

Ill-Posed Image Reconstruction Without an Image Prior

Oscar Leong^{*1}, Angela F. Gao^{*1}, He Sun², and Katherine L. Bouman¹

¹California Institute of Technology

²Peking University

Abstract—We consider solving ill-posed imaging inverse problems without access to an image prior or ground-truth examples. An overarching challenge in these inverse problems is that an infinite number of images, including many that are implausible, are consistent with the observed measurements. Thus, image priors are required to reduce the space of possible solutions to more desirable reconstructions. However, in many applications it is difficult or potentially impossible to obtain example images to construct an image prior. Hence inaccurate priors are often used, which inevitably result in biased solutions. Rather than solving an inverse problem using priors that encode the spatial structure of any one image, we propose to solve a set of inverse problems jointly by incorporating prior constraints on the collective structure of the underlying images. The key assumption of our work is that the underlying images we aim to reconstruct share common, low-dimensional structure. We show that such a set of inverse problems can be solved simultaneously without the use of a spatial image prior by instead inferring a shared image generator with a low-dimensional latent space. The parameters of the generator and latent embeddings are found by maximizing a proxy for the Evidence Lower Bound (ELBO). Once identified, the generator and latent embeddings can be combined to provide reconstructed images for each inverse problem. The framework we propose can handle general forward model corruptions, and we show that measurements derived from only a small number of ground-truth images (≤ 150) are sufficient for “prior-free” image reconstruction. We demonstrate our approach on a variety of convex and non-convex inverse problems, ranging from denoising, phase retrieval, and black hole video reconstruction.

I. INTRODUCTION

In imaging inverse problems, the goal is to recover the underlying image from corrupted measurements, where the measurements and image are related via an understood forward model: $y = f(x) + \eta$. Here, y are measurements, x is the underlying image, f is a known forward model, and η is noise. Such problems are ubiquitous and include denoising [6, 17], super-resolution [7], compressed sensing [8, 16], phase retrieval [18], and deconvolution [32]. Due to corruption by the forward model and noise, these problems are

often ill-posed: there are many images that are consistent with the observed measurements, including ones that are implausible.

To combat the ill-posedness in imaging problems, solving for an image traditionally requires imposing additional structural assumptions to reduce the space of possible solutions. We encode these assumptions in an *image generation model (IGM)*, whose goal is to capture the desired properties of an image’s spatial structure. IGMs are general, as they can encompass probabilistic *spatial-domain* priors (e.g., that encourage smoothness or sparsity), but also deep image generators that are not necessarily probabilistic but are trained to primarily sample a certain class of images.

In order to define an IGM, it is necessary to have knowledge of the underlying image’s structure. If images similar to the underlying image are available, then an IGM can be learned directly [41, 52, 4]. However, an abundance of clean images is not available for many scientific imaging modalities (e.g., geophysical imaging and astronomical imaging). Collecting images in these domains can be extremely invasive, time-consuming, expensive, or even impossible. For instance, how should we define an IGM for black hole imaging without having ever seen a direct image of a black hole or knowing what one should look like? Moreover, classical approaches that utilize hand-crafted IGMs, such as total variation [23] or sparsity in a wavelet basis [34], are prone to human bias [31].

In this work, we show how one can solve a set of ill-posed image reconstruction tasks without prior information about an image’s spatial structure. The key insight of our work is that knowledge of common structure across multiple images is sufficient regularization alone. In particular, we suppose we have access to a collection of noisy measurements $\{y^{(i)}\}_{i=1}^N$ which are observed through (potentially different) forward models $y^{(i)} := f^{(i)}(x^{(i)}) + \eta^{(i)}$. The core assumption we make is that the underlying images $\{x^{(i)}\}_{i=1}^N$ are drawn from the same distribution (unknown *a priori*) and share common, low-dimensional structure. Thus, our “prior” is not at the

^{*}These authors contributed equally to this work.

spatial-level, but rather exploits the *collective structure* of the underlying images. This assumption is satisfied in a number of applications where there is not access to an abundance of clean images. For instance, although we might not know what a black hole looks like, we might expect it to be similar in appearance over time. We show that under this assumption, the image reconstruction posteriors $p(x|y^{(i)})$ can be learned jointly from a small number of examples $\{y^{(i)}\}_{i=1}^N$ due to the common, low-dimensional structure of the collection $\{x^{(i)}\}_{i=1}^N$. Specifically, our main result is that one can capitalize on this common structure by jointly solving for 1) a shared image generator G_θ and 2) N low-dimensional latent distributions $q_{\phi^{(i)}}$, such that the distribution induced by the push-forward of $q_{\phi^{(i)}}$ through G_θ approximately captures the image reconstruction posterior $p(x|y^{(i)})$ for each measurement example $i \in [N]$.

A. Our Contributions

We outline the main contributions of our work, which is an extension to our prior work presented in [19]:

- 1) We solve a collection of ill-posed inverse problems without prior knowledge of an image’s spatial structure by exploiting the common, low-dimensional structure shared across images. This common structure is exploited in inferring a shared IGM with a low-dimensional latent space.
- 2) To infer this IGM, we define a loss inspired by the evidence lower bound (ELBO). We motivate this loss by showing how it aids in “prior-free” image reconstruction by helping select one IGM from a collection of candidate IGMs using a single measurement example.
- 3) We apply our approach to convex and non-convex inverse problems, such as denoising, black hole compressed sensing, and phase retrieval. We establish that we can solve inverse problems without spatial-level priors and demonstrate good performance with only a small number of independent measurement examples (e.g., ≤ 150).
- 4) We theoretically analyze the inferred IGM in linear inverse problems under a linear image model to show that in this setting the inferred IGM performs dimensionality reduction akin to PCA on the collection of measurements.

II. BACKGROUND AND RELATED WORK

We now discuss related literature in model selection and learning-based IGMs. In order to highlight our key contributions, we emphasize the following assumptions in our framework:

- 1) We do not have access to a set of images from the same distribution as the underlying images.

- 2) We only have access to a collection of measurement examples, where each example comes from a different underlying image. The number of examples N is small, e.g., $N \leq 150$.
- 3) For each underlying image $x^{(i)}$ we wish to reconstruct, we only have access to a single measurement example $y^{(i)} = f^{(i)}(x^{(i)}) + \eta^{(i)}$. That is, we do not have multiple observations of the same underlying image. Note each $f^{(i)}$ can be potentially different.

A. Model selection

Model selection techniques seek to choose a model that best explains data by balancing performance and model complexity. In supervised learning problems with sufficiently large amounts of data, this can be achieved simply by evaluating the performance of different candidate models using reserved test data [46]. However, in image reconstruction or other inverse problems with limited data, one cannot afford to hold out data. In these cases, model selection is commonly conducted using probabilistic metrics. The simplest probabilistic metric used for linear model selection is adjusted R^2 [37]. It re-weights the goodness-of-fit by the number of linear model parameters, helping reject high-dimensional parameters that do not improve the data fitting accuracy. Similar metrics in nonlinear model selection are Bayesian Information Criterion (BIC) [42] and Akaike Information Criterion (AIC) [1]. AIC and BIC compute different weighted summations of a model’s log-likelihood and complexity, offering different trade offs between bias and variance to identify the best model for a given dataset.

In our work, we consider the use of the ELBO as a model selection criterion. In [9, 10], the use of the ELBO as a model selection criterion is theoretically analyzed and rates of convergence for variational posterior estimation are shown. Additionally, [48] proposes a generalized class of evidence lower bounds leveraging an extension of the evidence score. In [47], the ELBO is used for model selection to select a few, discrete parameters modeling a physical system (e.g., parameters that govern the orbit of an exoplanet). A significant difference in our context, however, is that we use the ELBO as a model selection criterion in a *high-dimensional* imaging context and we optimize the ELBO over a continuous space of possible parameters.

B. Learning IGMs

With access to a large corpus of example images it is possible to directly learn an IGM to help solve inverse problems. Seminal work along these lines utilizing generative networks showcased that a pre-trained generative adversarial network (GAN) can be used as an IGM in the problem of compressed sensing [4]. To

solve the inverse problem, the GAN was used to constrain the search space for inversion. This approach was shown to outperform sparsity-based techniques with 5-10x fewer measurements. Since then, this idea has been expanded to other inverse problems, including denoising [22], super-resolution [36], magnetic resonance imaging (MRI) [2, 44], and phase retrieval [20, 43]. However, the biggest downside to this approach is the requirement of a large dataset of example images similar to the underlying image, which is often difficult or impossible to obtain in practice. Hence we consider approaches that are able to directly solve inverse problems without example images.

Recently, methods that learn IGMs from only noisy measurements have been proposed. The main four distinctions between our work and these methods are that these works either: 1) require multiple independent observations of the *same* underlying image, 2) can only be applied to certain inverse problems, 3) require significantly more observations (either through more observations of each underlying image or by observing more underlying images), or 4) require significant hyperparameter tuning based on knowledge of example images.

Noise2Noise (N2N) [30] learns to denoise images by training on a collection of noisy, independent observations of the *same* image. To do so, N2N learns a neural network Φ_θ whose goal is to map between noisy images y and denoised images x . Since it has no denoised image examples to supervise training, it instead employs a loss that maps between noisy examples of the same underlying image. This objective is as follows:

$$\arg \min_{\theta} \sum_{i=1}^N \mathbb{E}_{y_1^{(i)} \sim \mathcal{Y}_1} \mathbb{E}_{y_2^{(i)} \sim \mathcal{Y}_2} [L(\Phi_\theta(y_1^{(i)}), y_2^{(i)})], \quad (1)$$

where $y_j^{(i)}$ corresponds to a noisy observation of the i -th underlying image $x^{(i)}$, and \mathcal{Y}_j is a distribution of noisy images where $\mathbb{E}_{y \sim \mathcal{Y}_j} [y] = x$. This N2N objective requires at least two observations of the same image and is limited by the assumption that the expected value of multiple observations of a single image is the underlying image. Thus, N2N is only applicable to denoising problems where the forward model is the identity matrix with independent noise on each pixel. Additionally, in practice N2N requires thousands of underlying images (i.e., N) to perform well. Thus, N2N’s main distinctions with our work are distinctions 1), 2), and 3).

Regularization by Artifact Removal (RARE) [33] generalizes N2N to perform image reconstruction from measurements under linear forward models. That is, the objective in Equation (1) is modified to include a pseudo-inverse. Nonetheless, multiple observations of the same underlying image are required, such that $\mathbb{E}_{y \sim \mathcal{Y}} [A^\dagger y] = x$ for the pseudo-inverse matrix A^\dagger . Thus, RARE suffers

from the same limiting distinctions as N2N (i.e., 1), 2), and 3)).

Noise2Void [27] and **Noise2Self** [3] assume that the image can be partitioned such that the measurement noise in one subset of the partition is independent conditioned on the measurements in the other subset. This is true for denoising, but not applicable to general forward models. For example, in black hole and MRI compressed sensing, it is not true that the measurement noise can be independently partitioned since each measurement is a linear combination of all pixels. While this makes Noise2Void and Noise2Self more restrictive in the corruptions they can handle compared to RARE, they also don’t require multiple observations of the same underlying image. Hence the main differences between these works and our own are distinctions 2) and 3).

AmbientGAN [5] and other similar approaches based on GANs [25] and VAEs [38, 35] have been proposed to learn an IGM directly from noisy measurements. For instance, AmbientGAN aims to learn a generator whose images lead to simulated measurements that are indistinguishable from the observed measurements; this generator can subsequently be used as a prior to solve inverse problems. However, AmbientGAN requires many measurement examples (on the order of 10,000) to produce a high quality generator. We corroborate this with experiments in Section IV to show that they require many independent observations and/or fine tuning of learning parameters to achieve good performance. Thus, the main distinctions between AmbientGAN and our work are 3) and 4).

Deep Image Prior (DIP) [51] uses a convolutional neural network as an implicit “prior”. DIP has shown strong performance across a variety of inverse problems to perform image reconstruction without explicit probabilistic priors. However, it is prone to overfitting and requires selecting a specific stopping criterion. While this works well when example images exist, selecting this stopping condition from noisy measurements alone introduces significant human bias that can negatively impact results. Thus, the main distinction between DIP and our work is 4).

III. APPROACH

In this work, we propose to solve a set of inverse problems without prior access to an IGM by assuming that the set of underlying images have common, low-dimensional structure. We motivate the use of optimizing the ELBO to infer an IGM by showing that it is a good criterion for generative model *selection* in Section III-A. Then, by optimizing the ELBO, we show in Section III-B that one can directly *infer* an IGM from corrupted measurements alone by parameterizing the image model as a deep generative network with a low-dimensional latent distribution. The IGM network

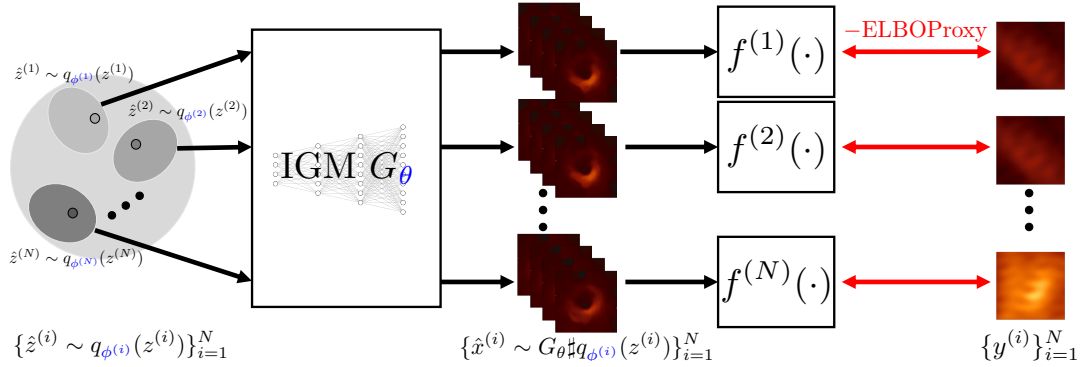


Fig. 1: We propose to solve a collection of N ill-posed inverse problems by exploiting the common, low-dimensional structure of the underlying images. Given a set of N measurement examples $\{y^{(i)}\}_{i=1}^N$ from N different underlying images, we propose to model each image posterior as the output of a shared IGM with a low-dimensional latent space. In particular, each posterior is approximated by $G_\theta \# q_{\phi^{(i)}}$, the push-forward of $q_{\phi^{(i)}}$ through G_θ , where G_θ is the shared, common generator to all N examples and $q_{\phi^{(i)}}$ is a low-dimensional, variational distribution particular to the i -th example. The inferred parameters, θ and $\{\phi^{(i)}\}_{i=1}^N$, are colored in blue and the loss is $-\text{ELBOProxy}$, which is given by Equation (4).

weights are shared across all images, capitalizing on the common structure present in the data, while the parameters of each latent distribution are learned jointly with the generator to model the image posteriors for each measurement example.

A. Motivation for ELBO as a model selection criterion

In order to accurately *infer* an IGM, we motivate the use of the ELBO as a loss by showing that it provides a principled criterion for *selecting* an IGM to use as a prior model. Suppose we are given noisy measurements from a single image: $y = f(x) + \eta$. In order to reconstruct the image x , we traditionally first require an IGM G that captures the distribution x was sampled from. A natural approach would be to find or select the model G that maximizes the model posterior distribution $p(G|y) \propto p(y|G)p(G)$. That is, conditioned on the noisy measurements, find the IGM of highest likelihood. Unfortunately computing $p(y|G)$ is intractable, as it requires marginalizing and integrating over all x encompassed by the IGM G . However, we show that this quantity can be well approximated using the ELBO.

To motivate our discussion, we first consider estimating the conditional posterior $p(x|y, G)$ by learning the parameters ϕ of a variational distribution $h_\phi(x)$. Observe that the definition of the KL-divergence followed by an application of Bayes' theorem gives

$$\begin{aligned} D_{\text{KL}}(h_\phi(x) \parallel p(x|y, G)) &:= \mathbb{E}_{x \sim h_\phi(x)} \left[\log \frac{h_\phi(x)}{p(x|y, G)} \right] \\ &= \mathbb{E}_{x \sim h_\phi(x)} \left[\log \frac{h_\phi(x)p(y|G)}{p(y|x, G)p(x|G)} \right] \\ &= -\mathbb{E}_{x \sim h_\phi(x)} [\log p(y|x, G) + \log p(x|G) - \log h_\phi(x)] \\ &\quad + \log p(y|G). \end{aligned}$$

The ELBO of an IGM G given measurements y under variational distribution h_ϕ is defined by

$$\begin{aligned} \text{ELBO}(G, h_\phi; y) &:= \mathbb{E}_{x \sim h_\phi(x)} [\log p(y|x, G) \\ &\quad + \log p(x|G) - \log h_\phi(x)]. \end{aligned} \quad (2)$$

Rearranging the previous equation, we see that by the non-negativity of the KL-divergence that

$$\begin{aligned} \log p(y|G) &= D_{\text{KL}}(h_\phi(x) \parallel p(x|y, G)) \\ &\quad + \text{ELBO}(G, h_\phi; y) \\ &\geq \text{ELBO}(G, h_\phi; y). \end{aligned}$$

Thus, we can lower bound the model posterior as

$$\log p(G|y) \geq \text{ELBO}(G, h_\phi; y) + \log p(G) - \log p(y).$$

Note that $\log p(y)$ is independent of the parameters of interest, ϕ . If the variational distribution $h_\phi(x)$ is a good approximation to the posterior $p(x|y, G)$, $D_{\text{KL}} \approx 0$. Thus, maximizing $\log p(G|y)$ with respect to G is approximately equivalent to maximizing $\text{ELBO}(G, h_\phi; y) + \log p(G)$.

Each term in the ELBO objective encourages certain properties of the IGM G . In particular, the first term in the ELBO, $\mathbb{E}_{x \sim h_\phi(x)} [\log p(y|x, G)]$, requires that G should lead to an image estimate that is consistent with our measurements y . The second term, $\mathbb{E}_{x \sim h_\phi(x)} [\log p(x|G)]$, encourages images sampled from $h_\phi(x)$ to have high likelihood under our model G . The final term is the entropy term, $\mathbb{E}_{x \sim h_\phi(x)} [-\log h_\phi(x)]$, which encourages a G that leads to “fatter” minima that are less sensitive to small changes in likely images x under G .

1) *ELBOProxy*: Some IGMs are explicit, which allows for direct computation of $\log p(x|G)$. For example, if our IGM models x as isotropic Gaussian with variance λ , then $-\log p(x|G) \propto \lambda^{-1} \|x\|_2^2$. In this case,

we can optimize the ELBO defined in Equation (2) directly and then perform model selection. However, an important class of IGMs that we are interested in are those given by deep generative networks. Such IGMs are not probabilistic in the usual Bayesian interpretation of a prior, but instead implicitly enforce structure in the data. A key characteristic of many generative network architectures (e.g., VAEs and GANs) that we leverage is that they generate high-dimensional images from low-dimensional latent representations. Bottlenecking helps the network learn global characteristics of the underlying image distribution while also respecting the low intrinsic dimensionality of natural images. However, this means that we can only compute $\log p(x|G)$ directly if we have an injective map [26]. This architectural requirement limits the expressivity of the network.

We instead consider a proxy of the ELBO that is especially helpful for deep generative networks. That is, suppose our IGM is of the form $x = G(z)$. Introducing a variational family for our latent representations $z \sim q_\phi(z)$ and choosing a latent prior distribution $\log p_Z(z|G)$, we arrive at the following proxy of the ELBO:

$$\text{ELBOProxy}(G, q_\phi; y) := \mathbb{E}_{z \sim q_\phi(z)} [\log p(y|G(z)) + \log p_Z(z|G) - \log q_\phi(z)]. \quad (3)$$

In our experiments, we chose $p_Z(z|G)$ to be an isotropic Gaussian prior. This is a common choice in many generative modeling frameworks and has shown to be a good choice of prior in the latent space.

To motivate this proxy, it is instructive to consider the case where our variational distribution $h_\phi := G\#q_\phi$ is the push-forward of a latent distribution q_ϕ through an injective or invertible function G . It follows from the definition of the push-forward that $x \sim G\#q_\phi$ if and only if $x = G(z)$ where $z \sim q_\phi$. While not all generators G will be injective, quality generators are largely injective over high likelihood image samples. In the case $h_\phi = G\#q_\phi$ for an injective function G , the ELBO and ELBOProxy are equivalent, as shown in the following proposition:

Proposition 1. *Suppose $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ is an injective function and let $p(x|G)$ be $G\#p_Z$. If $h_\phi = G\#q_\phi$ for a latent variational distribution q_ϕ , then*

$$\text{ELBO}(G, h_\phi; y) = \text{ELBOProxy}(G, q_\phi; y) \quad \forall y \in \mathbb{R}^m.$$

Proof. It suffices to show

$$\begin{aligned} \mathbb{E}_{x \sim h_\phi(x)} [\log p(x|G) - \log h_\phi(x)] \\ = \mathbb{E}_{z \sim q_\phi(z)} [\log p_Z(z|G) - \log q_\phi(z)] \end{aligned}$$

Let $J_G(z) \in \mathbb{R}^{n \times k}$ denote the Jacobian of G at an input z . Since G is injective, we can compute the likelihood

of a point in the range of G [26] via

$$\begin{aligned} \log p(x|G) &= \log p_Z(G^{-1}(x)|G) \\ &\quad - \frac{1}{2} \log |\det[J_G(G^{-1}(x))^T J_G(G^{-1}(x))]| \end{aligned}$$

where G^{-1} is the inverse of G on its range. Likewise, since h_ϕ is the push-forward of q_ϕ through G , we can compute the entropy as

$$\begin{aligned} \log h_\phi(x) &= \log q_\phi(G^{-1}(x)) \\ &\quad - \frac{1}{2} \log |\det[J_G(G^{-1}(x))^T J_G(G^{-1}(x))]|. \end{aligned}$$

Thus, for $x \in \text{range}(G)$, we have that

$$\begin{aligned} \log p(x|G) - \log h_\phi(x) \\ = \log p_Z(G^{-1}(x)|G) - \log q_\phi(G^{-1}(x)). \end{aligned}$$

Observe, that for $x \sim h_\phi(x) = G\#q_\phi(z)$, $x = G(z)$ for some $z \sim q_\phi(z)$, giving $G^{-1}(x) = G^{-1}(G(z)) = z$. Thus,

$$\begin{aligned} \mathbb{E}_{x \sim h_\phi(x)} [\log p(x|G) - \log h_\phi(x)] \\ = \mathbb{E}_{x \sim h_\phi(x)} [\log p_Z(G^{-1}(x)|G) - \log q_\phi(G^{-1}(x))] \\ = \mathbb{E}_{z \sim q_\phi(z)} [\log p_Z(G^{-1}(G(z))|G) - \log q_\phi(G^{-1}(G(z)))] \\ = \mathbb{E}_{z \sim q_\phi(z)} [\log p_Z(z|G) - \log q_\phi(z)]. \end{aligned}$$

□

An important consequence of this result is that for injective generators G , the inverse of G (on its range) is not required for computing the ELBO. In this case, the ELBOProxy is in fact equivalent to the ELBO. In Section III-A2 and Figure 2, we experimentally show that this proxy can aid in selecting potentially non-injective generative networks from corrupted measurements.

2) *Toy example:* To illustrate the use of the ELBOProxy as a model selection criterion, we conduct the following experiment that asks whether the ELBOProxy can identify the best model from a given set of image generation models. For this experiment, we use the MNIST dataset [29] and consider two inverse problems: denoising and phase retrieval. We train a generative model G_c on each class $c \in \{0, 1, 2, \dots, 9\}$ using the clean MNIST images directly. Hence, G_c generates images from class c via $G_c(z)$ where $z \sim \mathcal{N}(0, I)$. Then, given noisy measurements y_c from a single image from class c , we ask whether the generative model G_c from the appropriate class would achieve the best ELBOProxy. Each G_c is the decoder of a VAE with a low-dimensional latent space, with no architectural constraints to ensure injectivity. For denoising, our measurements are $y_c = x_c + \eta_c$ where $\eta_c \sim \mathcal{N}(0, \sigma^2 I)$ and $\sigma = \sqrt{0.5}$. For phase retrieval, $y_c = |\mathcal{F}(x_c)| + \eta_c$ where \mathcal{F} is the Fourier transform and $\eta_c \sim \mathcal{N}(0, \sigma^2 I)$ with $\sigma = \sqrt{0.05}$.

We construct 10×10 arrays for each problem, where in the i -th row and j -th column, we compute

the $-$ ELBOProxy obtained by using model G_{j-1} to reconstruct images from class $i-1$. We calculate $\text{ELBOProxy}(G_c, q_{\phi_c}; y_c)$ by parameterizing q_{ϕ_c} with a Normalizing Flow and optimizing network weights ϕ_c to maximize (3). The expectation in the ELBOProxy is approximated via Monte Carlo sampling. Results from the first 5 classes are shown in Fig. 2 and the full arrays are shown in the Supplemental. We note that all of the correct models are chosen in both denoising and phase retrieval. We also note some interesting cases where the ELBOProxy values are similar for certain cases, such as when recovering the 3 or 4 image. For example, when denoising the 4 image, both G_4 and G_9 achieve comparable ELBOProxy values. By carefully inspecting the noisy image of 4, one can see that both models are reasonable given the structure of the noise.

B. Simultaneously solving many inverse problems

As the previous section illustrates, the ELBOProxy provides a good criterion for choosing an appropriate IGM from noisy measurements. Here, we consider the task of directly inferring the IGM from a collection of measurement examples $y^{(i)} = f^{(i)}(x^{(i)}) + \eta^{(i)}$ for $i \in [N]$, where the parameters are found by optimizing the ELBOProxy. The key assumption we make is that common, low-dimensional structure is shared across the underlying images $\{x^{(i)}\}_{i=1}^N$. We propose to find a *shared* generator G_θ with weights θ along with latent distributions $q_{\phi^{(i)}}$ that can be used to reconstruct the full posterior of each image $x^{(i)}$ from its corresponding measurement example $y^{(i)}$. This approach is illustrated in Fig. 1. Having the generator be shared across all images helps capture their common collective structure. Each forward model corruption, however, likely induces its own complicated image posteriors. Hence, we assign each measurement example $y^{(i)}$ its own latent distribution to capture the differences in their posteriors. Note that because we optimize a proxy of the ELBO, the inferred distribution may not necessarily be the true image posterior, but it still captures a distribution of images that fit to the observed measurements.

a) Inference approach: More explicitly, given a collection of measurement examples $\{y^{(i)}\}_{i=1}^N$, we jointly infer a generator G_θ and a set of variational distributions $\{q_{\phi^{(i)}}\}_{i=1}^N$ by optimizing a Monte Carlo estimation of the ELBOProxy from Equation (3), described by:

$$\max_{\theta, \phi^{(i)}} \frac{1}{N} \sum_{i=1}^N \text{ELBOProxy}(G_\theta, q_{\phi^{(i)}}; y^{(i)}) + \log p(G_\theta). \quad (4)$$

In terms of choices for $\log p(G_\theta)$, we can add additional regularization to promote particular properties of the IGM G_θ , such as having a small Lipschitz constant.




Fig. 2: We consider two inverse problems: denoising and phase retrieval. Top: the two topmost rows correspond to the ground truth image x_c and the noisy measurements y_c . Center: in each column, we show the means of the distribution induced by the push-forward of G_j and each latent distribution $z \sim q_{\phi_j}$ for $j \in \{0, \dots, 9\}$. Bottom: each column of the array corresponds to the $-$ ELBOProxy achieved by each model in reconstructing the images. Here, lower is better. Boxes highlighted in green correspond to the best $-$ ELBOProxy values in each column. In all these examples, the correct model was chosen.

Here, we consider having sparse neural network weights as a form of regularization and use dropout [45] during training to represent $\log p(G_\theta)$.

Once a generator G_θ and variational parameters $\phi^{(i)}$ have been inferred, we solve the i -th inverse problem by simply sampling $\hat{x}^{(i)} = G_\theta(\hat{z}^{(i)})$ where $\hat{z}^{(i)} \sim q_{\phi^{(i)}}(z^{(i)})$ or computing an average $\bar{x}^{(i)} = \frac{1}{T} \sum_{t=1}^T G_\theta(\hat{z}_t^{(i)})$. Producing samples for each inverse problem can help visualize the range of uncertainty

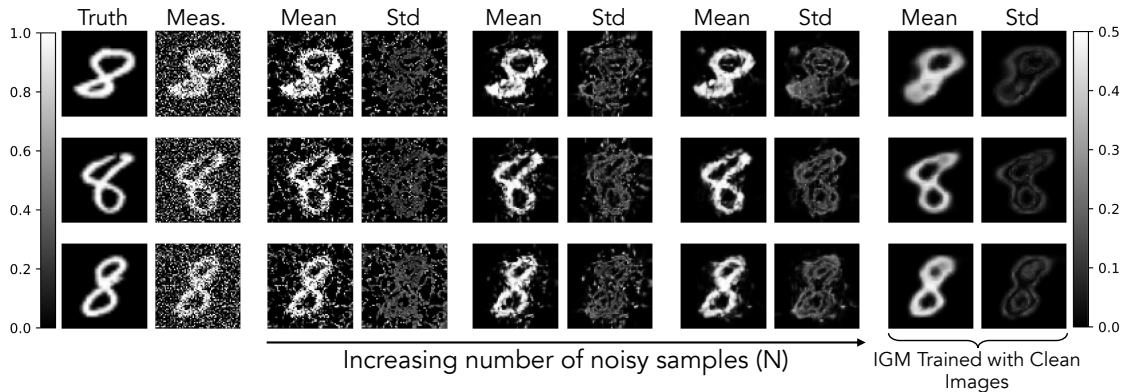


Fig. 3: **Improvement with an increasing number of noisy observations.** We demonstrate our method of inferring an IGM to perform denoising for increasing number of noisy MNIST images (5, 75, and 150 images from left to right). We showcase results on three randomly selected examples that appear in each collection of inverse problems. In each panel, we include the ground truth, noisy measurements, mean of the posterior, and standard deviation of the posterior. We also include the reconstructions using an IGM trained on the full clean MNIST 8’s class (6000 images). We observe that the mean reconstructions and standard deviations from our low-data IGMs become more similar to the full-data IGM with increasing data.



Fig. 4: **Denoising 95 images of a celebrity.** We demonstrate our method described in III-B using 95 noisy images of a celebrity. Here we show the underlying image (top), noisy measurements (middle), and mean reconstruction (bottom) for a subset of the 95 different noisy images. Our reconstructions are much less noisy than the measurements and exhibit sharper features that are hard to discern in the noisy images. Note that no spatial-domain prior/regularizer was used in denoising.

under the learned IGM G_θ , while the expected value of the distribution empirically provides clearer estimates with better metrics in terms of PSNR or MSE. We report PSNR outputs in our subsequent experiments and also visualize the standard deviation of our reconstructions.

IV. EXPERIMENTAL RESULTS

We now consider solving a set of inverse problems via the framework described in III-B. For each of these experiments, we use a multivariate Gaussian distribution to parameterize each of the posterior distributions $q_{\phi^{(i)}}$ and a Deep Decoder [21] with 6 layers of size 150, a latent size of 40, and a dropout of 10^{-4} as the IGM. The multivariate Gaussian distributions are parameterized by means and covariance matrices $\{\mu^{(i)}, \Lambda^{(i)} = U_i U_i^T + \varepsilon I\}_{i=1}^N$, where εI with $\varepsilon = 10^{-3}$ is added to the covariance matrix to help with stability of the optimization. We choose to parameterize the latent distributions using Gaussians for memory considerations. Note that

the same hyperparameters are used for all experiments demonstrating our proposed method.

In our experiments, we also compare to the following baseline methods: AmbientGAN [5], Deep Image Prior (DIP) [50], and regularized maximum likelihood using total variation (TV-RML). AmbientGAN is most similar to our setup, as it constructs an IGM directly from measurement examples; however, it doesn’t aim to estimate image reconstruction posteriors, but instead aims to learn an IGM that samples from the full underlying prior. TV-RML uses an explicit total variation regularization, while DIP uses an implicit convolutional neural network “prior”. As we will show, all these baseline methods require fine-tuning hyperparameters to each set of measurements in order to produce the best results. All methods can also be applied to a variety of inverse problems, making them appropriate choices as baselines.

A. Denoising

We show results on denoising a collection of noisy images of 8’s from the MNIST dataset in Fig. 3 and de-

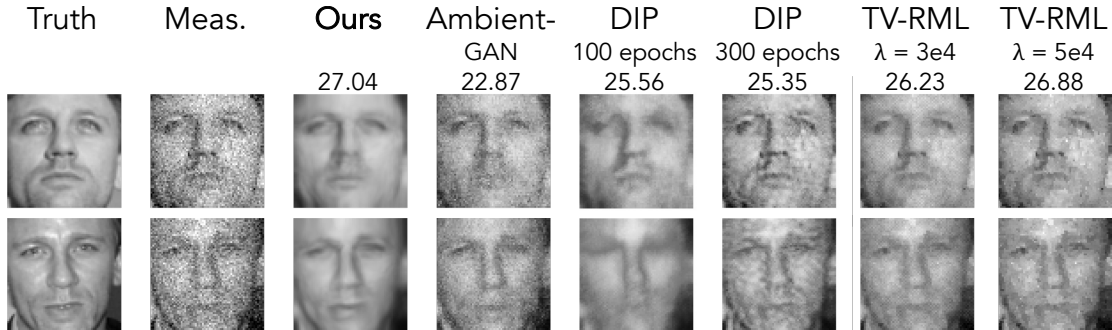


Fig. 5: **Denoising baseline comparisons.** We compare to various baselines (AmbientGAN, Deep Image Prior (DIP), and regularized maximum likelihood using TV (TV-RML) with weight λ), and we report the average PSNR across all 95 reconstructions. We show both early stopping and full training results using DIP. Our method exhibits higher PSNR than all other baselines while maintaining distinct features of the underlying images. We include a few potential results for baselines that require fine-tuning to demonstrate sensitivity to subjective stopping conditions.




Fig. 6: **Phase retrieval from MNIST 8's.** We demonstrate our method described in III-B to perform phase retrieval on 150 images. For the Fourier phase retrieval setting, we show examples from the two observed modes of the posterior. For the Gaussian phase retrieval setting, we show the mean and standard deviations of our reconstructions.

noising a collection of noisy images of a single face from the PubFig [28] dataset in Fig. 4. The measurements for both are defined by $y = x + \eta$ where $\eta \sim \mathcal{N}(0, \sigma^2 I)$ with an SNR of ~ 3 dB for the MNIST digits and an SNR of ~ 15 dB for the faces. Our method is able to remove much of the added noise and recovers small scale features, even with only 10's of observations. As shown in Fig. 3, the reconstructions achieved under the learned IGM improves as the number of independent observations increases. Our reconstructions also substantially outperform the baseline methods, as shown in Fig. 5. Unlike DIP, our method does not overfit and does not require early stopping. Our method does not exhibit noisy artifacts like those seen in all baseline methods, despite such methods being fine-tuned.

B. Phase retrieval

Here we consider solving non-convex inverse problems, and demonstrate our approach on phase retrieval. Our measurements are described by $y = |\mathcal{F}(x)| + \eta$ where $\mathcal{F}(\cdot)$ is a linear operator and $\eta \sim \mathcal{N}(0, \sigma^2 I)$. We

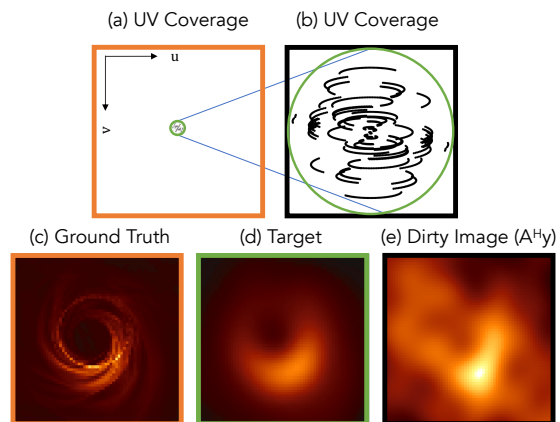


Fig. 7: **Visualization of the intrinsic resolution of the EHT compressed sensing measurements.** The EHT measures sparse spatial frequencies of the image (i.e., components of the image's Fourier transform). In order to generate the underlying image (c), all frequencies in the entire domain of (a) must be used. Restricting spatial frequencies to the ones in (a) and (b)'s green circle generates the target (d), where (b) is a zoom in of (a). The EHT samples a subset the interior of the green circle, indicated by the sparse black samples in (b). Naively recovering an image using only these frequencies results in the *dirty image* (e), which is computed by $A^H y$. The 2D spatial Fourier frequency coverage represented with (u, v) positions is referred to as the UV coverage.

consider two types of measurements, one where $\mathcal{F}(\cdot)$ is the Fourier transform and the other when $\mathcal{F}(\cdot)$ is an $m \times n$ complex Gaussian matrix with $m = \lceil 0.1n \rceil$. Since each measurement is the magnitude of complex linear measurements, there is an inherent phase ambiguity in the problem. Additionally, flipping and spatial-shifts are possible reconstructions when performing Fourier phase retrieval. Due to the severe ill-posedness of the problem, representing this complicated posterior that includes all spatial shifts is challenging. Thus, we incorporate an envelope (i.e., a centered rectangular mask) as the final

layer of G_θ to encourage the reconstruction to be centered. Nonetheless, flipping and shifts are still possible within this enveloped region.

We show results from a set of $N = 150$ noisy phase retrieval measurements from the MNIST 8’s class with a SNR of ~ 52 dB in Fig. 6. In the Gaussian case, we note that our mean reconstructions are nearly identical to the true digits and the standard deviations exhibit uncertainty in regions we would expect (e.g., around edges). In the Fourier case, our reconstructions have features similar to the digit 8, but contain artifacts. These artifacts are only present in the Fourier case due to additional ambiguities, which lead to a more complex posterior [24].

C. Black hole imaging

We consider a real-world inverse problem for which ground-truth data would be impossible to obtain. In particular, we consider a compressed sensing problem inspired by astronomical imaging of black holes with the Event Horizon Telescope (EHT): suppose we are given access to N measurement examples of the form $y^{(i)} = A^{(i)}x^{(i)} + \eta^{(i)}$ where $A^{(i)} \in \mathbb{C}^{m \times n}$ is a low-rank compressed sensing matrix arising from interferometric telescope measurements and $\eta^{(i)}$ denotes noise with known properties (e.g., distributed as a zero-mean Gaussian with known standard deviation). The collection of images $\{x^{(i)}\}_{i=1}^N$ are snapshots of an evolving black hole target. This problem is ill-posed and requires the use of priors or regularizers to recover a reasonable image [12]. Moreover, it is impossible to directly acquire example images of black holes, so any pixel-level prior defined *a priori* will exhibit human bias. Recovering an image, or movie, of a black hole with as little human bias as possible is essential for both studying the astrophysics of black holes as well as testing fundamental physics [11, 14]. We show how our proposed method can be used to tackle this important problem. In particular, we leverage knowledge that, although the black hole evolves, it will not change drastically from minute-to-minute or day-to-day. We study two black hole targets: the black hole at the center of the Messier 87 galaxy (M87*) and the black hole at the center of the Milky Way galaxy – Sagittarius A* (Sgr A*).

a) Imaging M87* using the current EHT array:

We first consider reconstructing the black hole at the center of the Messier 87 galaxy, which does not evolve noticeably within the timescale of a single day. The underlying images are from a simulated 60 frame video with a single frame for each day. We show results on $N = 60$ frames from an evolving black hole target with a diameter of ~ 40 microarcseconds, as was identified as the diameter of M87 according to [39, 13] in Fig. 8. In particular, the measurements are given by $\{y^{(i)} = Ax^{(i)} + \eta^{(i)}\}_{i=1}^N$, where $x^{(i)}$ is the underlying image on day i , A is the forward model that represents the

telescope array, which is static across different days, and the noise $\eta^{(i)} \sim \mathcal{N}(0, \Sigma)$ has a covariance of Σ , which is a diagonal matrix with realistic variances based on the telescope properties*. Measurements are simulated from black hole images with a realistic flux of 1 Jansky. We also visualize a reference “target” image, which is the underlying image filtered with a low-pass filter that represents the maximum resolution achievable with the telescope array used to collect measurements – in this case the EHT array consisting of 11 telescopes (see Fig. 7 and Section IX in the Supplement).

As seen in Fig. 8, our method is not only able to reconstruct the large scale features of the underlying image without any aliasing artifacts, but also achieves a level of super-resolution. Our reconstructions also achieve higher super-resolution as compared to our baselines (i.e., AmbientGAN, TV-RML, and DPI) in Fig. 8 and do not exhibit artifacts evident in the reconstructions from these baselines. The two AmbientGAN settings were qualitatively chosen to show that the final result is sensitive to the choice of hyperparameters. The default AmbientGAN parameters produce poor results, and even with fine-tuning to best fit the underlying images (e.g., cheating with knowledge of the ground-truth), the results still exhibit substantial artifacts. Our results demonstrate that we are able to capture the varying temporal structure of the black hole, rather than just recovering a static image. It is important to note that there is no explicit temporal regularization introduced; the temporal smoothness is implicitly inferred by the constructed IGM.

b) Imaging Sgr A* from multiple forward models:

The framework we introduce can also be applied to situations in which the measurements themselves are induced by different forward models. In particular, the measurements $\{y^{(i)} = f^{(i)}(x^{(i)}) + \eta^{(i)}\}_{i=1}^N$ are given by an underlying image $x^{(i)}$, a forward model $f^{(i)}$ that is specific to that observation, and noise $\eta^{(i)}$ with known properties.

As an illustrative example, we consider the problem of reconstructing a video of the black hole at the center of the Milky Way – Sagittarius A* (Sgr A*). Unlike M87*, Sgr A* evolves on the order of minutes. Therefore, we can only consider that the black hole is static for only a short time when only a subset of telescopes are able to observe the black hole. This results in a different measurement forward model for each frame of the black hole “movie.” [15]. In particular, the measurements are given by $\{y^{(i)} = A^{(i)}x^{(i)} + \eta^{(i)}\}_{i=1}^N$, where $x^{(i)}$ is the underlying image at time i , $A^{(i)}$ is the forward model that incorporates the telescope configuration at that time, and $\eta^{(i)} \sim \mathcal{N}(0, \Sigma^{(i)})$ is noise where $\Sigma^{(i)}$ is a diagonal matrix with realistic standard deviations derived from the

*We leave the more challenging atmospheric noise that appears on measurements to future work.




Fig. 8: **Recovering a 60-day video of the M87 black hole with the EHT telescope array.** Left: We demonstrate our method described in III-B using simulated measurements from 60 frames of an evolving black hole target with a forward model described in Fig. 7. Here we show the underlying images, dirty images ($A^H y$, see Fig. 7), and mean reconstruction, respectively. Additionally, we show the unwrapped space \times time image, which is taken along the overlaid white ring illustrated in the day 1 underlying image. The bright-spot’s temporal trajectory of our reconstruction matches that of the truth. Right: We compare our method to various baselines methods. Our results are much sharper and exhibit less artifacts than AmbientGAN and TV-RML with weight λ . Note that include results using AmbientGAN using both default parameters and fine-tuned parameters.

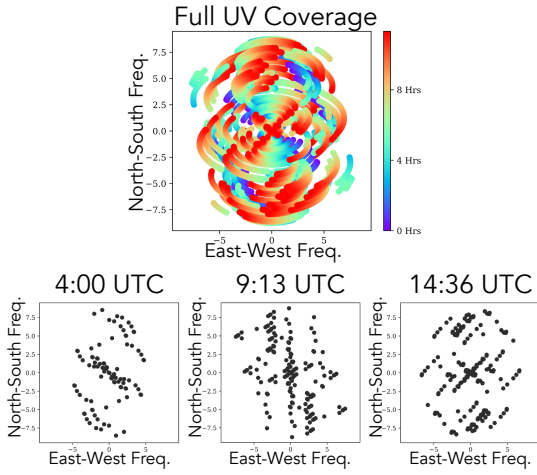


Fig. 9: **Visualization of the EHT forward model of Sgr A* over time.** The spatial frequency space coordinates of Sgr A* from the ngEHT [40] array measured over a course of one night is shown on the top. Snapshots of the coordinates that are measured at different times are shown on the bottom. Note that the location and sparsity of the measurement coordinates changes over time.

telescopes’ physical properties. The measurement noise is consistent with a black hole with a flux of 2 Janskys. The measurement operator is illustrated in Fig. 9.

We show examples of reconstructing 60 frames of Sgr A* with a diameter of ~ 50 microarcseconds using measurements simulated from a proposed future next-generation EHT (ngEHT) [40] array, which consists of 23 telescopes. These results are shown in Fig. 10. Our reconstructions remove much of the aliasing artifacts evident in the dirty images and reconstructs the primary features of the underlying image without any form of temporal regularization. These results have high fidelity especially considering that the measurements are very

sparsely sampled. Although these results come from simulated measurements that do not account for all forms of noise we expect to encounter, the high-quality movie reconstructions obtained without the use of a spatio-temporal prior show great promise towards scientific discovery that could come from a future telescope array paired with our proposed reconstruction approach.

V. THEORY FOR LINEAR IGMs

We now introduce theoretical results on the inferred IGM for linear inverse problems. Specifically, we consider the case when the IGM G_θ is linear and the latent variational distributions are Gaussian. The goal of this section is to develop intuition for the inferred IGM in a simpler setting. While our results may not generalize to non-linear generators parameterized by deep neural networks, our results aim to provide an initial understanding on the inferred generator.

More concretely, suppose we are given N measurement examples of noisy linear measurements of the form

$$y^{(i)} := Ax^{(i)} + \eta^{(i)}, \quad \eta^{(i)} \sim \mathcal{N}(0, \sigma^2 I)$$

where $A \in \mathbb{R}^{m \times n}$ with $m \leq n$ and $x^{(i)} \in \mathbb{R}^n$. We aim to infer $G_\theta \in \mathbb{R}^{n \times k}$ with $k < n$ and latent distributions $q_{\phi^{(i)}} = \mathcal{N}(\mu_i, U_i U_i^T)$ where $\phi^{(i)} = \{\mu_i, U_i\}$, $\mu_i \in \mathbb{R}^k$, and $U_i \in \mathbb{R}^{k \times k}$ by minimizing the negative ELBOProxy:

$$\mathcal{L}_N(G_\theta, \{\phi^{(i)}\}) := \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{z \sim q_{\phi^{(i)}}(z)} [-\log p(y^{(i)} | G_\theta(z)) - \log p_Z(z | G_\theta) + \log q_{\phi^{(i)}}(z)]. \quad (5)$$

We characterize the generator G_* and latent parameters $\phi_*^{(i)} = \{\mu_i^*, U_i^*\}$ that are stationary points of (5). The result is proven in Section X:

Theorem 1. Fix $\sigma > 0$ and let $y^{(i)} = Ax^{(i)} + \eta^{(i)} \in \mathbb{R}^m$ for $i \in [N]$ where $\eta^{(i)} \sim \mathcal{N}(0, \sigma^2 I)$ and $A \in \mathbb{R}^{m \times n}$ has




Fig. 10: **Recovering a video of the Sgr A* black hole with a futuristic telescope array over the course of a night.** We demonstrate our method described in III-B to perform video reconstruction on 60 frames of a video where the measurements for each frame is generated by different forward models given by imaging Sgr A*. Here we show the underlying image, the dirty image ($A^H y$), the mean of the reconstructed posterior, and a visualization of the measurement coverage in frequency (UV) space. Additionally, we show the unwrapped space \times time image, which is taken along the white ring illustrated in the 4:00 UTC underlying image. The bright-spot’s temporal trajectory of our reconstruction matches that of the truth. The measurement noise is consistent with a black hole with a flux of 2 Janskys.

$k < \text{rank}(A) = m \leq n$. Define the sample covariance of the measurements $Y_N := \frac{1}{N} \sum_{i=1}^N y^{(i)}(y^{(i)})^T$. Then, with probability 1, G_* , μ_i^* , and U_i^* that satisfy

$$\begin{aligned} AG_* &:= E_k(L_k - \sigma^2 I)^{1/2} R, \\ \mu_i^* &:= (G_*^T A^T AG_* + \sigma^2 I)^{-1} G_*^T A^T y^{(i)}, \text{ and} \\ U_i^* &:= \sigma(G_*^T A^T AG_* + \sigma^2 I)^{-1/2} \tilde{R}, \end{aligned}$$

for $i \in [N]$ are stationary points of the objective (5) over $\mathbb{R}^{n \times k} \times (\mathbb{R}^k \times \text{GL}(k, \mathbb{R}))^N$ where $\text{GL}(k, \mathbb{R})$ denotes the set of real invertible $k \times k$ matrices. Here, $E_k \in \mathbb{R}^{m \times k}$ denotes the matrix whose columns correspond to the top- $\min\{\text{rank}(Y_N), k\}$ eigenvectors of Y_N , $L_k \in \mathbb{R}^{k \times k}$ contains the corresponding eigenvalues of Y_N , and $R, \tilde{R} \in \mathbb{R}^{k \times k}$ are arbitrary orthogonal matrices. If $\text{rank}(Y_N) < k$, then for $\text{rank}(Y_N) < i \leq k$, the i -th column of E_k can be arbitrary and $L_{k,ii} = \sigma^2$.

The Theorem establishes the precise form of stationary points of the objective (5). In particular, it shows that this inferred IGM performs dimensionality reduction akin to PCA [49] on the collection of measurement examples. To gain further intuition about the Theorem, in Section X-B of the Supplemental, we analyze this result in a context where the underlying images we wish to reconstruct explicitly lie in a low-dimensional subspace. We show that, in that setting, our estimator returns an approximation of the solution found via MAP estimation, which can only be computed with complete prior knowledge of the underlying image structure. While the Theorem focused on linear IGMs, it will be interesting to theoretically analyze non-linear IGMs parameterized by deep networks. We leave this for future work.

VI. CONCLUSION

In this work, we showcased how one can solve a set of inverse problems without a pre-defined IGM (e.g., image prior) by leveraging common structure present across the underlying images. We demonstrated that even with a small set of corrupted measurements, one can jointly solve these inverse problems by directly inferring an IGM that maximizes a proxy of the ELBO. We demonstrate our method on a number of convex and non-convex imaging problems, including the challenging problem of black hole video reconstruction from interferometric measurements. Overall, our work showcases the possibilities of solving inverse problems in a “prior-free” fashion, free from significant human bias typical of ill-posed image reconstruction. We believe our approach can aid automatic discovery of novel structure from scientific measurements, leading to potentially new avenues for scientific discovery.

REFERENCES

- [1] Hirotugu Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.
- [2] Marcus Alley, Joseph Y. Cheng, William Dally, Enhao Gong, Song Hann, Morteza Mardani, John M. Pauly, Neil Thakur, Shreyas Vasanawala, Lei Xing, and Greg Zaharchuk. Deep generative adversarial networks for compressed sensing (gancs) automates mri. *arXiv preprint arXiv:1706.00051*, 2017.
- [3] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning*, pages 524–533. PMLR, 2019.

- [4] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros Dimakis. Compressed sensing using generative models. *International Conference on Machine Learning (ICML)*, 2017.
- [5] Ashish Bora, Eric Price, and Alexandros G Dimakis. Ambientgan: Generative models from lossy measurements. In *International conference on learning representations*, 2018.
- [6] Harold C. Burger, Christian J. Schuler, and Stefan Harmeling. Image denoising: Can plain neural networks compete with bm3d? *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [7] Emmanuel J. Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2013.
- [8] 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.
- [9] Badr-Eddine Chérif-Abdellatif. Consistency of elbo maximization for model selection. *Proceedings of The 1st Symposium on Advances in Approximate Bayesian Inference*, 96:1–21, 2018.
- [10] Badr-Eddine Chérif-Abdellatif and Pierre Alquier. Consistency of variational bayes inference for estimation and model selection in mixtures. *Electronic Journal of Statistics*, 12(2):2995–3035, 2018.
- [11] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. i. the shadow of the supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L1, Apr 2019.
- [12] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. iv. imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019.
- [13] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. v. physical origin of the asymmetric ring. *The Astrophysical Journal Letters*, 875(1):L5, 2019.
- [14] The Event Horizon Telescope Collaboration. First sagittarius a* event horizon telescope results. i. the shadow of the supermassive black hole in the center of the milky way. *The Astrophysical Journal Letters*, 930(2):L12, May 2022.
- [15] The Event Horizon Telescope Collaboration. First sagittarius a* event horizon telescope results. iii. imaging of the galactic center supermassive black hole. *The Astrophysical Journal Letters*, 930(2):L14, 2022.
- [16] David Donoho. For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59(6), 2006.
- [17] Laura Dyer, Andrew Parker, Keanu Paphiti, and Jeremy Sanderson. Lightsheet microscopy. *Current Protocols*, 2(7):e448, 2022.
- [18] Albert Fannjiang and Thomas Strohmer. The numerics of phase retrieval. *Acta Numerica*, 29:125 – 228, 2020.
- [19] Angela F. Gao, Oscar Leong, He Sun, and Katie L. Bouman. Image reconstruction without explicit priors. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2023.
- [20] Paul Hand, Oscar Leong, and Vladislav Voroninski. Phase retrieval under a generative prior. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [21] Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. *International Conference on Learning Representations (ICLR)*, 2019.
- [22] Reinhard Heckel, Wen Huang, Paul Hand, and Vladislav Voroninski. Rate-optimal denoising with deep neural networks. *Information and Inference: A Journal of the IMA*, 2020.
- [23] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992.
- [24] Kishore Jaganathan, Samet Oymak, and Babak Hassibi. Sparse phase retrieval: Convex algorithms and limitations. *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 1022–1026, 2013.
- [25] Maya Kabkab, Pouya Samangouei, and Rama Chellappa. Task-aware compressed sensing with generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [26] Konik Kothari, AmirEhsan Khorashadizadeh, Maarten de Hoop, and Ivan Dokmanić. Trumpets: Injective flows for inference and inverse problems. In *Uncertainty in Artificial Intelligence*, pages 1269–1278. PMLR, 2021.
- [27] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void - learning denoising from single noisy images. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2129–2137, 2019.
- [28] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar. Attribute and simile classifiers for face verification. In *2009 IEEE 12th international conference on computer vision*, pages 365–372. IEEE, 2009.
- [29] Yann LeCun, Leon Bottou, Yoshua Bengio, and

- Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [30] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. *Proceedings of the 35th International Conference on Machine Learning*, 80:2965–2974, 2018.
- [31] Anat Levin, Yair Weiss, Fredo Durand, and William T. Freeman. Understanding and evaluating blind deconvolution algorithms. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1964–1971, 2009.
- [32] Xiaodong Li, Shuyang Ling, Thomas Strohmer, and Ke Wei. Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *arXiv preprint*, arXiv:1606.04933, 2016.
- [33] Jiaming Liu, Yu Sun, Cihat Eldeniz, Weijie Gan, Hongyu An, and Ulugbek S Kamilov. Rare: Image reconstruction using deep priors learned without groundtruth. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1088–1099, 2020.
- [34] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.
- [35] Ray Mendoza, Minh Nguyen, Judith Weng Zhu, Vincent Dumont, Talita Perciano, Juliane Mueller, and Vidya Ganapati. A self-supervised approach to reconstruction in sparse x-ray computed tomography. In *NeurIPS Machine Learning and the Physical Sciences Workshop*, 2022.
- [36] Sachit Menon, Alex Damian, McCourt Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [37] Jeremy Miles. R squared, adjusted r squared. *Encyclopedia of Statistics in Behavioral Science*, 2005.
- [38] Andrew Olsen, Yolanda Hu, and Vidya Ganapati. Data-driven computational imaging for scientific discovery. In *NeurIPS AI for Science Workshop*, 2022.
- [39] Oliver Porth, Koushik Chatterjee, Ramesh Narayan, Charles F Gammie, Yosuke Mizuno, Peter Anninos, John G Baker, Matteo Bugli, Chi-kwan Chan, Jordy Davelaar, et al. The event horizon general relativistic magnetohydrodynamic code comparison project. *The Astrophysical Journal Supplement Series*, 243(2):26, 2019.
- [40] Alexander W Raymond, Daniel Palumbo, Scott N Paine, Lindy Blackburn, Rodrigo Córdoba Rosado, Sheperd S Doleman, Joseph R Farah, Michael D Johnson, Freek Roelofs, Remo PJ Tilanus, et al. Evaluation of new submillimeter vlbi sites for the event horizon telescope. *The Astrophysical Journal Supplement Series*, 253(1):5, 2021.
- [41] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- [42] Gideon E. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.
- [43] Fahad Shamshad and Ali Ahmed. Compressed sensing-based robust phase retrieval via deep generative priors. *IEEE Sensors Journal*, 21(2):2286 – 2298, 2017.
- [44] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *International Conference on Learning Representations (ICLR)*, 2022.
- [45] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [46] Mervyn Stone. Cross-validators choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)*, 36(2):111–133, 1974.
- [47] He Sun, Katherine L. Bouman, Paul Tiede, Jason J. Wang, Sarah Blunt, and Dimitri Mawet. alpha-deep probabilistic inference (alpha-dpi): efficient uncertainty quantification from exoplanet astrometry to black hole feature extraction. *The Astrophysical Journal*, 932(2):99, 2022.
- [48] Chenyang Tao, Liqun Chen, Ruiyi Zhang, Ricardo Henao, and Lawrence Carin. Variational inference and model selection with generalized evidence bounds. *International Conference on Machine Learning (ICML)*, 80:893–902, 2018.
- [49] Michael E. Tipping and Christopher M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [50] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [51] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *International Journal of Computer Vision*, 128:1867–1888, 2020.
- [52] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE, 2013.

CONTENTS

I	Introduction	1
I-A	Our Contributions	2
II	Background and Related Work	2
II-A	Model selection	2
II-B	Learning IGMs	2
III	Approach	3
III-A	Motivation for ELBO as a model selection criterion	4
III-A1	ELBOProxy	4
III-A2	Toy example	5
III-B	Simultaneously solving many inverse problems	6
IV	Experimental Results	7
IV-A	Denoising	7
IV-B	Phase retrieval	8
IV-C	Black hole imaging	9
V	Theory for Linear IGMs	10
VI	Conclusion	11
VII	Experimental details	14
VIII	Inference Details	14
VIII-A	Inferring the IGM	14
IX	Event Horizon Telescope Array	15
IX-A	Array used for the M87* results	15
IX-B	Array used for the Sgr A* results	15
X	Supplemental Theoretical Results	15
X-A	Proof of Theorem 1	15
X-B	Application: Denoising Data From a Subspace	17
X-C	Additional Theoretical Results	18
X-D	Further Discussion of ELBOProxy	20
XI	Full Model Selection Results	20
XII	Additional results	20
XII-A	Denoising	20
XII-B	Phase Retrieval	21
XIII	IGM as a Generator	22

VII. EXPERIMENTAL DETAILS

a) *Model selection*: We parameterize the latent variational distributions $q_\phi(z)$ with normalizing flows

l_ϕ . In particular, suppose our latent variables $z = l_\phi(s)$ where $s \sim \mathcal{N}(0, I)$. Then,

$$\log q_\phi(z) = \log \pi(s) - \log \left| \det \frac{dl_\phi(s)}{ds} \right|$$

where $\pi(\cdot)$ denotes the likelihood of the base distribution of s , which in this case is a Gaussian. For our prior on z given G_θ , we use an isotropic Gaussian prior. Then the ELBOProxy equation becomes

$$\begin{aligned} \text{ELBOProxy}(G_\theta, q_\phi; y) &:= \mathbb{E}_{z \sim l_\phi(z)} [\log p(y|G(z)) \\ &\quad + \log p(z|G) - \log q_\phi(z)] \\ &= \mathbb{E}_{s \sim \mathcal{N}(0, I)} \left[\log p(y|G_\theta(l_\phi(s))) \right. \\ &\quad + \frac{1}{2} \|l_\phi(s)\|_2^2 - \log \pi(s) \\ &\quad \left. + \log \left| \det \frac{dl_\phi(s)}{ds} \right| \right]. \end{aligned}$$

Note that the expectation over $\log \pi(s)$ is constant with respect to the parameters ϕ , so we only optimize the remaining terms.

VIII. INFERENCE DETAILS

a) *Optimization details*: Given a generator G_c for a class c , we optimized the ELBOProxy with respect to the parameters ϕ of the Normalizing Flow. For denoising, we optimized for 5000 epochs with Adam [3] and a learning rate of 10^{-4} and set $T = 200$. For phase retrieval, we optimized for 7500 epochs with Adam and a learning rate of 10^{-4} and set $T = 1000$. The entire experiment took approximately 10/15 hours for denoising/phase retrieval, respectively, on a single NVIDIA V100 GPU. When optimizing the flow h_ϕ for a particular generator G_c , we record the parameters with the best ELBOProxy value. This is the final value recorded in the figure. To generate the mean and standard deviation, we used 100 samples.

b) *Network details*: The generator G_c corresponds to the decoder part of a VAE [4] trained on the training set of class c of MNIST. We used a convolutional architecture, where the encoder and decoder architecture contained 4 convolutional layers. There are also linear, fully-connected layers in between to learn the latent representation, which had dimension 16. The Normalizing Flows used the RealNVP [1] architecture. In both problems, the network had 16 affine coupling layers with linear layers, Leaky ReLU activation functions, and batch normalization. After each affine-coupling block, the dimensions were randomly permuted.

A. Inferring the IGM

a) *Variational distribution*: We parameterize the posterior using Gaussian distributions due to memory

constraints; the memory of the whole setup with Normalizing Flows using 75 images is around 5 times larger than the setup using a Gaussian variational family. However, the Gaussian variational family is likely not as good of a variational distribution, so there will be a wider margin between $p(y|G)$ and the ELBO. Additionally, the Gaussian variational family will not be as good at capturing multi-modal distributions. Note that we do not notice overfitting so there is no need for early stopping.

b) IGM network details.: We use a Deep Decoder [2] generator architecture as the IGM with a dropout of 10^{-4} . The deep decoder has 6 layers of hidden size 150, and has a latent size of 150. There is a sigmoid activation as the final layer of the generator.

c) Optimization details.: We train for 99,000 epochs using the Adam [3] optimizer with a learning rate of 10^{-4} , which takes around 54 hours for 75 images and 4 hours for 4 images on a single NVIDIA A100 GPU. We use a batch size of 20 samples per measurement i , and we perform batch gradient descent based on all measurements $i \in \{1, \dots, N\}$.

IX. EVENT HORIZON TELESCOPE ARRAY

A. Array used for the M87* results

NAME	X	Y	Z
PDB	4523998	468045	4460310
PV	5088968	-301682	3825016
SMT	-1828796	-5054407	3427865
SMA	-5464523	-2493147	2150612
LMT	-768714	-5988542	2063276
ALMA	2225061	-5440057	-2481681
APEX	2225040	-5441198	-2479303
JCMT	-5464585	-2493001	2150654
CARMA	-2397431	-4482019	3843524
KP	-1995679	-5037318	3357328
GLT	1500692	-1191735	6066409

B. Array used for the Sgr A* results

NAME	X	Y	Z
ALMA	2225061	-5440057	-2481681
APEX	2225040	-5441198	-2479303
GLT	1500692	-1191735	6066409
JCMT	-5464585	-2493001	2150654
KP	-1995679	-5037318	3357328
LMT	-768714	-5988542	2063276
SMA	-5464523	-2493147	2150612
SMT	-1828796	-5054407	3427865
BAJA	-2352576	-4940331	3271508
BAR	-2363000	-4445000	3907000
CAT	1569000	-4559000	-4163000
CNI	5311000	-1725000	3075000
GAM	5648000	1619000	-2477000
GARS	1538000	-2462000	-5659000
HAY	1521000	-4417000	4327000
NZ	-4540000	719000	-4409000
OVRO	-2409598	-4478348	3838607
SGO	1832000	-5034000	-3455000
CAS	1373000	-3399000	-5201000
LLA	2325338	-5341460	-2599661
PIKE	-1285000	-4797000	3994000
PV	5088968	-301682	3825016

X. SUPPLEMENTAL THEORETICAL RESULTS

We present proofs of some of the theoretical results that appeared in the main body of the paper, along with new results that were not discussed. Throughout the proofs, we let $\|\cdot\|_2$, $\|\cdot\|$, and $\|\cdot\|_F$ denote the Euclidean norm for vectors, the spectral norm for matrices, and the Frobenius norm for matrices, respectively. For a matrix $M \in \mathbb{R}^{m \times n}$, let $\sigma_i(M)$ and $\lambda_i(M)$ denote the i -th singular value and eigenvalue, respectively (ordered in decreasing order). Let I_n denote the $n \times n$ identity matrix. We may use I for simplicity when the size of the matrix is clear from context. To avoid cumbersome notation, we sometimes use the shorthand notation $\mathbb{E}_z[\cdot]$ for $\mathbb{E}_{z \sim \mathcal{D}}[\cdot]$.

A. Proof of Theorem 1

We first focus on proving Theorem 1. To show this, we first establish that the ELBOProxy loss can be evaluated in closed-form:

Proposition 2. Fix $\sigma > 0$ and suppose $y^{(i)} := Ax^{(i)} + \eta^{(i)}$, $\eta^{(i)} \sim \mathcal{N}(0, \sigma^2 I)$ for $i \in [N]$ where $A \in \mathbb{R}^{m \times n}$. Then the loss $\mathcal{L}_N : \mathbb{R}^{n \times k} \times (\mathbb{R}^k \times \mathbb{R}^{k \times k})^N \rightarrow \mathbb{R}$ in (5)

can be written as

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{2\sigma^2} \left(\|y^{(i)}\|_2^2 + \|AG_\theta U_i\|_F^2 + \|AG_\theta \mu_i\|_2^2 \right) \right. \\ & \quad - \frac{1}{\sigma^2} \langle y^{(i)}, AG_\theta \mu_i \rangle + \frac{1}{2} (\|U_i\|_F^2 + \|\mu_i\|_2^2) \\ & \quad \left. - \log \det(U_i) \right] - \frac{k}{2} - \frac{k}{2} \log(2\pi). \end{aligned}$$

Proof of Proposition 2. Since the loss is separable in N and expectation is linear, we can consider computing the loss for a single $i \in [N]$. Dropping the indices for notational simplicity, consider a fixed $y = Ax + \eta \in \mathbb{R}^m$ where $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$. We aim to calculate

$$\mathbb{E}_v [-\log p(y|G_\theta(v)) - \log p(v|G_\theta) + \log q_\phi(v)]$$

where $p(y|G_\theta(v))$ is Gaussian with variance σ^2 , $p(v|G_\theta)$ is isotropic Gaussian, and $v \sim q_\phi = \mathcal{N}(\mu, UU^T)$. Note that $v \sim \mathcal{N}(\mu, UU^T)$ is equivalent to $v = Uz + \mu$ for $z \sim \mathcal{N}(0, I_k)$. We can calculate the individual terms in the loss as follows. We first collect some useful results, whose proofs are deferred until after this proof is complete:

Lemma 1. *For any matrix $U \in \mathbb{R}^{d \times k}$, we have $\mathbb{E}_z \|Uz\|_2^2 = \|U\|_F^2$ where $z \sim \mathcal{N}(0, I_k)$. Moreover, if UU^T is invertible, then $\mathbb{E}_z [z^T U^T (UU^T)^{-1} Uz] = d$.*

Now, for the first term, since the noise is Gaussian with variance σ^2 , it is equal to

$$\begin{aligned} \mathbb{E}_z \frac{1}{2\sigma^2} \|y - AG_\theta(Uz + \mu)\|_2^2 &= \frac{1}{2\sigma^2} \|y\|_2^2 \\ & \quad + \mathbb{E}_z \frac{1}{2\sigma^2} \|AG_\theta(Uz + \mu)\|_2^2 \\ & \quad - \frac{1}{\sigma^2} \mathbb{E}_z \langle y, AG_\theta(Uz + \mu) \rangle \\ &= \frac{1}{2\sigma^2} (\|y\|_2^2 + \|AG_\theta U\|_F^2 + \|AG_\theta \mu\|_2^2) \\ & \quad - \frac{1}{\sigma^2} \langle y, AG_\theta \mu \rangle \end{aligned}$$

where we used the first part of Lemma 1 in the third equality. For the prior term, we can calculate using Lemma 1 again

$$\begin{aligned} \mathbb{E}_v [-\log p(v|G_\theta)] &= \mathbb{E}_z \frac{1}{2} \|Uz + \mu\|_2^2 \\ &= \frac{1}{2} (\|U\|_F^2 + \|\mu\|_2^2). \end{aligned}$$

Finally, we can calculate the entropy term as

$$\begin{aligned} \mathbb{E}_v [\log q_\phi(v)] &= -\frac{1}{2} \mathbb{E}_v [(v - \mu)^T (UU^T)^{-1} (v - \mu)] \\ & \quad - \frac{k}{2} \log(2\pi) - \frac{1}{2} \log |\det(UU^T)| \\ &= -\frac{1}{2} \mathbb{E}_z [(Uz + \mu - \mu)^T (UU^T)^{-1} (Uz + \mu - \mu)] \\ & \quad - \frac{k}{2} \log(2\pi) - \frac{1}{2} \log |\det(UU^T)| \\ &= -\frac{1}{2} \mathbb{E}_z [z^T U^T (UU^T)^{-1} Uz] \\ & \quad - \frac{k}{2} \log(2\pi) - \frac{1}{2} \log |\det(UU^T)| \\ &= -\frac{k}{2} - \frac{k}{2} \log(2\pi) - \log \det(U) \end{aligned}$$

where we used the second half of Lemma 1 in the last equality and $\log |\det(UU^T)| = 2 \log \det(U)$ in the last equality. Adding the terms achieves the desired result. \square

Proof of Lemma 1. For the first claim, let $U = ESV^T$ denote the singular value decomposition of U with E and V square and unitary. Then by the rotational invariance of the Gaussian distribution and the ℓ_2 -norm, we have

$$\begin{aligned} \mathbb{E}_z \|Uz\|_2^2 &= \mathbb{E}_z \|ESV^T z\|_2^2 = \mathbb{E}_z \|Sz\|_2^2 \\ &= \sum_{i=1}^{\text{rank}(U)} \sigma_i(U)^2 \mathbb{E}_{z_i} [z_i^2] \\ &= \sum_{i=1}^{\text{rank}(U)} \sigma_i(U)^2 = \|U\|_F^2. \end{aligned}$$

For the second claim, recall that $z^T U^T (UU^T)^{-1} Uz = \langle U^T (UU^T)^{-1} U, zz^T \rangle$ where $\langle A, B \rangle = \text{Tr}(A^T B)$ is the Hilbert-Schmidt inner product for matrices. Using linearity of the expectation and the cyclic property of the trace operator, we get

$$\begin{aligned} \mathbb{E}_z [z^T U^T (UU^T)^{-1} Uz] &= \mathbb{E}_z \langle U^T (UU^T)^{-1} U, zz^T \rangle \\ &= \langle U^T (UU^T)^{-1} U, \mathbb{E}_z [zz^T] \rangle \\ &= \langle U^T (UU^T)^{-1} U, I_n \rangle \\ &= \text{Tr}(U^T (UU^T)^{-1} U) \\ &= \text{Tr}(UU^T (UU^T)^{-1}) \\ &= \text{Tr}(I_d) = d. \end{aligned}$$

\square

We now prove Theorem 1. Our proof essentially computes the stationary conditions of the loss \mathcal{L}_N and computes the first-order critical points directly.

Proof of Theorem 1. Using standard matrix calculus, the gradient of the objective \mathcal{L}_N with respect to each

parameter is given by

$$\begin{aligned}\nabla_{G_\theta} \mathcal{L}_N &= \frac{1}{\sigma^2} \left[\frac{1}{N} \sum_{i=1}^N A^T A G_\theta (U_i U_i^T + \mu_i \mu_i^T) \right. \\ &\quad \left. - \frac{1}{N} \sum_{i=1}^N A^T y^{(i)} \mu_i^T \right], \\ \nabla_{U_i} \mathcal{L}_N &= \frac{1}{\sigma^2} G_\theta^T A^T A G_\theta U_i + U_i - (U_i^T)^{-1}, \text{ and} \\ \nabla_{\mu_i} \mathcal{L}_N &= \frac{1}{\sigma^2} (G_\theta^T A^T A G_\theta \mu_i - G_\theta^T A^T y^{(i)}) + \mu_i.\end{aligned}$$

For a fixed G_θ , we can directly compute which μ_i and U_i satisfy the stationary equations $\nabla_{\mu_i} \mathcal{L}_N = 0$ and $\nabla_{U_i} \mathcal{L}_N = 0$:

$$\begin{aligned}\mu_i &= (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T y^{(i)} \text{ and} \\ U_i &= \sigma (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1/2} \tilde{R}\end{aligned}$$

where \tilde{R} is an arbitrary orthogonal matrix.

We now calculate, given such μ_i and U_i , which G_θ satisfies the condition $\nabla_{G_\theta} \mathcal{L}_N = 0$. For notational simplicity, let $\Sigma_\sigma := G_\theta^T A^T A G_\theta + \sigma^2 I$. Plugging in μ_i and U_i into the stationary condition $\nabla_{G_\theta} \mathcal{L}_N = 0$, this equation reads

$$\begin{aligned}A^T A G_\theta (\sigma^2 I_k + \Sigma_\sigma^{-1} G_\theta^T A^T Y_N A G_\theta) \Sigma_\sigma^{-1} \\ = A^T Y_N A G_\theta \Sigma_\sigma^{-1}.\end{aligned}$$

Since A has rank m , AA^T is invertible, so we can multiply both sides of the equation on the left by $(AA^T)^{-1}A$ and the right by Σ_σ to obtain

$$(I_m - A G_\theta \Sigma_\sigma^{-1} G_\theta^T A^T) Y_N A G_\theta = \sigma^2 A G_\theta. \quad (6)$$

Now, consider the SVD of $A G_\theta = E S V^T$ where $E \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{k \times k}$ are orthogonal and $S \in \mathbb{R}^{m \times k}$ is rectangular with non-zero diagonal submatrix containing the singular values of $A G_\theta$ denoted by $\tilde{S} \in \mathbb{R}^{k \times k}$. First, observe that

$$\begin{aligned}A G_\theta \Sigma_\sigma^{-1} G_\theta^T A^T &= E S V^T (V S^T S V^T + \sigma^2 I_k)^{-1} V S^T E^T \\ &= E S V^T (V \tilde{S}^2 V^T + \sigma^2 I_k)^{-1} V S^T E^T \\ &= E S V^T V (\tilde{S}^2 + \sigma^2 I_k)^{-1} V^T V S^T E^T \\ &= E S (\tilde{S}^2 + \sigma^2 I_k)^{-1} S^T E^T \\ &= E \begin{bmatrix} \tilde{S}^2 (\tilde{S}^2 + \sigma^2 I_k)^{-1} & 0 \\ 0 & 0 \end{bmatrix} E^T.\end{aligned}$$

This implies that $I_m - A G_\theta \Sigma_\sigma^{-1} G_\theta^T A^T = E \hat{S}_\sigma E^T$ where

$$\hat{S}_\sigma := \begin{bmatrix} I_k - \tilde{S}^2 (\tilde{S}^2 + \sigma^2 I_k)^{-1} & 0 \\ 0 & I_{m-k} \end{bmatrix}.$$

Note that for $\sigma > 0$, \hat{S}_σ is invertible. Hence, our stationary equation (6) reads

$$E \hat{S}_\sigma E^T Y_N E S V^T = \sigma^2 E S V^T.$$

Using the orthogonality of E and V and invertibility of

\hat{S}_σ , we can further simplify the above equation to

$$Y_N E S = \sigma^2 E \hat{S}_\sigma^{-1} S.$$

Now, suppose $S_{ii} = \sigma_i (A G_\theta) = \sigma_i \neq 0$. Then we have that the i -th column of E , denoted e_i , satisfies

$$Y_N e_i = \sigma^2 \left(1 - \frac{\sigma_i^2}{\sigma_i^2 + \sigma^2} \right)^{-1} e_i.$$

That is, e_i is an eigenvector of Y_N with eigenvalue $\lambda_i = \sigma^2 (1 - \sigma_i^2 / (\sigma_i^2 + \sigma^2))^{-1}$. Observe that this equation is satisfied if and only if $\sigma_i = \pm \sqrt{\lambda_i - \sigma^2}$. Since singular values are non-negative, we must have $\sigma_i = \sqrt{\lambda_i - \sigma^2}$. Note that if $S_{ii} = 0$, then e_i can be arbitrary. Thus, we obtain that in order for G_θ to satisfy the stationary condition, it must satisfy

$$A G_\theta = E_k (L_k - \sigma^2 I)^{1/2} R$$

where E_k contains the top- $\min\{\text{rank}(Y_N), k\}$ eigenvectors of Y_N , L_k is a diagonal matrix whose i -th entry is λ_i with corresponding eigenvector e_i and σ^2 otherwise, and R is an arbitrary orthogonal matrix. \square

B. Application: Denoising Data From a Subspace

In order to gain further intuition about Theorem 1, it will be useful to consider an example where the data explicitly lies in a low-dimensional subspace. We will then compare our inferred estimator with an estimator that has access to complete prior knowledge about the data.

Consider the denoising problem $A = I$ and suppose that each ground-truth image is drawn at random from a k -dimensional subspace as follows: $x^{(i)} = G z^{(i)}$ where $G \in \mathbb{R}^{n \times k}$ and $z^{(i)} \sim \mathcal{N}(0, I)$. If we had access to the ground-truth generating matrix G and knew the underlying distribution in the latent space $\mathcal{N}(0, I)$, then one could denoise the i -th image $y^{(i)}$ by returning the following MAP estimate $\hat{x}^{(i)}$ in the range of G :

$$\hat{x}^{(i)} := G \hat{z}^{(i)},$$

$$\begin{aligned}\text{where } \hat{z}^{(i)} &:= \arg \min_{z \in \mathbb{R}^k} \frac{1}{2\sigma^2} \|y^{(i)} - Gz\|_2^2 + \frac{1}{2} \|z\|_2^2 \\ &= (G^T G + \sigma^2 I)^{-1} G^T y^{(i)}.\end{aligned}$$

Note that this is the mean of the posterior $p(x|y^{(i)}, G)$. Now, to compare solutions, observe that our (mean) estimator given by Theorem 1 is

$$G_* \mu_i^* = G_* (G_*^T G_* + \sigma^2 I)^{-1} G_*^T y^{(i)}.$$

Recall that G_* contains the top eigenvectors and eigenvalues of Y_N , the sample covariance of the measurements. In the case that our data $x^{(i)} = G z^{(i)}$, note that

for sufficiently many examples N , the sample covariance Y_N is close to its mean $\mathbb{E}_{z,\eta} Y_N$. That is, for large N ,

$$Y_N \approx \mathbb{E}_{z,\eta} Y_N = GG^T + \sigma^2 I.$$

To see this more concretely, let $G = ELV^T$ denote the SVD of G . Then $GG^T + \sigma^2 I = E\hat{L}_\sigma E^T$ where \hat{L}_σ has $\sigma_i(G)^2 + \sigma^2$ in the first k diagonal entries and σ^2 in the remaining $n - k$ entries. So in the limit of increasing data, G_* estimates the top- k eigenvectors of $\mathbb{E}_{z,\eta} Y_N$, which are precisely the top- k principal components of G . Moreover, the singular values of G_* are $\sqrt{\lambda_i - \sigma^2}$ where $\lambda_i = \lambda_i(Y_N)$. As data increases, $\lambda_i(Y_N) \approx \lambda_i(\mathbb{E}_{z,\eta} Y_N) = \sigma_i(G)^2 + \sigma^2$ so that $\sigma_i(G_*) = \sqrt{\lambda_i - \sigma^2} \approx \sigma_i(G)$. Thus the singular values of G_* also approach those of G .

A natural question to consider is how large the number of examples N would need to be in order for such an approximation to hold. Ideally, this approximation would hold as long as the number of examples N scales like the intrinsic dimensionality of the underlying images. For example, the following result shows that in the noiseless case, this approximation holds as long as the number of examples N scales like k , the dimension of the subspace. The proof follows from standard matrix concentration bounds, e.g., Theorem 4.6.1 in [5].

Proposition 3. Fix $G \in \mathbb{R}^{n \times k}$ and let $y^{(i)} = Gz^{(i)}$ where $z^{(i)} \sim \mathcal{N}(0, I)$ for $i \in [N]$. Then, for any $\varepsilon > 0$, there exists positive absolute constants C_1 and C_2 such that the following holds: if $N \geq C_1 \varepsilon^{-2} k$, then with probability $1 - 2e^{-k}$ the sample covariance $Y_N := \frac{1}{N} \sum_{i=1}^N y^{(i)} (y^{(i)})^T$ obeys

$$\|Y_N - GG^T\| \leq C_2 \|G\|^2 \varepsilon.$$

Hence, in the case of data lying explicitly in a low-dimensional subspace, our estimator returns an approximation of the solution found via MAP estimation, which can only be computed with complete prior knowledge of the underlying image structure. In contrast, our estimator infers this subspace from measurement information only and its performance improves as the number of examples increases. The approximation quality of the estimator is intimately related to the number of examples and the intrinsic dimensionality of the images, which we show can be quantified in simple settings. While our main result focused on IGMs inferred with linear models, it is an interesting future direction to theoretically understand properties of non-linear IGMs given by deep networks and what can be said for more complicated forward models. We leave this for future work.

C. Additional Theoretical Results

We can further analyze properties of the model found by minimizing \mathcal{L}_N from Theorem 1. In particular, we

can characterize specific properties of the inferred mean μ_i^* as a function of the amount of noise in our measurements and a fixed IGM $G_\theta \in \mathbb{R}^{n \times k}$. Specifically, given measurements y_i and an IGM G_θ , recall that the inferred mean computed in Theorem 1 is given by

$$\mu_i^*(\sigma) := \sigma^{-2} (\sigma^{-2} \cdot G_\theta^T A^T A G_\theta + I)^{-1} G_\theta^T A^T y_i$$

We will give an explicit form for $\mu_i^*(\sigma)$ as the amount of noise goes to 0 and show that, in the case our IGM is correct ($G_\theta = G$ for some ground-truth G), the mean recovers the ideal latent code with vanishing noise.

Proposition 4. Suppose $A \in \mathbb{R}^{m \times n}$ and $G \in \mathbb{R}^{n \times k}$ with $k \leq m \leq n$. Let $y_i := AGz_i + \eta_i$ where $z_i \in \mathbb{R}^k$ and $\eta_i \sim \mathcal{N}(0, \sigma^2 I)$ for $i = 1, \dots, N$. Then, with probability 1, for any $G_\theta \in \mathbb{R}^{n \times k}$, the estimator $\mu_i^*(\sigma)$ for z_i defined above satisfies

$$\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) = (AG_\theta)^\dagger AGz_i \quad \forall i \in [N].$$

In the special case that $G_\theta = G$ and $\text{rank}(AG) = k$, then

$$\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) = z_i \quad \forall i \in [N].$$

Moreover, if $\text{rank}(AG_\theta) = \text{rank}(AG) = k$, then we get that the following error bound holds for all $i \in [N]$ with probability at least $1 - Ne^{-\gamma m}$ for some positive absolute constant γ :

$$\begin{aligned} \|\mu_i^*(\sigma) - z_i\|_2 &\leq \|AGz_i\|_2 \sigma_k(AG_\theta)^{-3} \cdot \sigma^2 \\ &\quad + 2\sigma_k(AG_\theta)^{-1} \sqrt{m} \cdot \sigma \\ &\quad + \sqrt{2} \|(AG_\theta)^\dagger\| \|(AG)^\dagger\| \cdot \|A(G_\theta - G)\|. \end{aligned}$$

Proof. For notational simplicity, set $\Sigma_{\theta,\sigma} := \sigma^{-2} \cdot G_\theta^T A^T A G_\theta + I$. Then

$$\begin{aligned} \mu_i^*(\sigma) &= \sigma^{-2} \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T y_i \\ &= \sigma^{-2} \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T (AGz_i + \eta_i) \\ &= \sigma^{-2} \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T AGz_i + \sigma^{-2} \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T \eta_i \\ &= \sigma^{-2} \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T AGz_i + \sigma^{-1} \cdot \Sigma_{\theta,\sigma}^{-1} G_\theta^T A^T \tilde{\eta}_i \end{aligned}$$

where we note that any $\eta_i \sim \mathcal{N}(0, \sigma^2 I)$ can be written as $\eta_i = \sigma \cdot \tilde{\eta}_i$ where $\tilde{\eta}_i \sim \mathcal{N}(0, I)$. Observe that

$$\begin{aligned} \Sigma_{\theta,\sigma}^{-1} &= (\sigma^{-2} G_\theta^T A^T A G_\theta + \sigma^{-2} \cdot \sigma^2 I)^{-1} \\ &= \sigma^2 (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} \end{aligned}$$

since $(cA)^{-1} = c^{-1}A^{-1}$ for $c \neq 0$. This simplifies $\mu_i^*(\sigma)$ to

$$\begin{aligned} \mu_i^*(\sigma) &= (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T AGz_i \\ &\quad + \sigma \cdot (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i. \end{aligned}$$

To compute the limit of the expression above, we recall that the pseudoinverse of a matrix M is a limit of the

form

$$\begin{aligned} M^\dagger &:= \lim_{\delta \rightarrow 0^+} (M^T M + \delta I)^{-1} M^T \\ &= \lim_{\delta \rightarrow 0^+} M^T (M M^T + \delta I)^{-1}. \end{aligned}$$

Note that this quantity is defined even if the inverses of $M M^T$ or $M^T M$ exist. Applying this to $M = A G_\theta$, we observe that

$$\lim_{\sigma \rightarrow 0^+} (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T A G z_i = (A G_\theta)^\dagger A G z_i$$

and

$$\begin{aligned} &\lim_{\sigma \rightarrow 0^+} \sigma (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i \\ &= \left(\lim_{\sigma \rightarrow 0^+} \sigma \right) \left(\lim_{\sigma \rightarrow 0^+} (G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i \right) \\ &= 0 \cdot (A G_\theta)^\dagger = 0. \end{aligned}$$

Combining the above two limits gives $\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) = (A G_\theta)^\dagger A G z_i$ as claimed.

For the special case, since $A G \in \mathbb{R}^{m \times k}$ has $\text{rank}(A G) = k$, the pseudoinverse of $A G$ is given by $(A G)^\dagger = (G^T A^T A G)^{-1} G^T A^T$. Applying this to $M = A G$, we observe that

$$\begin{aligned} &\lim_{\sigma \rightarrow 0^+} (G^T A^T A G + \sigma^2 I)^{-1} G^T A^T A G z_i \\ &= (A G)^\dagger A G z_i = z_i. \end{aligned}$$

Now, for the error bound, we note that for $\sigma \geq 0$, we have by the triangle inequality,

$$\begin{aligned} \|\mu_i^*(\sigma) - z_i\|_2 &\leq \|\mu_i^*(\sigma) - \lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma)\|_2 \\ &\quad + \|\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) - z_i\|_2 \end{aligned}$$

so we bound each term separately. Recall that we have shown $\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) = (A G_\theta)^\dagger A G z_i$. For the first term, observe that

$$\begin{aligned} &\|\mu_i^*(\sigma) - (A G_\theta)^\dagger A G z_i\|_2 \\ &\leq \|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T A G z_i - (A G_\theta)^\dagger A G z_i\|_2 \\ &\quad + \sigma \cdot \|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i\|_2 \\ &\leq \|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T - (A G_\theta)^\dagger\| \|A G z_i\|_2 \\ &\quad + \sigma \cdot \|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i\|_2 \end{aligned}$$

Let $A G_\theta = U S V^T$ denote the singular value decomposition of $A G_\theta$. Then $G_\theta^T A^T A G_\theta = V S U^T U S V^T = V S^2 V^T$. This implies that

$$\begin{aligned} &(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \\ &= (V S^2 V^T + \sigma^2 I)^{-1} V S U^T. \end{aligned}$$

Moreover, since V is orthogonal, we have that

$$\begin{aligned} (V S^2 V^T + \sigma^2 I)^{-1} &= (V S^2 V^T + \sigma^2 V V^T)^{-1} \\ &= V (S^2 + \sigma^2 I)^{-1} V^T \end{aligned}$$

Thus, for $\sigma, \alpha \geq 0$, we have that

$$\begin{aligned} &[(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} - (G_\theta^T A^T A G_\theta + \alpha^2 I)^{-1}] G_\theta^T A^T \\ &= [(V S^2 V^T + \sigma^2 I)^{-1} - (V S^2 V^T + \alpha^2 I)^{-1}] V S U^T \\ &= V [(S^2 + \sigma^2 I)^{-1} - (S^2 + \alpha^2 I)^{-1}] V^T V S U^T \\ &= V [(S^2 + \sigma^2 I)^{-1} - (S^2 + \alpha^2 I)^{-1}] S U^T. \end{aligned}$$

Thus, we conclude by rotational invariance of the spectral norm, we obtain

$$\begin{aligned} &\|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \\ &\quad - (G_\theta^T A^T A G_\theta + \alpha^2 I)^{-1} G_\theta^T A^T\| \\ &= \|[(S^2 + \sigma^2 I)^{-1} - (S^2 + \alpha^2 I)^{-1}] S\|. \end{aligned}$$

Recall that S is a diagonal matrix containing the singular values of $A G_\theta$. Let σ_i denote the i -th singular value of $A G_\theta$ for $i = 1, \dots, k$. Then by direct calculation, we have that

$$\begin{aligned} &\|[(S^2 + \sigma^2 I)^{-1} - (S^2 + \alpha^2 I)^{-1}] S\| \\ &= \max_{i \in [k]} \left| \sigma_i \cdot \left(\frac{1}{\sigma_i^2 + \sigma^2} - \frac{1}{\sigma_i^2 + \alpha^2} \right) \right| \\ &= \max_{i \in [k]} \left| \frac{\sigma_i (\sigma^2 - \alpha^2)}{(\sigma_i^2 + \sigma^2)(\sigma_i^2 + \alpha^2)} \right|. \end{aligned}$$

Taking the limit as $\alpha \rightarrow 0^+$, we obtain

$$\begin{aligned} &\|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T - (A G_\theta)^\dagger\| \\ &= \max_{i \in [k]} \frac{\sigma^2}{\sigma_i (\sigma_i^2 + \sigma^2)} \\ &\leq \sigma_k^{-3} \cdot \sigma^2. \end{aligned}$$

This additionally shows that for any $\sigma \geq 0$,

$$\begin{aligned} &\|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T\| \\ &= \|(S^2 + \sigma^2 I)^{-1} S\| \\ &\leq \sigma_k^{-1}. \end{aligned}$$

Let E_N be the event that $\max_{i \in [N]} \|\tilde{\eta}_i\|_2 \leq 2\sqrt{m}$. By a single-tailed variant of Corollary 5.17 in [5], we have $\mathbb{P}(E_N) \geq 1 - N e^{-\gamma m}$ for some absolute constant $\gamma > 0$. Thus, with the same probability, we have that the following bound holds:

$$\sigma \cdot \|(G_\theta^T A^T A G_\theta + \sigma^2 I)^{-1} G_\theta^T A^T \tilde{\eta}_i\|_2 \leq 2\sigma \sigma_k^{-1} \sqrt{m}.$$

For the term $\|\lim_{\sigma \rightarrow 0^+} \mu_i^*(\sigma) - z_i\|_2$, note that we have

$$\begin{aligned} \|(A G_\theta)^\dagger A G z_i - z_i\|_2 &= \|(A G_\theta)^\dagger A G z_i - (A G)^\dagger A G z_i\|_2 \\ &\leq \|(A G_\theta)^\dagger - (A G)^\dagger\| \cdot \|A G z_i\|_2 \end{aligned}$$

where $\|\cdot\|$ is the spectral norm. Assuming $\text{rank}(A G_\theta) = \text{rank}(A G) = k$, we can apply Theorem 4.1 in [6] to bound

$$\begin{aligned} &\|(A G_\theta)^\dagger - (A G)^\dagger\| \\ &\leq \sqrt{2} \|(A G_\theta)^\dagger\| \|(A G)^\dagger\| \|A(G_\theta - G)\|. \end{aligned}$$

Combining this bound with the previous equation yields the desired inequality. \square

D. Further Discussion of ELBOProxy

In Section III-A, we showed that the ELBO is equal to the ELBOProxy for injective or invertible generators G with approximate posteriors h_ϕ induced by G . While this equality may not hold for non-injective networks G , we expect this proxy to be a reasonable estimate for the ELBO via the following heuristic argument. Consider the distribution $G\#p(z|G)$ convolved with isotropic Gaussian noise: $p_\sigma(x|G) = \int p_\sigma(x|z, G)p(z|G)dz$ where $p_\sigma(x|z, G) = \mathcal{N}(G(z), \sigma^2 I)$ for $\sigma \ll 1$. If $x \sim h_\phi = G\#q_\phi$, then $x = G(z_0)$ for some $z_0 \sim q_\phi$. For any small $\delta > 0$, observe that the likelihood for such an x can be decomposed into two terms:

$$\begin{aligned} p_\sigma(x|G) &= \int p_\sigma(G(z_0)|z, G)p(z|G)dz \\ &= \int_{\|z-z_0\|_2 \leq \delta} p_\sigma(G(z_0)|z, G)p(z|G)dz \\ &\quad + \int_{\|z-z_0\|_2 > \delta} p_\sigma(G(z_0)|z, G)p(z|G)dz. \end{aligned}$$

For smooth generators G and δ sufficiently small, we would expect $G(z) \approx G(z_0)$ for $z \approx z_0$ so that the density $p_\sigma(G(z_0)|z, G)p(z|G)$ is constant in a neighborhood of z_0 , making the first term $\approx C_\sigma p(z_0|G)$ for some constant C_σ . Moreover, if $G(z_0)$ is of sufficiently high likelihood, we would expect the second term to be negligible for $\sigma \ll 1$ so that $\mathbb{E}_{x \sim h_\phi}[\log p_\sigma(x|G)] \approx \mathbb{E}_{z \sim q_\phi}[\log p(z|G)]$ (up to a constant).

XI. FULL MODEL SELECTION RESULTS

In this section, we expand upon the model selection results using the ELBOProxy in Section III-A in the main text. The experimental setup is the same, where we aim to select a generative model for a class of MNIST in two different inverse problems: denoising and phase retrieval. The full array of ELBOProxy values for denoising is given in Fig. 11 and for phase retrieval in Fig. 12.

XII. ADDITIONAL RESULTS

In this section, we expand upon the results for inferring the IGM using the methods described in Section III-B of the main text.

A. Denoising

In Fig. 4 of the main body, we showed results for denoising 95 images of celebrity A. We include additional visualizations of the posterior by showing samples from multiple posteriors in Fig. 13, which are generated from

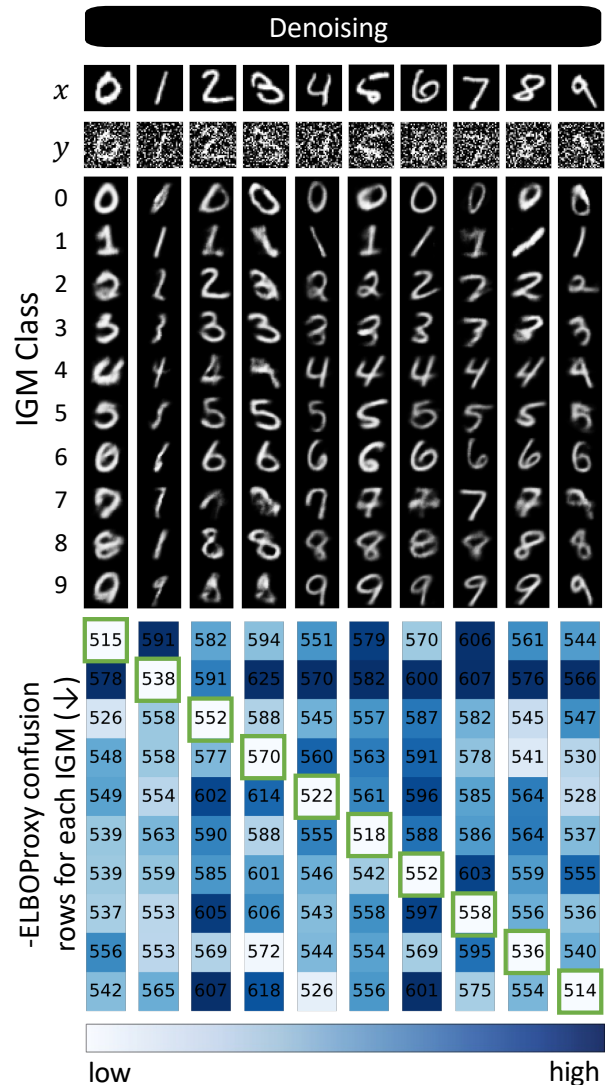


Fig. 11: **Full model selection results for denoising.** We show case further results on model selection with generative networks for denoising. Left: the two leftmost columns correspond to the ground-truth image x_c and the noisy measurements y_c . Center: we show the means of the variational distributions given by the IGM trained on a particular class. Interesting examples in which the ELBOProxy values are close to one another are highlighted in red. Right: each row of the array corresponds to the $-\text{ELBOProxy}$ achieved by each model in reconstructing the images. Here, lower is better. We highlight entries with green boxes with the best $-\text{ELBOProxy}$ values in each row.

the IGM trained as shown in Fig. 4 in the main body. We show additional results on denoising 95 images of celebrity face B in Fig. 14 along with samples from the posterior in Fig. 15. In both cases, our reconstructions are much less noisy than the measurements and sharpen the primary facial features.

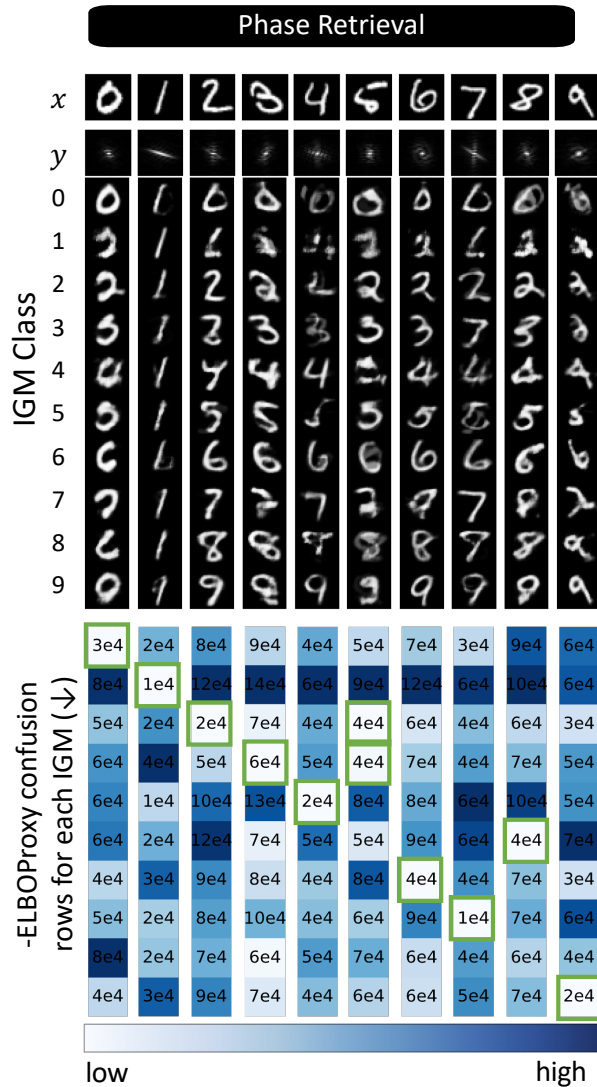


Fig. 12: **Full model selection results for phase retrieval.** We showcase further results on model selection with generative networks for phase retrieval. Left: the two leftmost columns correspond to the ground-truth image x_c and the noisy measurements y_c . Center: we show the means of the variational distributions given by the IGM trained on a particular class. Interesting examples in which the ELBOProxy values are close to one another are highlighted in red. Right: each row of the array corresponds to the $-\text{ELBOProxy}$ achieved by each model in reconstructing the images. Here, lower is better. We highlight entries with green boxes with the best $-\text{ELBOProxy}$ values in each row.

B. Phase Retrieval

We show additional results from phase retrieval with 150 images of MNIST 8's in Fig. 16 and Fig. 17. The examples shown in Fig. 16 are from the Fourier phase retrieval measurements whereas the ones in Fig. 17 are from Gauss phase retrieval measurements. These are in addition to those shown in Fig. 6 in the main text. Note that the simple, uni-modal Gaussian variational

distribution is not expressive enough to capture the multi-modal structure of the true posterior in the Fourier phase retrieval problem.

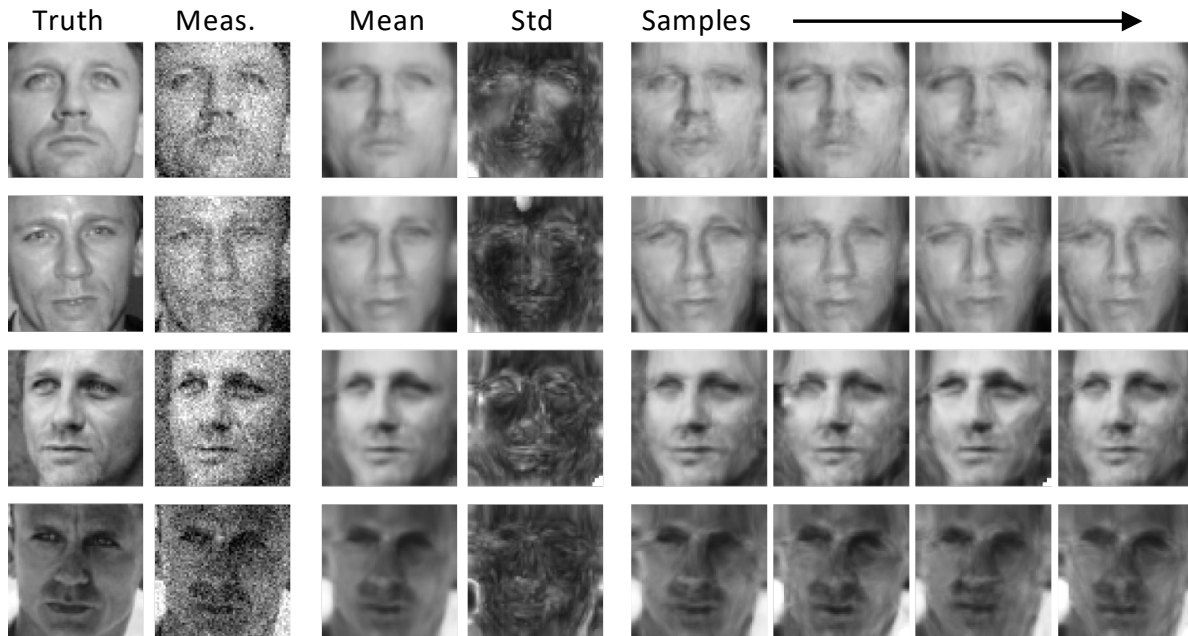


Fig. 13: **Posterior samples from denoising 95 images of celebrity A.** We show additional results for denoising 95 images of celebrity A, which corresponds to the same results shown in Fig. 4 of the main body. Here we show the target image, the noisy measurements, mean of the posterior, standard deviation of the posterior, and samples from the posterior.



Fig. 14: **Denoising 95 images of celebrity B.** We demonstrate our method described in Section III-B of the main body using 95 noisy images of a celebrity. These results are of celebrity B, but correspond to Fig. 4 in the main body. Here we show the ground truth (top), noisy measurements (middle), and mean reconstruction (bottom) for a subset of the 95 different noisy images. Our reconstructions are much less noisy, exhibiting sharper features that are hard to discern in the noisy images. Note that no predefined prior/regularizer was used in denoising.

XIII. IGM AS A GENERATOR

While the goal of our method is to improve performance in solving the underlying inverse problem, we can inspect what the inferred IGM has learned by generating samples. We show samples from different inferred IGMs in Fig. 18 for a variety of datasets and inverse problems. From left to right: IGM from denoising MNIST 8's (Fig. 3 from the main body), IGM inferred from phase retrieval measurements of MNIST 8's, IGM from denoising celebrity face A (Fig. 4 from the main body), IGM from denoising celebrity face B (Fig. 14), and IGM from video reconstruction of a black hole from

the black hole compressed sensing problem (Fig. 8 from the main body).

REFERENCES

- [1] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. 2017.
- [2] Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. *International Conference on Learning Representations (ICLR)*, 2019.
- [3] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua




Fig. 15: **Posterior samples from denoising 95 images of a celebrity.** We demonstrate our method described in Section III-B of the main body to denoise 95 images. Here we show the target image, the noisy measurements, mean of the posterior, standard deviation of the posterior, and samples from the posterior.

Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

- [4] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations (ICLR)*, 2014.
- [5] Roman Vershynin. High-dimensional probability: An introduction with applications in data science. *Cambridge University Press*, 2018.
- [6] Per-Ake Wedin. Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics volume*, 13:217–232, 1973.




Fig. 16: **Random examples from Fourier phase retrieval on MNIST 8's.** We demonstrate our method described in Section III-B of the main body to perform phase retrieval on 75 images. Here we show the target image, the log magnitude of the Fourier transform, mean of the posterior, standard deviation of the posterior, and samples from the posterior. Note that the simple Gaussian posterior distribution is unable to capture the full multi-modal structure of the true posterior.




Fig. 17: **Random examples from Gauss phase retrieval on MNIST 8's.** We demonstrate our method described in Section III-B of the main body to perform phase retrieval on 150 images. Here we show the target image, mean of the posterior, standard deviation of the posterior, and samples from the posterior.

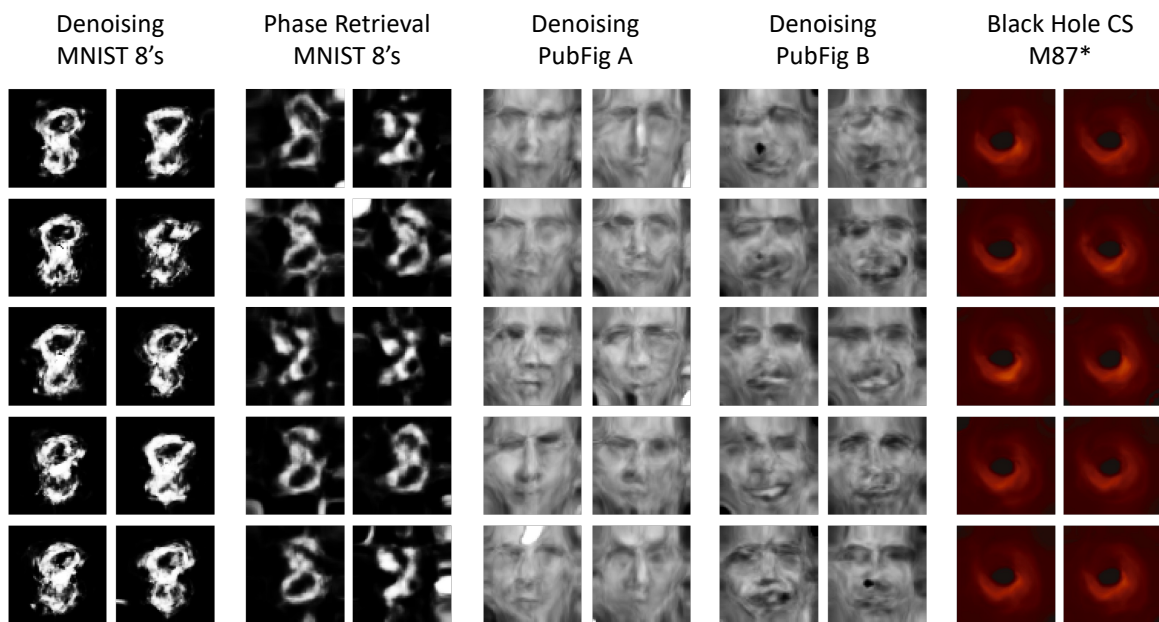


Fig. 18: **Samples from IGM.** Here we show samples of the IGM defined by $w = G(z)$ where $z \sim \mathcal{N}(0, \sigma^2 I)$. We use a $\sigma = 1$ for all cases except for phase retrieval, which uses a $\sigma = 5$. We show samples from a variety of IGMs using different datasets and inverse problems. From top to bottom: denoising MNIST 8's, measurements of MNIST 8's from the black hole compressed sensing problem, denoising celebrity face A, denoising celebrity face B, video reconstruction of a black hole from the black hole compressed sensing problem. The range of these images are all 0 to 1.