-squares solutions, the iterative trajectories are curved, and always arrive at solutions on the signal manifold.

## E Convergence

Figure 3 illustrates the convergence of our iterative sampling algorithm, expressed in terms of the effective noise standard deviation  $\sigma = \frac{\|d_t\|}{\sqrt{N}}$  averaged over synthesis of three images, for three different levels of the stochasticity parameter  $\beta$ . Convergence is well-behaved and efficient in all cases. As expected, with smaller  $\beta$  (larger amounts of injected noise), effective standard deviation falls more slowly, and convergence takes longer.




Figure 3: Convergence of  $\sigma$  for three synthesized patches with different values of  $\beta$ . Injecting more noise in each iteration (i.e., smaller  $\beta$ ) slows down the convergence as expected.




Figure 4: Comparison of the computed effective noise,  $\sigma_t = \frac{\|d_t\|}{\sqrt{N}}$ , and the noise expected from the schedule  $\sigma_t = (1 - \beta h_t)\sigma_{t-1}$ , where  $h_t = h_0 \frac{t}{1+h_0(t-1)}$ . When  $\beta = 1$ , the effective noise estimated by the denoiser drops almost with the same rate as expected by the schedule. As  $\beta$  decreases, i.e. for non-zero additive noise, convergence to  $\sigma_L$  becomes faster than the scheduled rate. This accelerated convergence is probably because a portion of the additive noise is in line with the manifold, so does not contribute to calculated variance.

In addition to the total effective noise, we can compare the evolution of the removed noise versus injected noise. Figure 5 shows the reduction in effective standard deviation,  $h_t \sigma_t = h_t \frac{\|d_t\|}{\sqrt{N}}$  in each iteration, along with the standard deviation of the added noise,  $\gamma_t$ . The amount of noise added relative to the amount removed is such that effective noise drops as  $\sigma_t = (1 - \beta h_t)\sigma_{t-1}$ . When  $\beta = 1$ , the additive noise is zero,  $\gamma_t = 0$ , and the convergence of  $\sigma_t$  is the fastest. When  $\beta = 0.01$ , a lot of noise is added in each iteration, and the convergence is the slowest.




Figure 5: Temporal evolution of the amplitude of removed (solid) and injected (dashed) noise for the same patches as in Figure 3.  $\beta = 1$  corresponds to zero added noise ( $\gamma_t = 0$ ), hence the fastest convergence, while  $\beta = 0.01$  corresponds to a high level of added noise, hence slower convergence.

## F Image synthesis examples

We examined the properties of images synthesized using this stochastic coarse-to-fine ascent algorithm. For a denoiser, we used BF-CNN [31], a bias-free variant of DnCNN [28], trained on  $40 \times 40$  patches cropped from Berkeley segmentation training dataset (see Appendix A for further details). We obtained similar results (not shown) using other bias-free CNN denoisers (see [31]). For the sampling algorithm, we chose parameters  $\sigma_0 = 1$ ,  $\sigma_L = 0.01$ , and  $h_0 = 0.01$  for all experiments.

Figure 6 illustrates the iterative generation of two images, starting from different random initializations,  $y_0$ , with no additional noise injected (i.e.  $\beta = 1$ ), demonstrating the way that the algorithm amplifies and "hallucinates" structure found in the initial (noise) images. Convergence is typically achieved in less than 40 iterations with stochasticity disabled ( $\beta = 1$ ). The left panel in Figure 7 shows samples drawn with different initializations,  $y_0$ , using a moderate level of injected noise ( $\beta = 0.5$ ). Images contain natural-looking features, with sharp contours, junctions, shading, and




Figure 6: Sampling from the implicit prior. Each row shows a sequence of images,  $y_t, t = 1, 9, 17, 25, \dots$ , from the iterative sampling procedure, with two different initializations,  $y_0$ , and no added noise ( $\beta = 1$ ).




Figure 7: Samples arising from different initializations,  $y_0$ . Left: A moderate level of noise ( $\beta = 0.5$ ) is injected in each iteration. Right: A high level of injected noise ( $\beta = 0.1$ ).

in some cases, detailed texture regions. The right panel in Figure 7 shows a set of samples drawn with more substantial injected noise ( $\beta = 0.1$ ). The additional noise helps to avoid local maxima, and arrives at images that are smoother and higher probability, but still containing sharp boundaries. As expected, the additional noise also lengthens convergence time (see Figure 3). Figure 8 shows a set of samples drawn from the implicit prior of a denoiser trained on the MNIST dataset of handwritten digits.

## G Examples of inverse problems

We demonstrate the results of applying our method to several linear inverse problems. The same algorithm and parameters are used on all problems - only the measurement matrix  $M$  and measured values  $M^T x$  are altered. In particular, as in section F, we used BF-CNN [31], and chose parameters  $\sigma_0 = 1, \sigma_L = 0.01, h_0 = 0.01, \beta = 0.01$ . For each example, we show a row of original images ( $x$ ), a row of direct least-squares reconstructions ( $MM^T x$ ), and a row of restored images generated by our algorithm. For these applications, comparisons to ground truth are not particularly meaningful, at least when the measurement matrix is low-rank. In these cases, the algorithm relies heavily on the prior to "hallucinate" the missing information, and the goal is not so much to reproduce the original image, but to create an image that looks natural while being consistent with the measurements. Thus, the best measure of performance is a judgement of perceptual quality by a human observer.

**Inpainting.** A simple example of a linear inverse problem involves restoring a block of missing pixels, conditioned on the surrounding content. Here, the columns of the measurement matrix  $M$  are a subset of the identity matrix, corresponding to the measured (outer) pixel locations. We choose a missing block of size  $30 \times 30$  pixels, which is less than the size of the receptive field of the BF-CNN network ( $40 \times 40$ ), the largest extent over which this denoiser can be expected to capture joint statistical relationships. There is no single correct solution for this problem: Figure 9 shows




Figure 8: Training BF-CNN on the MNIST dataset of handwritten digits [43] results in a different implicit prior (compare to Figure 7). Each panel shows 16 samples drawn from the implicit prior, with different levels of injected noise (increasing from left to right,  $\beta \in \{1.0, 0.3, 0.01\}$ ).

multiple solutions, resulting from different initializations. Each appears plausible and consistent with the surrounding content. Figure 10 shows additional examples.

**Random missing pixels.** Consider a measurement process that discards a random subset of pixels.  $M$  is a low rank matrix whose columns consist of a subset of the identity matrix corresponding to the randomly chosen set of preserved pixels. Figure ?? shows examples with either 30% or 10% of pixels retained. Despite the significant number of missing pixels, the recovered images are remarkably similar to the originals.

**Spatial super-resolution.** In this problem, the goal is to construct a high resolution image from a low resolution (i.e. downsampled) image. Downsampling is typically performed after lowpass filtering, which determines the measurement model,  $M$ . Here, we use a  $4 \times 4$  block-averaging filter. Figure 11 shows examples.

**Deblurring (spectral super-resolution).** The previous applications were based on partial measurements in the pixel domain. Consider a blurring operator that operates by retaining a set of low-frequency coefficients in the Fourier domain, discarding the rest. In this case,  $M$  consists of the preserved low-frequency columns of the discrete Fourier transform, and  $MM^T x$  is a blurred version of  $x$ . Examples are shown in Figure 12.

**Compressive sensing.** Compressive sensing [44, 45] provides a set of theoretical results regarding recovery of sparse signals from a small number of linear measurements. Specifically, if one assumes that signals can be represented with at most  $k \ll N$  non-zero coefficients in a known basis, they can be recovered from a measurements obtained by projecting onto a surprisingly small number of axes (approaching  $k \log(N/k)$ ), far fewer than expected from traditional Shannon-Whitaker sampling theory. The theory relies on the sparsity property, which corresponds to a “union of subspaces” prior [46], and on the measurement axes being incoherent (essentially, weakly correlated) with the sparse basis. Typically, one chooses a sensing matrix containing a set of  $n \ll N$  random orthogonal axes. Recovery is achieved by solving the sparse inverse problem, using any of a number of methods.

Photographic images are not truly sparse in any fixed linear basis, but they can be reasonably approximated by low-dimensional subsets of Fourier or wavelet basis functions, and compressive sensing results are typically demonstrated using one of these. The manifold prior embedded within our CNN denoiser corresponds to a nonlinear form of sparsity, and analogous to sparse inverse algorithms used in compressed sensing, our stochastic coarse-to-fine ascent algorithm can be used to recover an image from a set of linear projections onto a set of random basis functions. Figure 13 shows four examples of images recovered from random projections using our denoiser-induced manifold prior, versus a sparse discrete cosine transform (DCT) prior. In all cases, the denoiser-recovered images exhibit sharper edges, more detail, and fewer artifacts. Numerical performance, in terms of both PSNR and SSIM, is consistent with the perceptual comparison.



Figure 9: Inpainting example. Left: original image. Next: corrupted image. Right three images: Restored examples, with different random initializations,  $y_0$ .




Figure 10: Inpainting examples. Top row: original images ( $x$ ). Middle: Images corrupted with blanked region ( $MM^T x$ ). Bottom: Images restored using our algorithm.



Figure 11: Spatial super-resolution. Resolution reduced by averaging over  $4 \times 4$  blocks (dimensionality reduction to 6.25%).




Figure 12: Deblurring (spectral super-resolution). Left three columns: Images blurred by retaining only 30% of low frequencies. Right three columns: retaining 10% of low frequencies.




Figure 13: Compressive sensing. Measurement matrix  $M$  contains random, orthogonal unit vectors, with dimensionality reduced to 30% for the first three columns, and 10% for last three columns. As in previous figures, top row shows original images ( $x$ ), and second row is linear pseudo-inverse ( $MM^T x$ ). Third row: images recovered using our method. Fourth row: standard compressive sensing solutions, assuming a sparse DCT signal model. Last row: standard compressive sensing solutions, assuming a sparse wavelet (db3) signal model. Numbers indicate performance, in PSNR(dB), and SSIM [47].