
Invertible generative models for inverse problems: mitigating representation error and dataset bias

Muhammad Asim^{*1} Mara Daniels^{*2} Oscar Leong³ Ali Ahmed^{†1} Paul Hand^{†2}

Abstract

Trained generative models have shown remarkable performance as priors for inverse problems in imaging – for example, Generative Adversarial Network priors permit recovery of test images from 5-10x fewer measurements than sparsity priors. Unfortunately, these models may be unable to represent any particular image because of architectural choices, mode collapse, and bias in the training dataset. In this paper, we demonstrate that invertible neural networks, which have zero representation error by design, can be effective natural signal priors at inverse problems such as denoising, compressive sensing, and inpainting. Given a trained generative model, we study the empirical risk formulation of the desired inverse problem under a regularization that promotes high likelihood images, either directly by penalization or algorithmically by initialization. For compressive sensing, invertible priors can yield higher accuracy than sparsity priors across almost all undersampling ratios, and due to their lack of representation error, invertible priors can yield better reconstructions than GAN priors for images that have rare features of variation within the biased training set, including out-of-distribution natural images. We additionally compare performance for compressive sensing to unlearned methods, such as the deep decoder, and we establish theoretical bounds on expected recovery error in the case of a linear invertible model.

Equal contributions are denoted by ^{*} and [†]. ¹Department of Electrical Engineering, Information Technology University, Lahore, Pakistan ²Department of Mathematics and Khoury College of Computer Sciences, Northeastern University, Boston, MA ³Department of Computational and Applied Mathematics, Rice University, Houston, TX. Correspondence to: Mara Daniels <daniels.g@northeastern.edu>.

1. Introduction



Figure 1. We train an invertible generative model with CelebA images (including those shown). When used as a prior for compressive sensing, it can yield higher quality image reconstructions than Lasso and a trained DCGAN, even on out-of-distribution images. Note that the DCGAN reflects biases of the training set by removing the man’s glasses and beard, whereas our invertible prior does not.

Generative deep neural networks have shown remarkable performance as natural signal priors in imaging inverse problems, such as denoising, inpainting, compressive sensing, blind deconvolution, and phase retrieval. These generative models can be trained from datasets consisting of images of particular natural signal classes, such as faces, fingerprints, MRIs, and more (Karras et al., 2018; Minaee & Abdolrashidi, 2018; Shin et al., 2018; Chen et al., 2018). Some such models, including variational autoencoders (VAEs) and generative adversarial networks (GANs), learn an explicit low-dimensional manifold that approximates a natural signal class (Goodfellow et al., 2014; Kingma & Welling, 2014; Rezende et al., 2014). We will refer to such models as GAN priors. These priors can be used for inverse problems by attempting to find the signal in the range of

the generative model that is most consistent with provided measurements. When the GAN has a low dimensional latent space, this allows for a low dimensional optimization problem that operates directly on the natural signal class. Consequently, generative priors can obtain significant performance improvements over classical methods. For example, GAN priors have been shown to outperform sparsity priors at compressive sensing with 5-10x fewer measurements in some cases. Additionally, GAN priors have led to novel theory for signal recovery in the linear compressive sensing and nonlinear phase retrieval problems (Bora et al., 2017; Hand & Voroninski, 2018; Hand et al., 2018), and they have also shown promising results for the nonlinear blind image deblurring problem (Asim et al., 2019).

A significant drawback of GAN priors for solving inverse problems is that they can have large representation error or bias due to architecture and training. That is, a desired image may not be in or near the range of a particular trained GAN. Representation error can occur both for in-distribution and out-of-distribution images. For in-distribution, it can be caused by inappropriate latent dimensionality and mode collapse. For out-of-distribution images, representation error can be large in part because the GAN training process explicitly discourages such images due to the presence of the concurrently trained discriminator network. For many imaging inverse problems, it is important to be able to recover signals that are out-of-distribution relative to training data. For example, in scientific and medical imaging, novel objects or pathologies may be expressly sought. Additionally, desired signals may be out-of-distribution because a training dataset has bias and is unrepresentative of the true underlying distribution. As an example, the CelebA dataset (Liu et al., 2015) is biased toward people who are young, who do not have facial hair or glasses, and who have a light skin tone. As we will see, a GAN prior trained on this dataset learns these biases and exhibits image recovery failures because of them.

Several recent priors have been developed that have lower representation error than GANs. One class of approaches are unlearned neural network priors, such as the Deep Image Prior and the Deep Decoder (Ulyanov et al., 2018; Heckel & Hand, 2019). These are neural networks that are randomly initialized, and whose weights are optimized at inversion time to best fit provided measurements. They have practically zero representation error for natural images. Because they are untrained, there is no training set or training distribution, and hence there is no notion of in- or out-of-distribution natural images. Another class of approaches include updating the weights of the trained GAN at inversion time in an image adaptive way, such as the IAGAN (Hussein et al., 2020). Such an approach could be interpreted as using the GAN as a warm start for a Deep Image Prior. A further approach is Latent Convolutional Models (Athar et al., 2019),

in which a generative prior is trained using high dimensional latent representations which are structured as the parameters of a randomly initialized convolutional neural network.

In this paper, we study flow-based invertible neural networks as signal priors. These networks are mathematically invertible (one-to-one and onto) by architectural design (Dinh et al., 2017; Gomez et al., 2017; Jacobsen et al., 2018; Kingma & Dhariwal, 2018). Consequently, they have zero representation error and are capable of recovering any image, including those significantly out-of-distribution relative to a training set; see Figure 1. We call the domain of an invertible generator the latent space, and we call the range of the generator the signal space. These must have equal dimensionality. The strengths of these invertible models include: their architecture allows exact and efficient latent-variable inference, direct log-likelihood evaluation, and efficient image synthesis; they have the potential for significant memory savings in gradient computations; and they can be trained by directly optimizing the likelihood of training images. This paper emphasizes an additional strength: *because they lack representation error, invertible models can mitigate dataset bias and improve recovery performance on inverse problems, including for signals that are out-of-distribution relative to training data.*

We present a method for using pretrained generative invertible neural networks as priors for imaging inverse problems. An invertible generator, once trained, can be used for a wide variety of inverse problems, with no specific knowledge of those problems used during the training process. As an invertible net permits a likelihood estimate for all images, image recovery can be posed as seeking the highest likelihood image that is consistent with provided measurements.

As a proxy for the image log-likelihood, we pose an optimization of squared data-fit over the latent space under regularization by likelihood of latent representations. In the case of denoising, we explicitly penalize log-likelihood of latent codes, while in compressive sensing and inpainting, regularization is achieved algorithmically. This is due in part to initialization with latent code zero.

The contributions of this paper are as follows. We train a generative invertible model using the CelebA dataset. With this fixed model as a signal prior, we study its performance at multiple inverse problems.

- For image denoising, invertible neural network priors can yield sharper images with higher PSNRs than BM3D (Dabov et al., 2007).
- For compressive sensing of in-distribution images, invertible neural network priors can yield higher PSNRs than GANs with low-dimensional latent dimensionality (both DCGAN and PG-GAN), Image Adaptive

GANs, sparsity priors, and a Deep Decoder across a wide range of undersampling ratios.

- Invertible neural networks exhibit graceful performance decay for compressive sensing on out-of-distribution images. They can yield significantly higher PSNRs than GANs with low latent dimensionality across a wide range of undersampling ratios. They can additionally yield higher or comparable PSNRs than the Deep Decoder when there are sufficiently many measurements.
- We introduce a likelihood-based theoretical analysis of compressive sensing under invertible generative priors in the case that the generator is linear. Given m linear measurements of an n -dimensional signal, we prove a theorem establishing upper and lower bounds of expected squared recovery error in terms of the sum of the squares of the smallest $n - m$ singular values of the model.

2. Method

We assume that we have access to a pretrained generative Invertible Neural Network (INN), $G : \mathbb{R}^n \rightarrow \mathbb{R}^n$. We write $x = G(z)$ and $z = G^{-1}(x)$, where $x \in \mathbb{R}^n$ is an image that corresponds to the latent representation $z \in \mathbb{R}^n$. We will consider a G that has the Glow architecture introduced in (Kingma & Dhariwal, 2018). For a short introduction to the Glow model, see section 2.1.

We consider recovering an image x from possibly-noisy linear measurements of the form,

$$y = Ax + \eta,$$

where $A \in \mathbb{R}^{m \times n}$ is a measurement matrix and $\eta \in \mathbb{R}^m$ models noise. Given a pretrained invertible generator G , we have access to likelihood estimates for all images $x \in \mathbb{R}^n$. Hence, it is natural to attempt to solve the above inverse problem by a maximum likelihood formulation given by

$$\min_{x \in \mathbb{R}^n} \|Ax - y\|^2 - \gamma \log p_G(x),$$

where p_G is the likelihood function over x induced by G , $\|\cdot\|$ is the Euclidean norm, and γ is a hyperparameter. We have found this formulation to be empirically challenging to optimize, and instead we solve an optimization problem over latent space that encourages high likelihood latent representations. In the case of denoising, we solve

$$\min_{z \in \mathbb{R}^n} \|AG(z) - y\|^2 + \gamma \|z\|^2 \quad (1)$$

In the case of compressive sensing, we fix $\gamma = 0$ and solve

$$\min_{z \in \mathbb{R}^n} \|AG(z) - y\|^2 \quad (2)$$

Unless otherwise stated, we initialize both formulations at $z_0 = 0$.

The motivation for formulations (1) and (2) is as follows. As a proxy for the likelihood of an image $x \in \mathbb{R}^n$, we will use the likelihood of its latent representation $z = G^{-1}(x)$. Because the invertible network G was trained to map a standard normal in \mathbb{R}^n to a distribution over images, the log-likelihood of a latent representation z is proportional to $\|z\|^2$. The model induces a probability distribution over the affine space of images consistent with some given measurements, and so our proxy turns the likelihood maximization task over an affine space in x into the geometric task of finding the point on a manifold in z -space that is closest to the origin with respect to the Euclidean norm. In order to approximate that point, we run a gradient descent in z down the data misfit term starting at $z_0 = 0$.

In principle, this proxy is imperfect in that some high likelihood latent codes may correspond to low likelihood images. We find that the set of such images has low total probability and they are inconsistent with enough provided measurements. For further discussion of our choice of proxy and initialization at $z_0 = 0$, see the Supplemental Materials.

In all experiments that follow, we use an invertible Glow model, as in (Kingma & Dhariwal, 2018). Due to computational considerations, we run most of our experiments on 64×64 px color images with the pixel values scaled between $[0, 1]$. For some compressive sensing experiments, we additionally trained a 128×128 px Glow model in order to replicate results at this larger size. Once trained, the Glow prior is fixed for use in each of the inverse problems below.

For comparison to GAN architectures, we train a 64×64 DCGAN architecture (Radford et al., 2016) and a 128×128 PGGAN architecture (Karras et al., 2018). To use these priors in inverse problems, we use the formulation from (Bora et al., 2017), which is the formulation above in the case where the optimization is performed over \mathbb{R}^k , $\gamma = 0$, and initialization is selected randomly. For comparison to an unlearned neural image prior, we implement an overparameterized variant of a Deep Decoder prior at both resolutions as in (Heckel & Hand, 2019). In all experiments, we solve (2) using L-BFGS for Glow, and Adam (Kingma & Ba, 2015) for the Deep Decoder, DCGAN, and PGGAN architectures. DCGAN and PGGAN results are reported for an average of 3 runs because we observed some variance due to random initialization. To measure the quality of recovered images, we use Peak Signal-to-Noise Ratio (PSNR). For more information on the training algorithms, hyperparameters, and parameter counts for each of the tested models, see the Supplemental Materials.

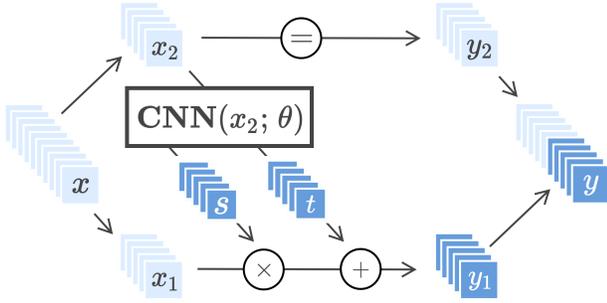


Figure 2. An Affine Coupling layer applies an affine transformation to half of the input data, here x_1 . The parameters of the affine transformations, s and t , can depend in a complex, learned way on the other half of the input data. The model can be inverted, even though s and t themselves are not invertible.

2.1. Details of the Glow Architecture

The Glow architecture (Kingma & Dhariwal, 2018) belongs to the class of *normalizing flow models*. A normalizing flow models output signals using a composition of many flow steps which are each individually invertible. In the Glow model, flow steps use an Affine Coupling layer, in which half of the input data is used to determine the scale and translation parameters of an affine transformation applied to the other half of the input data. This operation is shown schematically in Figure 2. Each affine transformation is invertible and has an upper-triangular Jacobian, making it computationally tractable to compute the Jacobian determinant of the entire normalizing flow. In turn, given a simple prior over the latent space, the model can be efficiently trained to sample structured, high-dimensional data by directly maximizing likelihood of the training data.

To ensure that each input component can affect each output component, the Glow models incorporate a pixelwise reshuffling. In other normalizing flows, such as RealNVP (Dinh et al., 2017), this is achieved by a fixed permutation, whereas in Glow it is achieved by a learned 1×1 convolution. Both models consist of multiscale architectures based on affine coupling and pixelwise reshuffling. We refer the reader to (Dinh et al., 2017) and (Kingma & Dhariwal, 2018) for more details.

3. Applications

3.1. Denoising

We consider the denoising problem with $A = I_n$ and $\eta \sim \mathcal{N}(0, \sigma^2 I_n)$, as given by formulation equation 1. We evaluate the performance of a Glow prior, a DCGAN prior, and BM3D for two different noise levels on 64px in-distribution images ($n = 64 \times 64 \times 3 = 12288$).

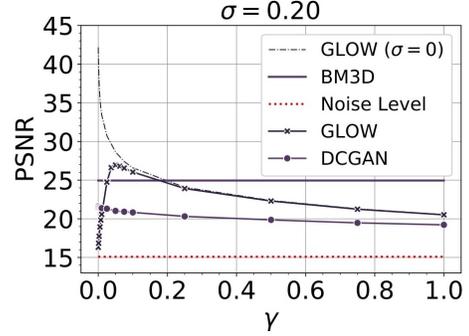


Figure 3. Recovered PSNR values as a function of γ for denoising by the Glow and DCGAN priors. Denoising results are averaged over $N = 50$ in-distribution test set images. For reference, we show the average PSNRs of the original noisy images, after applying BM3D, and under the Glow prior in the noiseless case ($\sigma = 0$).

Figure 3 shows the recovered PSNR values as a function of γ for denoising by the Glow and DCGAN priors, along with the PSNR by BM3D. The figure shows that the performance of the regularized Glow prior increases with γ , and then decreases. If γ is too low, then the network fits to the noise in the image. If γ is too high, then data fit is not enforced strongly enough. We study this effect for an extensive range of γ and noise levels, which may be found in the Supplemental Materials. We see in Figure 3 that an appropriately regularized Glow prior can outperform BM3D by almost 2 dB. The experiments also reveal that appropriately regularized Glow priors outperform the DCGAN prior, which suffers from representation error and is not aided by the regularization. A visual comparison of the recoveries at the noise level $\sigma = 0.1$ using Glow, DCGAN priors, and BM3D can be seen in Figure 4. Note that the recoveries with Glow are sharper than BM3D. See the Supplemental Materials for more quantitative and qualitative results.

3.2. Compressive Sensing

In compressive sensing, one is given undersampled linear measurements of an image, and the goal is to recover the image from those measurements. In our notation, $A \in \mathbb{R}^{m \times n}$ with $m < n$. As the image x is undersampled, there is an affine space of images consistent with the measurements, and an algorithm must select which is most ‘natural.’ A common proxy for naturalness in the literature has been sparsity with respect to the DCT or wavelet bases. With a GAN prior, an image is considered natural if it lies in or near the range of the GAN. For an invertible prior, image likelihood is a proxy for naturalness, and under our proxy for likelihood, we consider an image to be natural if it has a latent representation of small norm.

We study compressive sensing in the case that A is an $m \times n$ matrix of i.i.d. $\mathcal{N}(0, 1/m)$ entries and where η is standard

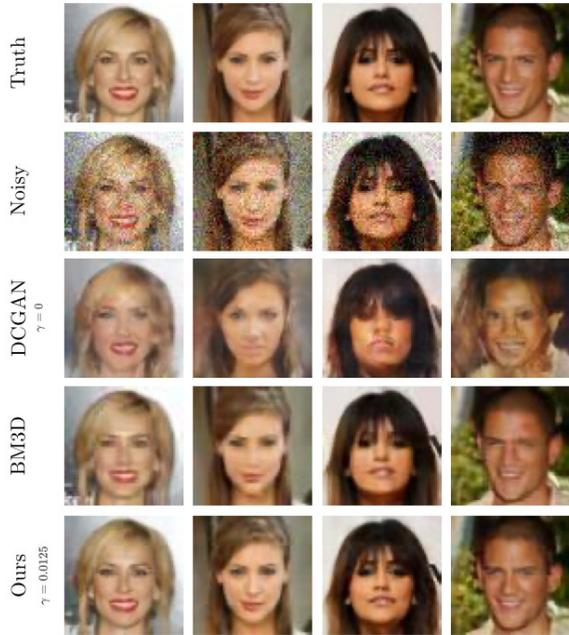


Figure 4. Denoising results using the Glow prior, the DCGAN prior, and BM3D at noise level $\sigma = 0.1$. Note that the Glow prior gives a sharper image than BM3D in these cases.

i.i.d. Gaussian noise normalized such that $\sqrt{\mathbb{E}\|\eta\|^2} = 0.1$. We present our results under formulation (2), the $\gamma = 0$ simplification of (1).

We compare the Glow prior to various other learned and unlearned image priors. This includes GANs, as used in (Bora et al., 2017), and IAGAN, as used in (Hussein et al., 2020). In the 64px case ($n = 64 \times 64 \times 3 = 12288$), we compare to a DCGAN, and in the 128px case ($n = 128 \times 128 \times 3 = 49152$) we compare to a PGGAN, both trained on the CelebA-HQ dataset. We also compare to unlearned image priors, including an overparameterized Deep Decoder and a sparsity prior in the DCT basis¹. To assess the performance of each image prior, we report the mean PSNR of recovered test set images from both the training distribution of the learned models ("in-distribution" images) and other datasets ("out-of-distribution" images). Our in-distribution images are sampled from a test set of CelebA-HQ images which were withheld from all learned models during their training procedures. Out-of-distribution images are sampled randomly from the Flickr Faces High Quality dataset, which provides images with features of variation that are rare among CelebA images (eg. skin tone, age, beards, and glasses) (Karras et al., 2019). The 64px and 128px recovery experiments have test sets with $N = 1000$ and $N = 100$ images respectively.

¹The inverse problems with Lasso were solved by $\min_z \|A\Phi z - y\|_2^2 + 0.01\|z\|_1$ using coordinate descent. We observe similar performance between the DCT basis and a Wavelet basis.

Surprisingly, the Glow prior exhibits superior performance on compressive sensing tasks with no likelihood penalization in the objective (2). We find that $z_0 = 0$ is a particularly good choice of initialization, for which one does not benefit from likelihood penalization, while for other choices of initialization direct penalization of likelihood may improve performance. We provide additional experiments exploring this phenomena in the Supplemental Materials.

As shown in Figure 5, we find that on its training distribution, the Glow prior outperforms both the learned and unlearned alternatives for a wide range of undersampling ratios. Surprisingly, in the case of extreme undersampling, Glow substantially outperforms these methods even though it does not maintain a direct low-dimensional parameterization of the signal manifold. In both the 64px and 128px cases, the GAN architectures quickly saturate due to their representation error. For out-of-distribution images, the Glow prior exhibits graceful performance decay, and is still highly performant in a large measurement regime. See figures 6 and 7 for a visual comparison of recovered images from the CelebA and FFHQ test sets for the 128px case. The PGGAN’s performance reveals biases of the underlying dataset and limitations of low-dimensional modeling, as the PGGAN fails completely to represent features like darker skin tones or accessories, which are uncommon in CelebA. In contrast, the Glow prior mitigates this bias, demonstrating image recovery for natural images that are not representative of the CelebA training set.

3.3. Inpainting

In inpainting, one is given a masked image of the form $y = M \odot x$, where M is a masking matrix with binary entries and $x \in \mathbb{R}^n$ is an n -pixel image. As in compressive sensing, we solve (2) to try to recover x , among the affine space of images consistent with the measurements. As before, using the minimizer \hat{z} , the estimated image is given by $G(\hat{z})$. We show qualitative inpainting results in Fig. 8. Our experiments reveal the same story as for compressive sensing. If initialized at $z_0 = 0$, then Glow model under the empirical risk formulation with $\gamma = 0$ exhibits high PSNRs on test images while the DCGAN is limited by its representation error.

4. Theory

We now introduce a likelihood-based theory for compressive sensing under invertible priors in the case of a linear invertible model. We will provide an estimate on the expected error of the recovered signal in terms of the singular values of the model. Specifically, we consider a fixed invertible linear $G : \mathbb{R}^n \rightarrow \mathbb{R}^n$. Assume signals are generated by a distribution $p_G(x)$, given by $x = Gz$, where $z \sim \mathcal{N}(0, I_n)$. Equivalently, $p_G = \mathcal{N}(0, GG^T)$.

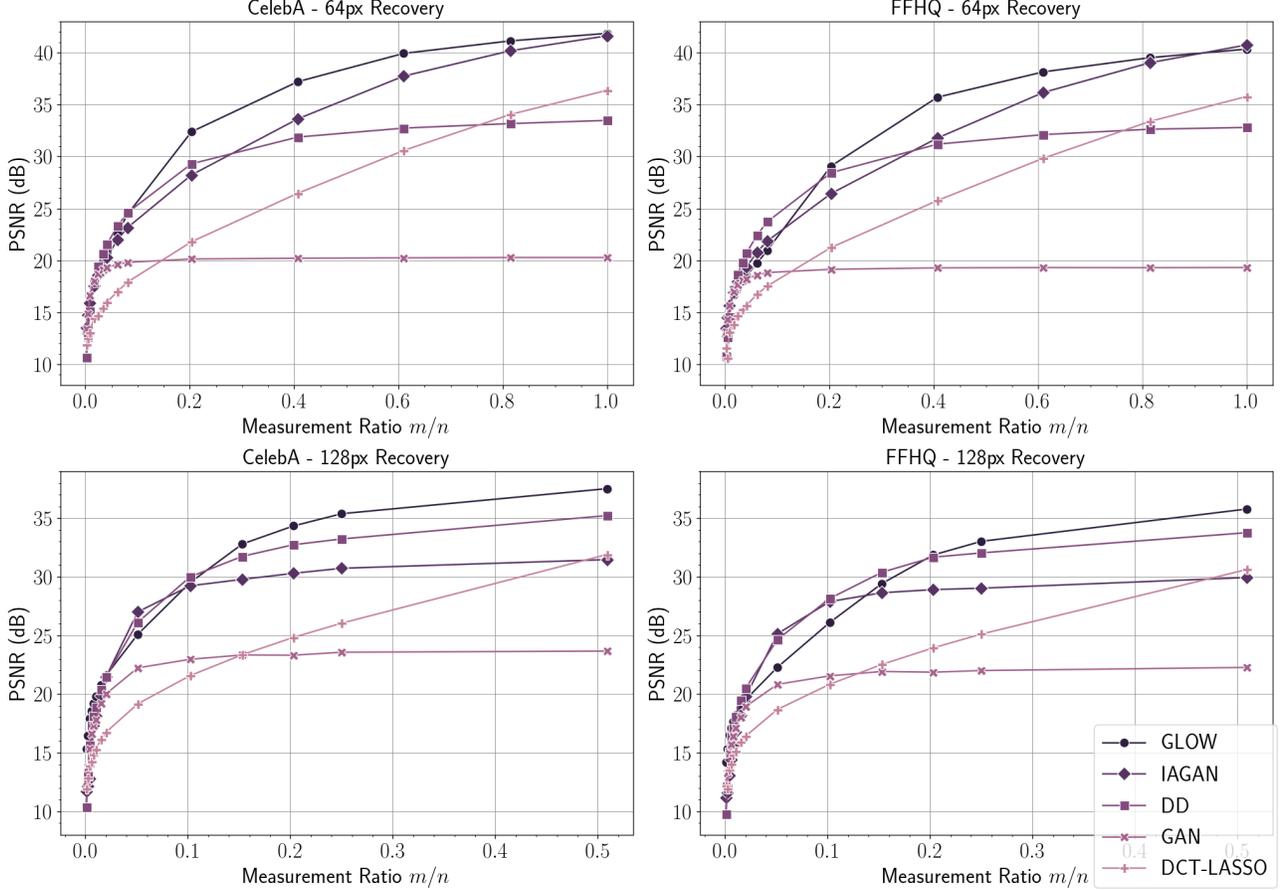


Figure 5. Performance of Glow, GAN, and IAGAN priors (learned) and the Deep Decoder and Lasso-DCT priors (unlearned) across various undersampling ratios in the 64px and 128px case. The 64px and 128px experiments use $N = 1000$ and $N = 100$ test set images respectively.

For an unknown sample $x_0 \sim p_G$, suppose we are provided noiseless measurements Ax_0 , where $A \in \mathbb{R}^{m \times n}$ has i.i.d. $\mathcal{N}(0, 1)$ entries. We consider the maximum likelihood estimate of x_0 :

$$\hat{x} := \arg \max_{x \in \mathbb{R}^n} p_G(x) \text{ s.t. } Ax = Ax_0. \quad (3)$$

The following theorem provides both upper and lower bounds on the absolute expected squared error in terms of the singular values of G .

Theorem 1. Suppose $x_0 \sim p_G$ where $p_G = \mathcal{N}(0, GG^T)$ and $G \in \mathbb{R}^{n \times n}$ has singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$. Let $A \in \mathbb{R}^{m \times n}$ have i.i.d. $\mathcal{N}(0, 1)$ entries where $4 \leq m < n$. Then the maximum likelihood estimator \hat{x} obeys

$$\sum_{i>m} \sigma_i^2 \leq \mathbb{E}_A \mathbb{E}_{x_0 \sim p_G} \|\hat{x} - x_0\|^2 \leq m \sum_{i>m-2} \sigma_i^2. \quad (4)$$

The relative expected squared error could be computed by $\frac{\mathbb{E}_A \mathbb{E}_{x_0 \sim p_G} \|\hat{x} - x_0\|^2}{\mathbb{E}_{x_0 \sim p_G} \|x_0\|^2}$, noting that $\mathbb{E}_{x_0 \sim p_G} \|x_0\|^2 = \sum_{k \geq 1} \sigma_k^2$.

The lower bound of this theorem establishes that under a linear invertible generative model, m Gaussian measurements give rise to at least as much error as would be given by the best m -dimensional signal model, i.e. the model corresponding to the only the top m directions of highest variance. The upper bound on this theorem establishes that up to a factor of m , expected square error is bounded above by the error given by the best $m - 2$ dimensional model.

Note that if the singular values decay quickly enough, then the expected recovery error under the linear invertible model decreases to a small value as $m \rightarrow n$. Specifically, this is achieved if $\sigma_i = o(i^{-1/2})$. This conclusion is in contrast with the theory for GANs with low latent dimensionality. In that literature, recovery error does not decrease to 0 as $m \rightarrow n$; instead it saturates at the representation error of the GAN model. In the present case of a linear model, the best k -dimensional model would have expected square recovery error at least $\sum_{i>k} \sigma_i^2$ regardless of m . Similar work in (Yu & Sapiro, 2011) also showed that m Gaussian measurements of a Gaussian signal give rise to an error proportional to



Figure 6. Compressive sensing on CelebA images with $m = 7,500$ ($\approx 20\%$) of measurements. Visual comparisons: CS under the Glow prior, PGGAN prior, and the overparameterized Deep Decoder prior. For images in-distribution, we observe qualitatively sharper recoveries from the Glow Prior than from the Deep Decoder.

the best m -term approximation. The difference between that work and our results is that we have an explicit upfront constant in the upper bound whereas in (Yu & Sapiro, 2011) it is estimated via Monte Carlo simulations.

The theorem is proved in the Supplemental Materials.

5. Discussion

We have demonstrated that pretrained generative invertible models can be used as natural signal priors in imaging inverse problems. Their strength is that *every desired image is in the range of an invertible model*, and the challenge that they overcome is that *every undesired image is also in the range of the model and no explicit low-dimensional representation is kept*. We demonstrate that this formulation can quantitatively and qualitatively outperform BM3D at denoising. For compressive sensing on in-distribution images, invertible priors can have lower recovery errors than Deep Decoder, GANs with low dimensional latent representations, and Lasso, across a wide range of undersampling ratios. We show that the performance of our invertible prior behaves gracefully with slight performance drops for out-of-distribution images. We additionally prove a theoretical upper and lower bound for expected squared recover error in the case of a linear invertible generative model.

The idea of analyzing inverse problems with invertible neural networks has appeared in Ardizzone et al. (2019). The

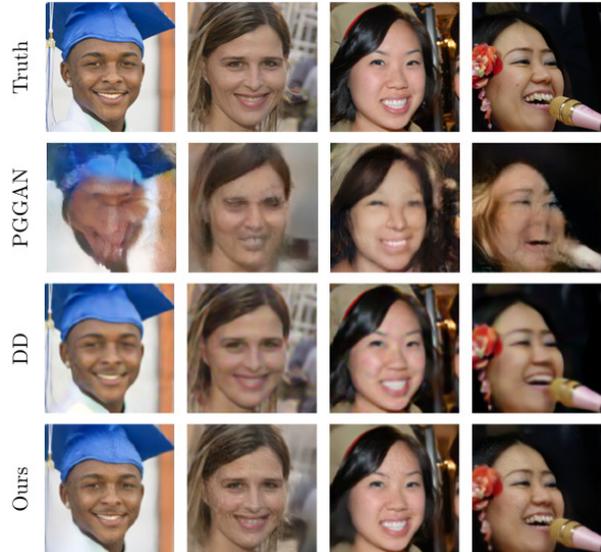


Figure 7. Compressive sensing on FFHQ images with $m = 7,500$ ($\approx 20\%$) of measurements. Visual comparisons: CS under the Glow prior, PGGAN prior, and the overparameterized Deep Decoder prior. For images out-of-distribution, the images recovered by the Deep Decoder and the Glow priors are both qualitatively and quantitatively (by PSNR) comparable.

authors study estimation of the complete posterior parameter distribution under a forward process, conditioned on observed measurements. Specifically, the authors approximate a particular forward process by training an invertible neural network. The inverse map is then directly available. In order to cope with information loss, the authors augment the measurements with additional variables. This work differs from ours because it involves training a separate model for every particular inverse problem.

In contrast, our work studies how to use a pretrained invertible generator for a variety of inverse problems not known at training time. Training invertible networks is challenging and computationally expensive; hence, it is desirable to separate the training of off-the-shelf invertible models from potential applications in a variety of scientific domains. In additional work by Putzky & Welling (2019), the authors also exploit the efficient gradient calculations of invertible nets for improved MR reconstruction.

Why do invertible neural networks perform well for both in-distribution and out-of-distribution images?

One reason that the invertible prior performs so well is because it has no representation error. Any image is potentially recoverable, even if the image is significantly outside of the training distribution. In contrast, methods based on projecting onto an explicit low-dimensional representation of a natural signal manifold (such as GAN priors) will have

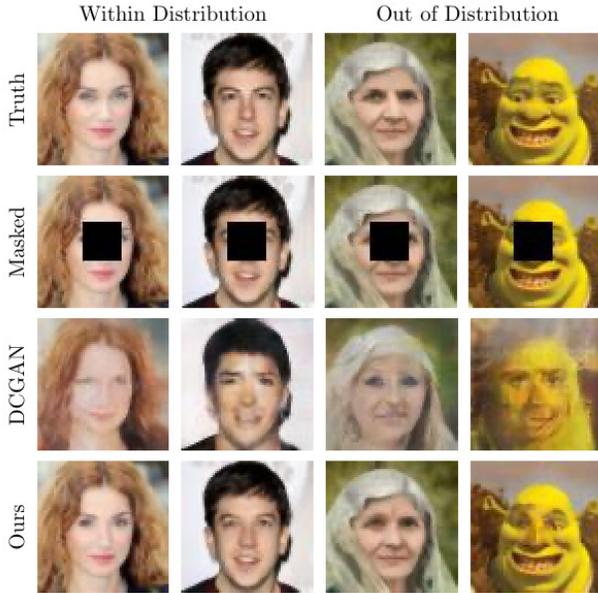


Figure 8. Inpainting on random images with the Glow and DCGAN priors. In the left columns, images are taken from the training distribution of the learned priors, and the right column includes two images from the wild.

representation error, perhaps due to modeling assumptions, mode collapse, or bias in a training set. Such methods will see performance prematurely saturate as the number of measurements increases. In contrast, an invertible prior would not see performance saturate. In the extreme case of having a full set of exact measurements, an invertible prior could in principle recover any image exactly.

How do invertible priors respect the low dimensionality of natural signals? And how does our theory inform this?

A surprising feature of invertible priors is that they perform well even though they do not maintain explicit low-dimensional representations of natural signals. Instead they have a fully dimensional representation of the natural signal class. Naturally, any signal class will truly be fully dimensional, for example due to sensor noise, but with different importances to different directions. Those directions in latent space corresponding to noise perturbations will have much smaller of an effect than corresponding perturbations in semantically meaningful directions. As an illustration, we observe with trained Glow models that the singular values of the Jacobian of G at a natural image exhibit significant decay, as we show in the Supplemental Materials. In principle, signal models of any given dimensionality could be extracted from G , though it is not obvious how to compute these. *The power of invertible priors is that each additional measurement acts to roughly increment the dimensionality of the modeled natural signal manifold. That is, more measurements permit exploiting a higher dimensional and,*

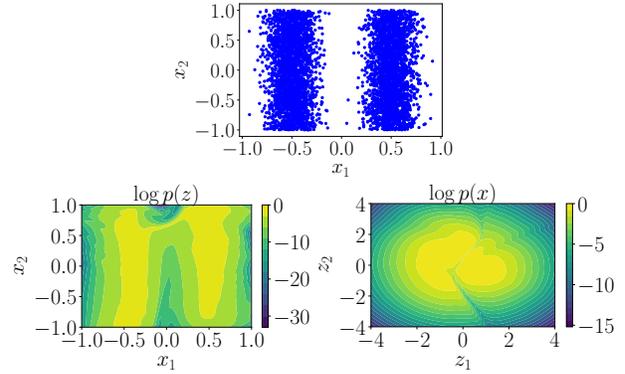


Figure 9. An invertible net was trained on the data points in x -space (top), resulting in the given plots of latent z -likelihood versus x (bottom left), and x -likelihood versus latent representation z (bottom right).

hence, lower-error signal model. In contrast, GAN-based recovery theory exploits a model of fixed dimensionality regardless of the number of measurements. Our theoretical analysis provides justification for this explanation based on linear invertible generators G . Each singular value σ_i of G quantifies the signal variation due to a particular direction in latent space, and the expected squared error given m random measurements is upper bounded by $m \sum_{i>m-2} \sigma_i^2$. Note that the best m -dimensional manifold (given by the top m singular values and vectors) would yield at best an expected squared error of $\sum_{i>m} \sigma_i^2$. The multiplicative m term and the sum's starting index may not be optimal, and improvement of this bound is left for future research.

Why is the likelihood of an image's latent representation a reasonable proxy for the image's likelihood?

The training process for an invertible generative model attempts to learn a target distribution in image space by directly maximizing the likelihood of provided samples from that distribution, given a standard Gaussian prior in latent space. High probability regions in latent space map to regions in image space of equal probability. Hence, broadly speaking, regions of small values of $\|z\|$ are expected to map to regions of large likelihoods in image space. There will be exceptions to this property. For example, natural image distributions have a multimodal character. The preimage of high probability modes in image space will correspond to high likelihood regions in latent space. Because the generator G is invertible and continuous, interpolation in latent space of these modes will provide images of high likelihood in z but low likelihood in the target distribution. To demonstrate this, we trained a Real-NVP (Dinh et al., 2017) invertible neural network on the two dimensional set of points depicted in Figure 9 (top panel). The lower left and right panels show that high likelihood regions in latent space

generally correspond to higher likelihood regions in image space, but that there are some regions of high likelihood in latent space that map to points of low likelihood in image space and in the target distribution. We see that the spurious regions are of low total probability and would be unlikely to be the desired outcomes of an inverse problem arising from the target distribution.

How can solving compressive inverse problems be successful without direct penalization of the image likelihood or its proxy?

If there are fewer linear measurements than the dimensionality of the desired signal, an affine space of images is consistent with the measurements. In our formulation, regularization does not occur by direct penalization by our proxy for image likelihood; instead, it occurs implicitly by performing the optimization in z -space with an initialization of $z_0 = 0$. The set of latent representations z that are consistent with the compressive measurements define a m -dimensional nonlinear manifold. As per the likelihood proxy mentioned above, the spirit of our formulation is to find the point on this manifold that is closest to the origin with respect to the Euclidean norm. Our specific way of estimating this point is to perform a gradient descent down a data misfit term in z -space, starting at the origin. While a gradient flow typically will not find the closest point on the manifold, it empirically finds a reasonable approximation of that point. This type of algorithmic regularization is akin to the linear invertible model setting where it is well known that running gradient descent to optimize an underdetermined least squares problem with initialization $z_0 = 0$ will converge to the minimum norm solution.

The results of this paper provide further evidence that reducing representational error of generators can significantly enhance the performance of generative models for inverse problems in imaging. This idea was also recently explored in (Athar et al., 2019), where the authors trained a GAN-like prior with a high-dimensional latent space. The high dimensionality of this space lowers representational error, though it is not zero. In their work, the high-dimensional latent space had a structure that was difficult to directly optimize, so the authors successfully modeled latent representations as the output of an untrained convolutional neural network whose parameters are estimated at test time.

Their paper and ours raises an important question: which generator architectures provide a good balance between low representation error, ease of training, and ease of inversion? This question is important, as solving (2) in our 128×128 px color images experiments took on the order of 11 minutes using an NVIDIA 2080 Ti GPU.

New developments are needed on architectures and frameworks in between low-dimensional generative priors and fully invertible generative priors. Such methods could lever-

age the strengths of invertible models while being much cheaper to train and use.

Acknowledgements

PH is partially supported by NSF CAREER Grant DMS-1848087. OL acknowledges support by the NSF Graduate Research Fellowship under Grant No. DGE-1450681.

References

- Ardizzone, L., Kruse, J., Rother, C., and Köthe, U. Analyzing inverse problems with invertible neural networks. In *International Conference on Learning Representations*, 2019.
- Asim, M., Shamshad, F., and Ahmed, A. Blind image deconvolution using deep generative priors. *arXiv:1802.04073 [cs]*, Feb 2019. arXiv: 1802.04073.
- Athar, S., Burnaev, E., and Lempitsky, V. Latent convolutional models. In *International Conference on Learning Representations*, 2019.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 537–546. PMLR, Aug 2017.
- Chen, Y., Shi, F., Christodoulou, A. G., Xie, Y., Zhou, Z., and Li, D. Efficient and accurate mri super-resolution using a generative adversarial network and 3d multi-level densely connected network. In Frangi, A. F., Schnabel, J. A., Davatzikos, C., Alberola-López, C., and Fichtinger, G. (eds.), *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018 - 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*, volume 11070 of *Lecture Notes in Computer Science*, pp. 91–99. Springer, 2018. doi: 10.1007/978-3-030-00928-1_11.
- Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8): 2080–2095, 2007.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. 2017.
- Eckart, C. and Young, G. The approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936.
- Gomez, A. N., Ren, M., Urtasun, R., and Grosse, R. B. The reversible residual network: Backpropagation without

- storing activations. In Guyon, I., Luxburg, U. v., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pp. 2214–2224, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. *Generative Adversarial Nets*, pp. 2672–2680. Curran Associates, Inc., 2014.
- Halko, N., Martinsson, P.-G., and Tropp, J. A. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53:217–288, 2011.
- Hand, P. and Voroninski, V. Global guarantees for enforcing deep generative priors by empirical risk. In Bubeck, S., Perchet, V., and Rigollet, P. (eds.), *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pp. 970–978. PMLR, 2018.
- Hand, P., Leong, O., and Voroninski, V. *Phase Retrieval Under a Generative Prior*, pp. 9136–9146. Curran Associates, Inc., 2018.
- Heckel, R. and Hand, P. Deep decoder: Concise image representations from untrained non-convolutional networks. In *International Conference on Learning Representations*, 2019.
- Hussein, S. A., Tirer, T., and Giryas, R. Image-adaptive gan based reconstruction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 3121–3129. AAAI Press, 2020.
- Jacobsen, J.-H., Smeulders, A. W. M., and Oyallon, E. i-revnet: Deep invertible networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of gans for improved quality, stability, and variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pp. 4401–4410. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00453.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 10236–10245, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In Bengio, Y. and LeCun, Y. (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 3730–3738. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.425.
- Minaee, S. and Abdolrashidi, A. Finger-gan: Generating realistic fingerprint images using connectivity imposed gan. *CoRR*, abs/1812.10482, 2018.
- Putzky, P. and Welling, M. Invert to learn to invert. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 444–454, 2019.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pp. 1278–1286. JMLR.org, 2014.

Shin, H.-C., Tenenholtz, N. A., Rogers, J. K., Schwarz, C. G., Senjem, M. L., Gunter, J. L., Andriole, K. P., and Michalski, M. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In Gooya, A., Goksel, O., Oguz, I., and Burgos, N. (eds.), *Simulation and Synthesis in Medical Imaging - Third International Workshop, SASHIMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings*, volume 11037 of *Lecture Notes in Computer Science*, pp. 1–11. Springer, 2018. doi: 10.1007/978-3-030-00536-8_1.

Ulyanov, D., Vedaldi, A., and Lempitsky, V. S. Deep image prior. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 9446–9454. IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00984.

Wang, Z., Bovik, A., Sheikh, H., and Simoncelli, E. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, Apr 2004. ISSN 1941-0042. doi: 10.1109/TIP.2003.819861.

Yu, G. and Sapiro, G. Statistical compressed sensing of gaussian mixture models. *IEEE Trans. Signal Process.*, 59(12):5842–5858, 2011. doi: 10.1109/TSP.2011.2168521.

6. Appendix

6.1. Proofs

We now proceed with the proof of Theorem 1. Throughout the proof, we use $\stackrel{d}{=}$ to denote equality in distribution.

Proof of Theorem 1. We first quantify the expected error over the distribution p_G . Observe that since $x_0 \sim p_G = \mathcal{N}(0, GG^T)$, it can be written as $x_0 = Gz_0$ where $z_0 \sim \mathcal{N}(0, I_n)$. Then

$$\mathbb{E}_{x_0 \sim p_G} \|\hat{x} - x_0\|^2 = \mathbb{E}_{z_0 \sim \mathcal{N}(0, I_n)} \|G\hat{z} - Gz_0\|^2$$

where \hat{z} is the maximum likelihood estimator in latent space:

$$\hat{z} := \arg \min_{z \in \mathbb{R}^n} \frac{1}{2} \|z\|^2 \text{ s.t. } AGz = AGz_0.$$

Since A has i.i.d. $\mathcal{N}(0, 1)$ entries, AG has full rank with probability 1 so \hat{z} is given explicitly by

$$\hat{z} = G^T A^T (AGG^T A^T)^{-1} AGz_0 =: \mathcal{P}_{G^T A^T} z_0$$

where $\mathcal{P}_{G^T A^T}$ is the orthogonal projection onto the range of $G^T A^T$. Thus

$$\mathbb{E}_{z_0 \sim \mathcal{N}(0, I_n)} \|G\hat{z} - Gz_0\|^2 = \mathbb{E}_{z_0 \sim \mathcal{N}(0, I_n)} \|G \mathcal{P}_{G^T A^T} z_0 - Gz_0\|^2 = \|G(\mathcal{P}_{G^T A^T} - I_n)\|_F^2$$

where the last equality follows from Lemma 1.

Without loss of generality, it suffices to consider the case when G is diagonal. To see this, let $U\Sigma V^T$ be the SVD of G . Observe that

$$\mathcal{P}_{G^T A^T} = \mathcal{P}_{V\Sigma U^T A^T} = V \mathcal{P}_{\Sigma U^T A^T} V^T \stackrel{d}{=} V \mathcal{P}_{\Sigma A^T} V^T$$

where we used Lemma 2 in the second equality and used the rotational invariance of A in the last equality. Thus we have

$$\|G(\mathcal{P}_{G^T A^T} - I_n)\|_F^2 \stackrel{d}{=} \|G(V \mathcal{P}_{\Sigma A^T} V^T - I_n)\|_F^2.$$

Moreover, note that

$$\|G(V \mathcal{P}_{\Sigma A^T} V^T - I_n)\|_F^2 = \|\Sigma V^T (V \mathcal{P}_{\Sigma A^T} V^T - I_n)\|_F^2 = \|\Sigma \mathcal{P}_{\Sigma A^T} V^T - \Sigma V^T\|_F^2 = \|\Sigma(\mathcal{P}_{\Sigma A^T} - I_n)\|_F^2$$

where we used the unitary invariance of the Frobenius norm in the first and last equality and the orthogonality of V in the second equality. Hence

$$\mathbb{E}_A \|G(\mathcal{P}_{G^T A^T} - I_n)\|_F^2 = \mathbb{E}_A \|\Sigma(\mathcal{P}_{\Sigma A^T} - I_n)\|_F^2 = \mathbb{E}_A \|(I_n - \mathcal{P}_{\Sigma A^T})\Sigma\|_F^2$$

so we assume without loss of generality that $G = \text{diag}(\sigma_1, \dots, \sigma_n)$ and consider $\mathbb{E}_A \|(I_n - \mathcal{P}_{GA^T})G\|_F^2$.

We now compute the lower bound of the expected error. Note that $\text{rank}(\mathcal{P}_{GA^T}) = m$ since A is full rank with probability 1. Also for any random draw of Gaussian A ,

$$\min_{S \in \mathbb{R}^{n \times n}, \text{rank}(S) \leq m} \|G - S\|_F^2 \leq \|(I_n - \mathcal{P}_{GA^T})G\|_F^2. \quad (5)$$

By the Eckart-Young Theorem (Eckart & Young, 1936),

$$\min_{S \in \mathbb{R}^{n \times n}, \text{rank}(S) \leq m} \|G - S\|_F^2 = \sum_{i > m} \sigma_i^2.$$

Taking the expectation with respect to A to the right hand-side of (5) establishes the lower bound in the theorem.

The upper bound comes from the following result in (Halko et al., 2011).

Theorem 2 (Minor variant of Theorem 10.5 in (Halko et al., 2011)). *Suppose G is a real $\ell \times n$ matrix with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min\{\ell, n\}} \geq 0$. Choose a target rank $k \geq 2$ and oversampling factor $m - k \geq 2$ where $m \leq \min\{\ell, n\}$. Draw an $n \times m$ Gaussian matrix A^T and construct the sample matrix GA^T . Then the expected approximation error*

$$\mathbb{E}_A \|(I_n - \mathcal{P}_{GA^T})G\|_F^2 \leq \left(1 + \frac{k}{m - k - 1}\right) \sum_{j>k} \sigma_j^2. \quad (6)$$

Theorem 10.5 in (Halko et al., 2011) literally states a bound on $\mathbb{E}_A \|(I_n - \mathcal{P}_{GA^T})G\|_F$ but in the proof, the authors show the stronger result (6). The upper bound in Theorem 1 follows by setting $\ell = n$ and $k = m - 2$ in Theorem 2 whereby the condition on k requires $4 \leq m \leq n$. □

The following are two supplementary lemmas used in the proof of Theorem 1. The first shows that the expected value of the ℓ_2 norm of a matrix acting on an isotropic gaussian vector is its frobenius norm.

Lemma 1. *Let $M \in \mathbb{R}^{m \times n}$. Then $\mathbb{E}_{z \sim \mathcal{N}(0, I_n)} \|Mz\|^2 = \|M\|_F^2$.*

Proof. Let $U\Sigma V^T$ be the SVD of M where $\Sigma \in \mathbb{R}^{m \times n}$ contains the singular values σ_i of M on the diagonal for $i = 1, \dots, r$ and $r = \text{rank}(M) \leq \min\{m, n\}$. Then

$$\|Mz\|^2 = \|U\Sigma V^T z\|^2 = \|\Sigma V^T z\|^2 \stackrel{d}{=} \|\Sigma z\|^2$$

where we used the unitary invariance of the ℓ_2 norm in the second equality and the rotational invariance of z in the last equality. The result follows by noting that

$$\mathbb{E}_{z \sim \mathcal{N}(0, I_n)} \|\Sigma z\|^2 = \sum_{i=1}^r \sigma_i^2 \mathbb{E}_{z_i \sim \mathcal{N}(0, 1)} z_i^2 = \sum_{i=1}^r \sigma_i^2 = \|M\|_F^2.$$

□

The second lemma asserts that unitary matrices exhibit a commutativity property when acting on the range of an orthogonal projector.

Lemma 2 (Proposition 8.4 in (Halko et al., 2011)). *Given a matrix M , let \mathcal{P}_M denote the orthogonal projection onto the range of M . Then for any unitary U , $U^T \mathcal{P}_M U = \mathcal{P}_{U^T M}$.*

6.2. Models

6.2.1. ARCHITECTURE AND TRAINING DETAILS

We train 128px and 64px variants of the Glow architecture (Kingma & Dhariwal, 2018). This model uses a sequence of invertible flow steps, each comprising of an activation normalization layer, a 1×1 invertible convolutional layer, and an affine coupling layer. Let K be the number of steps of flow before each splitting layer, and let L be the number of times the splitting is performed. For denoising, we use $K = 48$, $L = 4$ for 64px recovery. For compressive sensing, we use $K = 18$ and $L = 4$ for 64px recovery and $K = 32$, $L = 6$ for 128px recovery. All models are trained over 5-bit images using a learning rate of 0.0001 and 10,000 warmup iterations, as in (Kingma & Dhariwal, 2018). When solving inverse problems using Glow, original 8-bit images were used.

We observed numerical instability when solving inverse problems in the 128px case, in that activations of the Glow network could become too large to compute during inversion. To mitigate this problem, we train a modified version of the Glow reference implementation in which we add a small constant $\epsilon = 0.0005$ to computed scale parameters in each Actnorm layer, thus preventing a “division by zero error” during inversion. Additionally, in forward passes of the Glow model we clip activations to the range $[-40, 40]$.

We train a 64px DCGAN (Radford et al., 2016) and a 128px PGGAN (Karras et al., 2018) in order to compare the Glow model to traditional GAN architectures. The DCGAN model has $d = 5$ upsampling layers implemented through transpose convolutions. The PGGAN model has $d = 5$ upsampling layers implemented through nearest neighbor upsampling followed by a convolutional layer. In both cases, the first layer of each GAN includes $k_0 = 512$ activation channels and after each upsampling operation the number of activation channels is reduced by half. These architectures are used as-is for both the CSGM (Bora et al., 2017) and the Image Adaptive (Hussein et al., 2020) compressive sensing procedures.

Lastly, we also use an overparametrized variant of the Deep Decoder (Heckel & Hand, 2019) in order to compare to an unlearned neural network for compressive sensing. We use $d = 5$ and $d = 6$ upsampling layers for the 64px and 128px cases respectively, implemented in both cases with convolutional layers followed by bilinear upsampling. The number of activation channels in each layer is held fixed at $k = 250$ and $k = 700$ for the 64px and 128px cases respectively.

64px Models	Hyperparameters	Repr. Size	Model Size	Overparam. Ratio
Glow (Denoising)	$K = 48, L = 4$, affine coupling	12,288	67,172,224	1.000
Glow (CS)	$K = 18, L = 4$, affine coupling	12,288	25,783,744	1.000
DCGAN	$k_0 = 512, d = 5$	100	3,576,804	0.0081
IA-DCGAN	$k_0 = 512, d = 5$	3,576,804	3,576,804	291.1
Deep Decoder	$k = 250, d = 5$	253,753	254,753	20.65
128px Models	Hyperparameters	Repr. Size	Model Size	Overparam. Ratio
Glow	$K = 32, L = 6$, affine coupling	49152	129,451,520	1.000
PGGAN	$k_0 = 512, d = 5$	512	13,625,904	0.0104
IA-PGGAN	$k_0 = 512, d = 5$	13,625,904	13,625,904	277.2
Deep Decoder	$k = 700, d = 6$	2,462,603	2,473,803	50.10

Figure 10. Summary of the model parameters for all models used in our experiments. Hyperparameters refer to model-specific configurations as described in the text. Representation Size (Repr. Size) is the dimensionality of each model’s image representation, i.e. the total number of optimizable parameters available during inversion in each inverse problem. The model size is the total sum of all parameters of each model, including those of the image representation. The Overparametrization Ratio (Overparam. Ratio) describes the representation size as a fraction of the output image dimensionality.

6.2.2. UNCONDITIONAL OUTPUT SAMPLES

To demonstrate the successful training of our generative models, we provide output samples from each of the learned models used in our experiments.

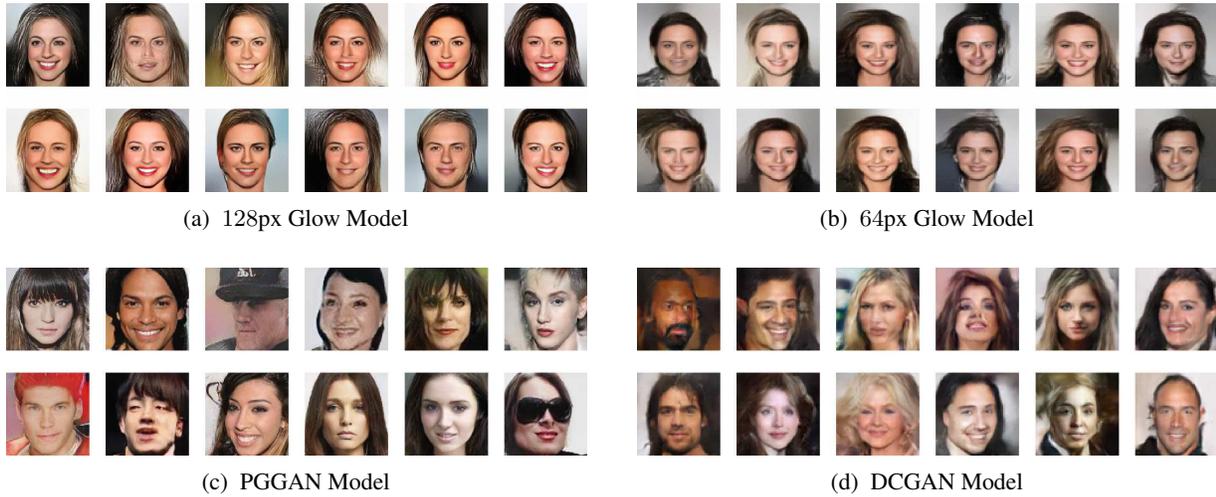


Figure 11. Samples from each of our trained generative models.

6.2.3. TEST SET IMAGES

Throughout our paper, we present experiments across two test sets, one from the same distribution of the trained generative models and another which is out-of-distribution. In particular, our in-distribution test set is sampled randomly from a validation split of the CelebA-HQ dataset (Karras et al., 2018) and our out-of-distribution test set is sampled randomly from the Flickr Faces High Quality (FFHQ) Dataset (Karras et al., 2019). None of the test set images are seen by the generative models during their training phases. A few samples from both test sets are shown below.

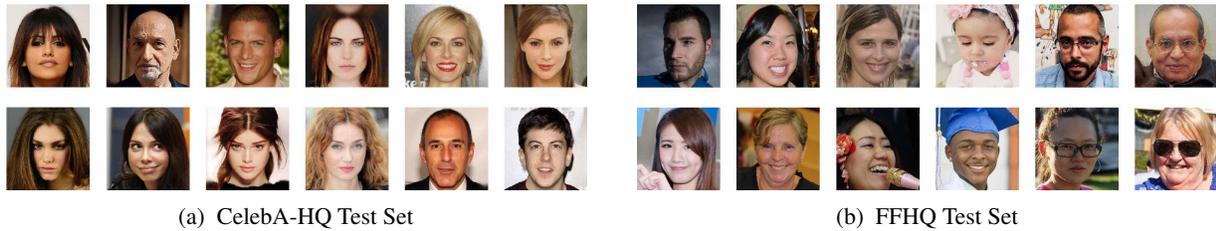


Figure 12. A few samples from each of the test sets used in our experiments.

6.3. Denoising: Additional Experiments

We additionally provide qualitative and quantitative denoising results over the test set of in-distribution CelebA images in Figures 13 and 14. Recall that we formulate denoising as the following empirical risk minimization problem:

$$\min_{z \in \mathbb{R}^n} \|G(z) - y\|^2 + \gamma \|z\|^2 \quad y = x_0 + \eta$$

Where $\eta \sim \mathcal{N}(0, \sigma^2 I_n)$ is additive Gaussian noise. PSNRs for varying choices of the penalization parameter γ under noise levels $\sigma = 0.01, 0.05, 0.1, \text{ and } 0.2$ are presented in Figure 13 below.

The central message is that the Glow prior outperforms the DCGAN prior uniformly across all γ due to the representation error of the DCGAN prior. In addition, when γ is chosen appropriately, regularization improves the performance of Glow, which can outperform the state-of-the-art BM3D algorithm at high noise levels such as $\sigma = 0.2$, and can offer comparable performance at lower noise levels. This is in contrast to the DCGAN prior, whose performance is harmed by increased regularization.

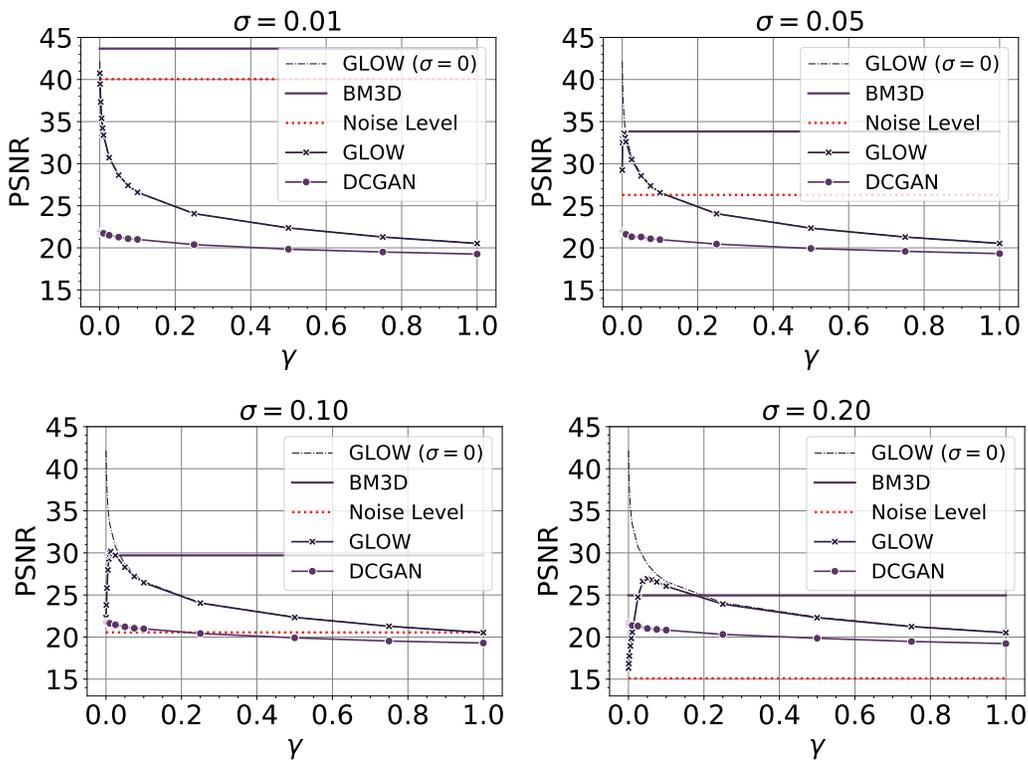


Figure 13. Image Denoising — Recovered PSNR values as a function of γ on $N = 50$ in-distribution test set CelebA images. We report the average PSNR after applying BM3D, under the DCGAN prior, and under the Glow prior, at noise levels $\sigma = 0.01, 0.05, 0.10, 0.20$. For reference, we also show the average PSNR of the original noisy images and the average PSNR of images recovered by the Glow prior in the noiseless case ($\sigma = 0$).

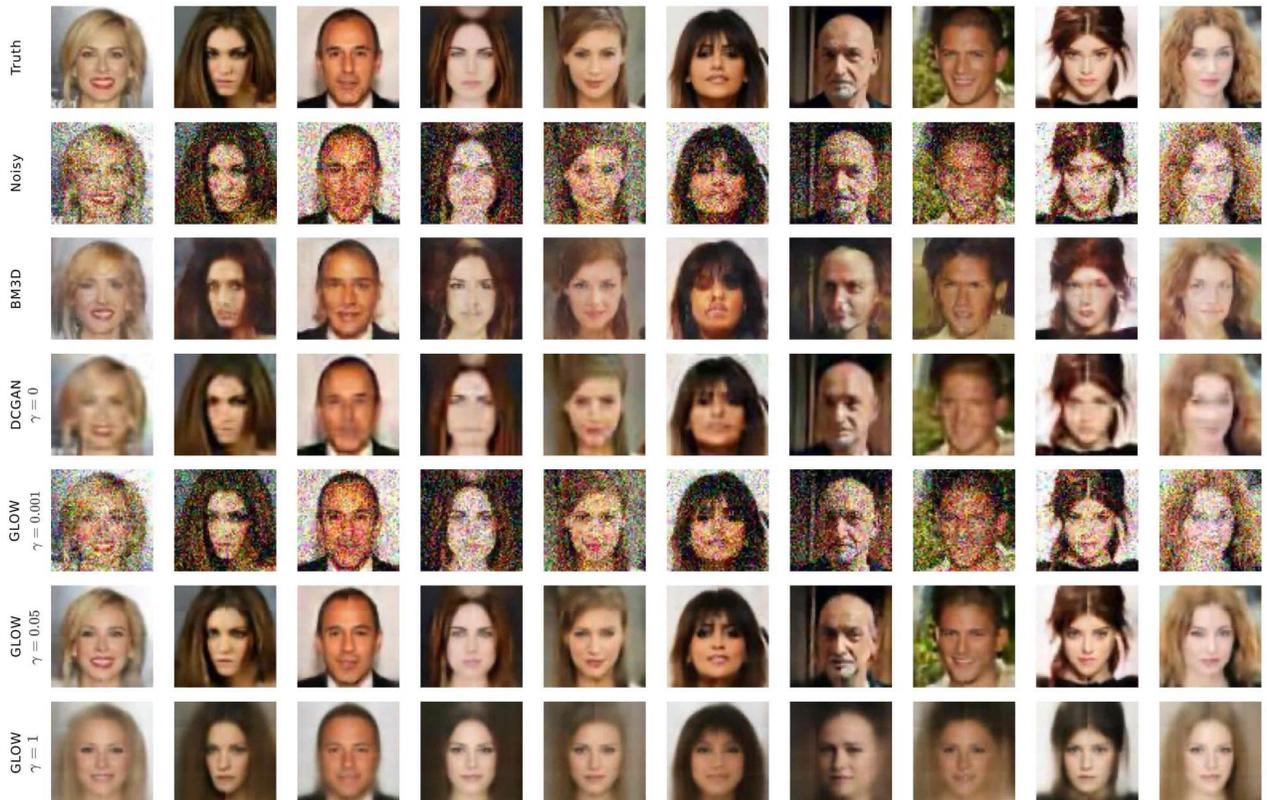


Figure 14. Image Denoising — Visual comparisons under the Glow prior, the DCGAN prior, and BM3D at noise level $\sigma = 0.2$ on a sample of CelebA in-distribution test set images. Under the DCGAN prior, we only show the case of $\gamma = 0$ as this consistently gives the best performance. Under the Glow prior, the best performance is achieved with $\gamma = 0.05$, overfitting of the image occurs with $\gamma = 0.001$ and underfitting occurs with $\gamma = 1$. Note that the Glow prior with $\gamma = 0.05$ also gives a sharper image than BM3D.

6.4. Compressive Sensing: Additional Experiments

We first analyze our formulation of compressive sensing as an empirical risk minimization problem. Recall that we solve the following optimization problem:

$$\min_{z \in \mathbb{R}^n} \|AG(z) - y\|^2 \quad y = Ax_0 + \eta$$

Where A is a Gaussian random measurement matrix with i.i.d. $\mathcal{N}(0, 1/m)$ entries, and η is noise normalized so that $\sqrt{\mathbb{E}\|\eta\|^2} = 0.1$. This formulation includes no explicit penalization on the likelihood of the latent code z , and instead relies on implicit regularization through the use of a gradient based optimization method initialized at $z_0 = 0$. To justify this formulation, we study various alternative methods of initialization: $z_0 = 0$, $z_0 \sim \mathcal{N}(0, 0.1^2 I_n)$, $z_0 \sim \mathcal{N}(0, 0.7^2 I_n)$, $z_0 = G^{-1}(x_0)$ with x_0 given by the solution to Lasso with respect to the wavelet basis, and $z_0 = G^{-1}(x_0)$ where x_0 is perturbed by a random point in the null space of A . For each initialization, we plot recovery performance on a compressive sensing task as a function of γ for the regularized objective:

$$\min_{z \in \mathbb{R}^n} \|AG(z) - y\|^2 + \gamma \|z\|^2$$

As shown in Fig. 15, the Glow model shows best performance with $\gamma = 0$ with initialization at the zero vector, despite there being no explicit penalization on the likelihood of recovered latent codes.

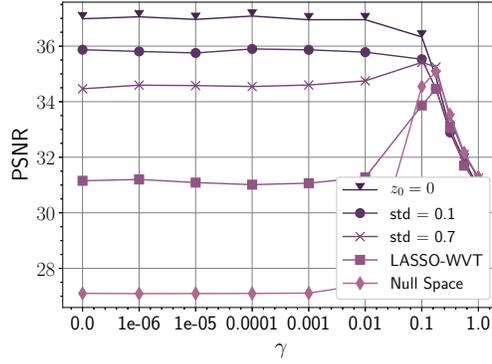


Figure 15. Performance of the Glow model in Compressive Sensing under various initialization strategies (as described in the text). The Glow model shows strongest performance when initialized with the zero vector and with no explicit penalization on the latent code likelihood. The task is 64px CelebA image recovery using $m = 5000$ ($\approx 50\%$) Gaussian measurements.

We show in Fig. 16 the results of our compressive sensing recovery experiments using the Structural Similarity index (SSIM) as a recovery quality metric (Wang et al., 2004). SSIM is designed to indicate perceptual quality by extracting structural information from images.

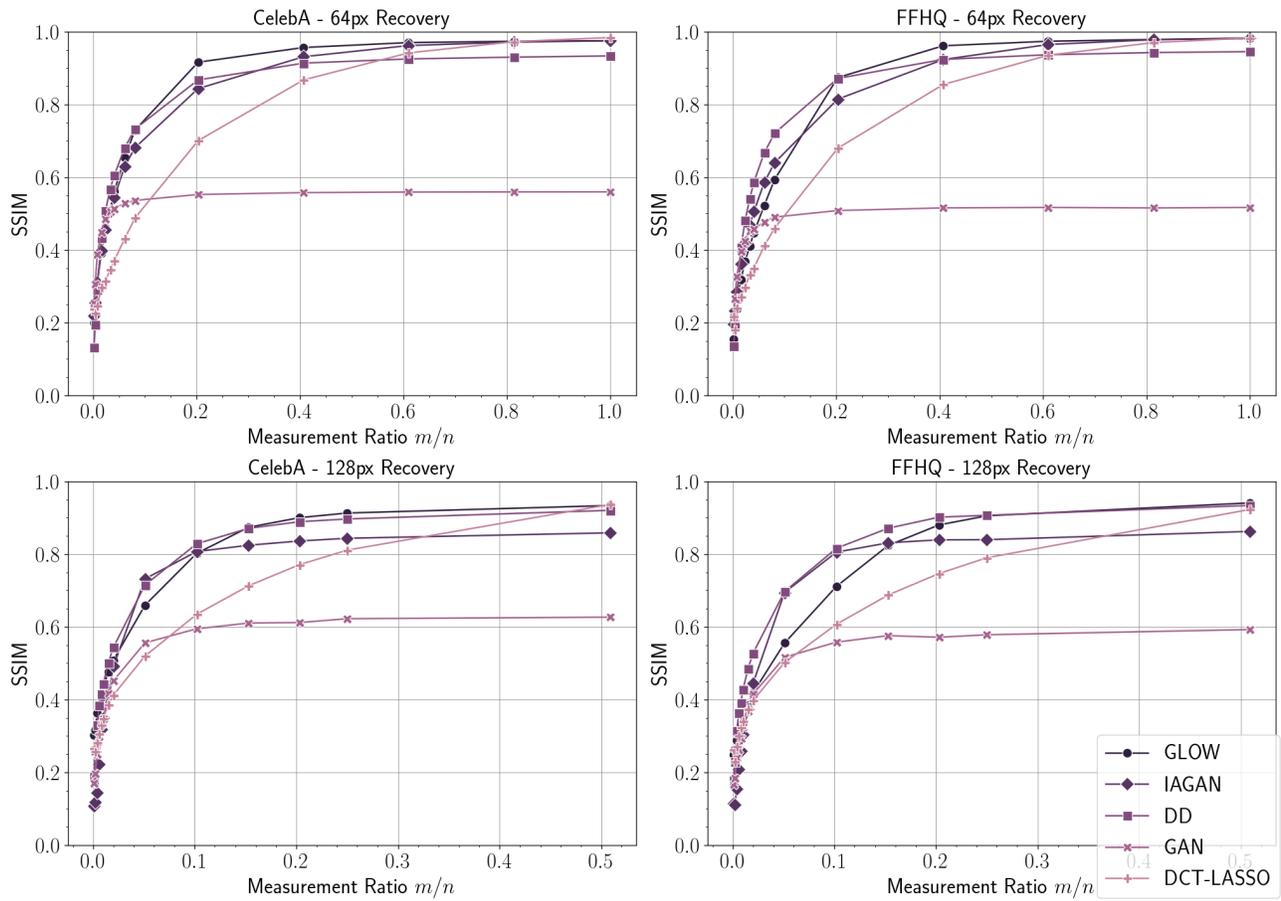


Figure 16. 64px and 128px Compressive Sensing results using $N = 1000$ and $N = 100$ test set images respectively. Each plot shows SSIM (higher is better, ranging in $[0, 1]$) across a variety of measurement ratios.

6.4.1. COMPRESSIVE SENSING 64PX SAMPLE SHEETS

We provide here additional sample sheets for our 64px experiments in compressive sensing. For in-distribution images, we show in Figures 17 and 18 qualitative examples of the image reconstructions for the Glow prior, the DCGAN prior, the IA-DCGAN prior, and the Deep Decoder. We replicate the same experiments for out-of-distribution FFHQ images in Figures 19 and 20.



Figure 17. Compressive sensing visual comparisons — Recoveries on a sample of in-distribution test set images with a number $m = 2500$ ($\approx 20\%$) of measurements under the Glow prior, the DCGAN prior, the IA-DCGAN prior, and the Deep Decoder.

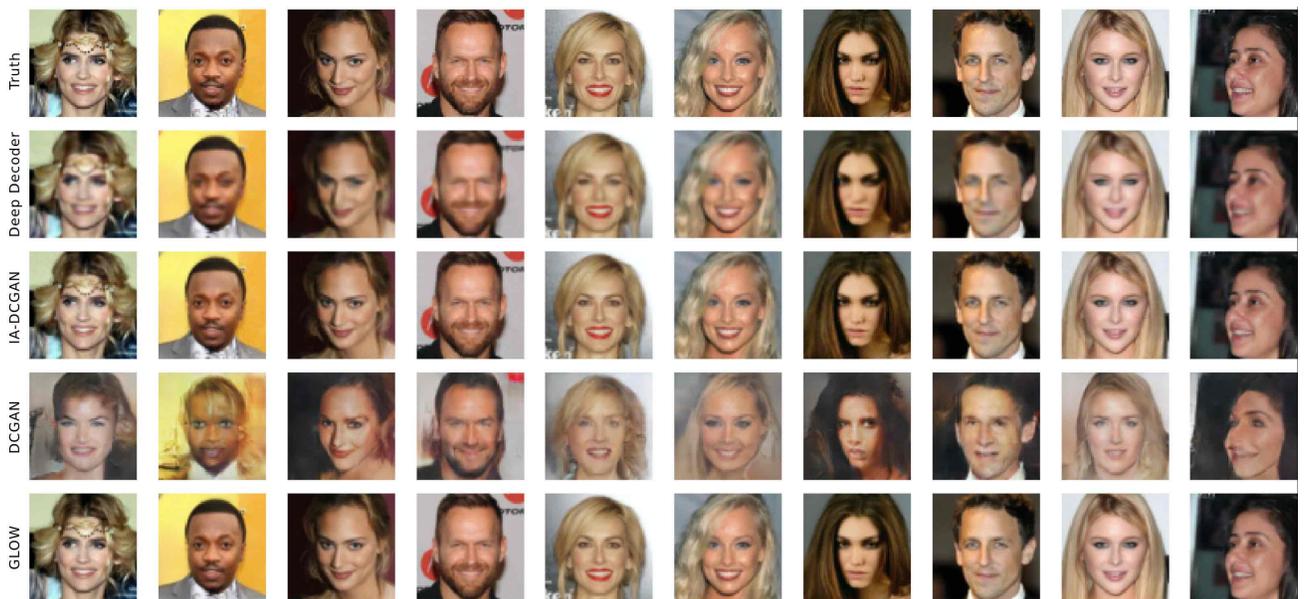


Figure 18. Compressive sensing visual comparisons — Recoveries on a sample of in-distribution test set images with a number $m = 10000$ ($\approx 80\%$) of measurements under the Glow prior, the DCGAN prior, the IA-DCGAN prior, and the Deep Decoder.

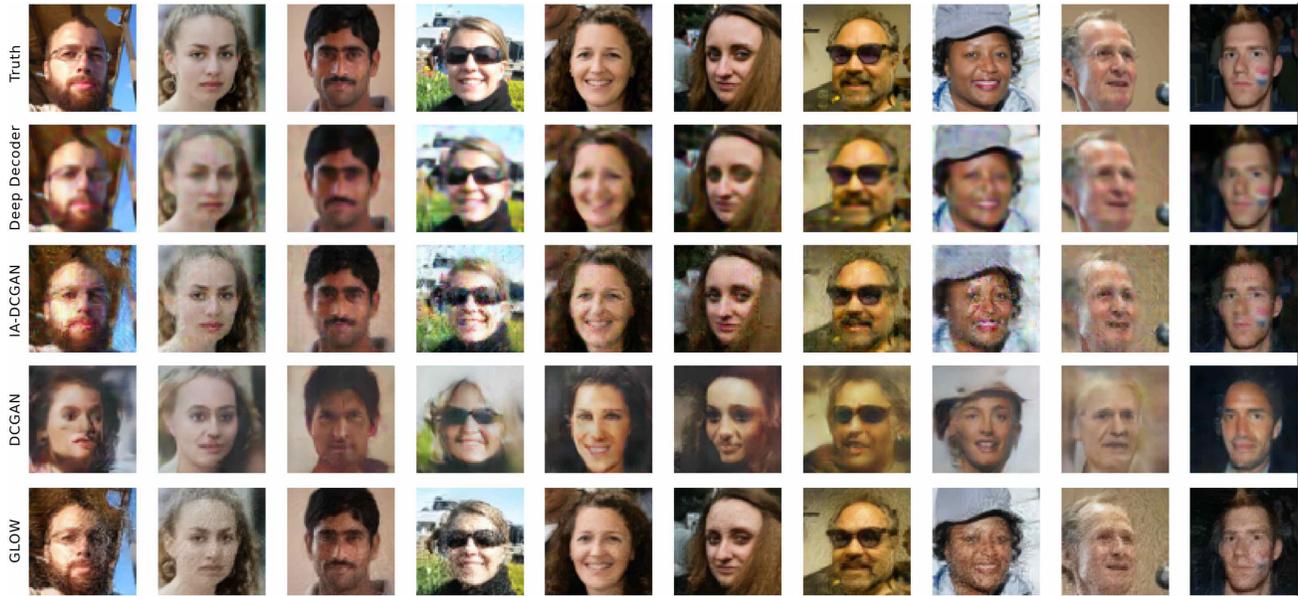


Figure 19. Compressive sensing visual comparisons — Recoveries on a sample of out-of-distribution (FFHQ) test set images with a number $m = 2500$ ($\approx 20\%$) of measurements under the Glow prior, the DCGAN prior, the IA-DCGAN prior, and the Deep Decoder.

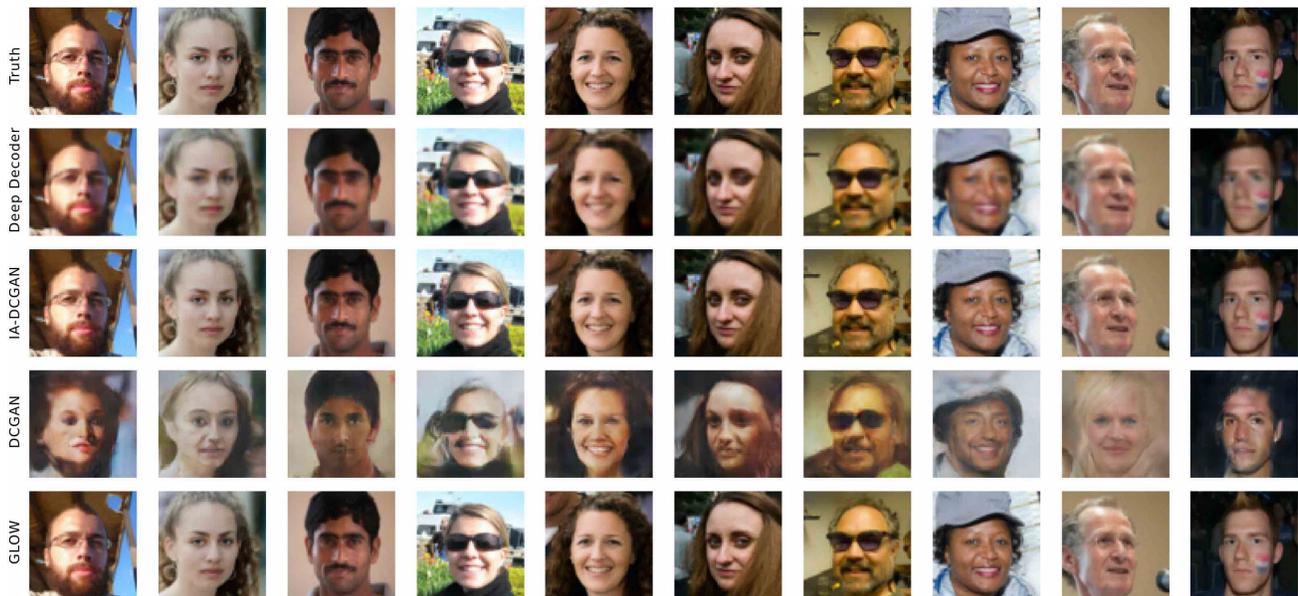


Figure 20. Compressive sensing visual comparisons — Recoveries on a sample of out-of-distribution (FFHQ) test set images with a number $m = 10000$ ($\approx 80\%$) of measurements under the Glow prior, the DCGAN prior, the IA-DCGAN prior, and the Deep Decoder.

To further investigate whether the Glow prior continues to be an effective proxy for arbitrarily out-of-distribution images, we tested arbitrary natural images such as a car, house door, and butterfly wings, which are all semantically unrelated to CelebA images. In general, we found that Glow is an effective prior at compressive sensing of out-of-distribution natural images, which are assigned a relatively high likelihood score (small normed latent representations) as compared to noisy, unnatural images. On these images, Glow also outperforms LASSO. This means that invertible networks have at least partially learned something more general about natural images from the CelebA dataset – there may be some high level features that face images share with other natural images, such as smooth regions followed by discontinuities, etc. This allows the Glow model to extend its effectiveness as a prior to other natural images beyond just the training set.

As compared to in-distribution training images, however, semantically unrelated images are assigned very low-likelihood scores by the Glow model, causing instability issues. In particular, an L-BFGS search for the solution of an inverse problem to recover a low-likelihood image leads the iterates into neighborhoods of low-likelihood representations that may induce instability. All the network parameters such as scaling in the coupling layers of Glow network are learned to behave stably with high likelihood representations. However, on very low-likelihood representations, unseen during the training process, the networks becomes unstable and outputs of network begin to diverge to very large values; this may be due to several reasons, such as normalization (scaling) layers not being tuned to the unseen representations. See Section 6.2.1 for details on our approach to handling these instabilities.

We show in Figure 21 a comparison of the performance of the LASSO-DCT, LASSO-WVT, DCGAN, and Glow priors on the compressive sensing of 64px out-of-distribution images for $m = 2500$ ($\approx 20\%$) measurements.

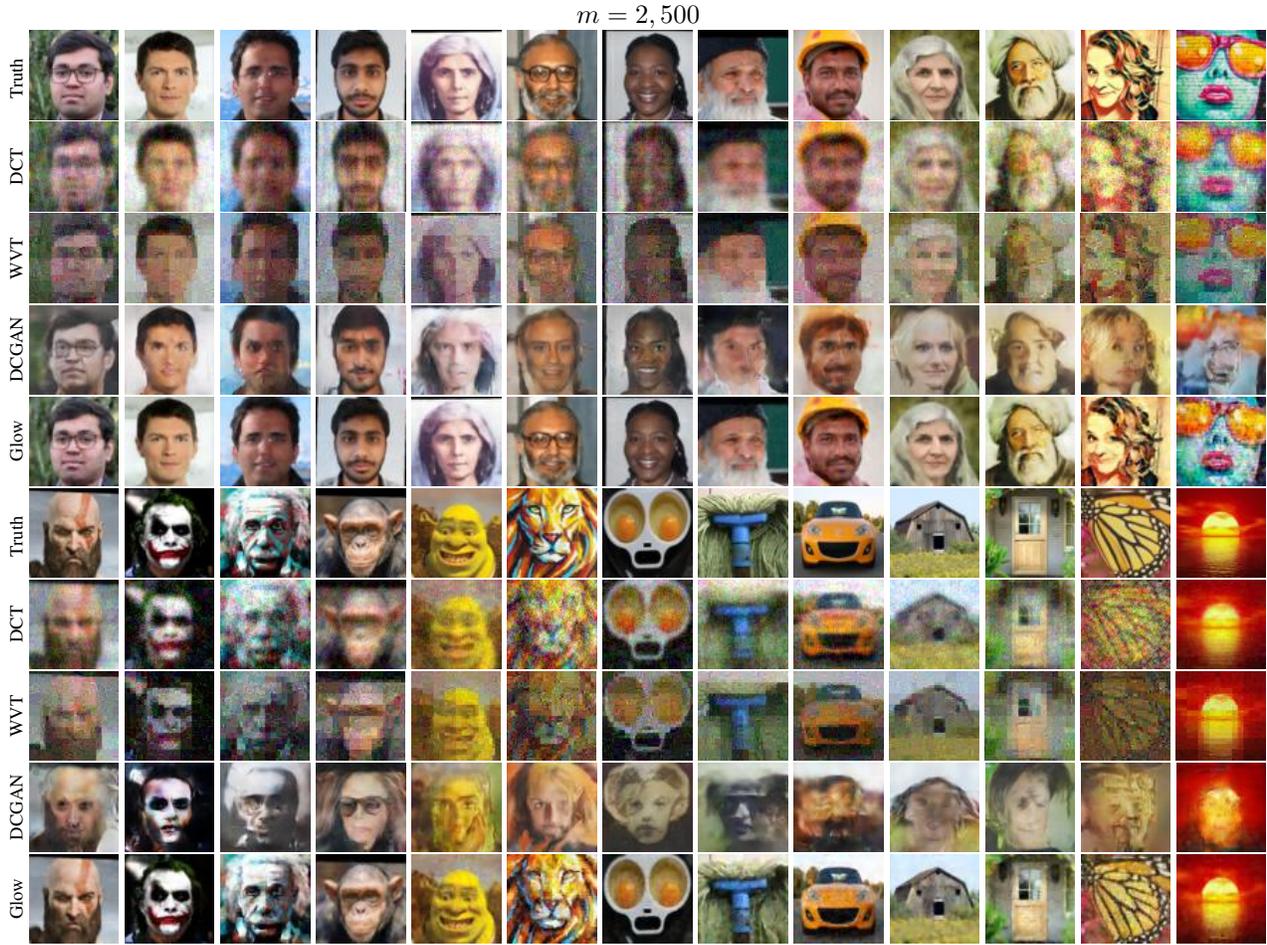


Figure 21. Compressive sensing of 64px out-of-distribution images given $m = 2500$ ($\approx 20\%$) measurements at a noise level $\sqrt{\mathbb{E}\|\eta\|^2} = 0.1$. We provide a visual comparison of recoveries of the LASSO-DCT, LASSO-WVT, DCGAN, and Glow priors, where the DCGAN and Glow priors are trained on CelebA images. In each case, we choose values of the penalization parameter γ to yield the best performance. We use $\gamma = 0$ for both the DCGAN and Glow priors, and optimize γ for each recovery using the LASSO-WVT and LASSO-DCT priors.

6.4.2. COMPRESSIVE SENSING 128PX SAMPLE SHEETS

We provide here additional sample sheets for our 128px experiments in compressive sensing. For in-distribution images, we show in Figures 22 and 23 qualitative examples of the image reconstructions for the Glow prior, the PGGAN prior, the IA-PGGAN prior, and the Deep Decoder. We replicate the same experiments for out-of-distribution (FFHQ) images in Figures 24 and 25.



Figure 22. Compressive sensing visual comparisons — Recoveries on a sample of in-distribution test set images with a number $m = 10000$ ($\approx 20\%$) of measurements under the Glow prior, the PGGAN prior, the IA-PGGAN prior, and the Deep Decoder.



Figure 23. Compressive sensing visual comparisons — Recoveries on a sample of in-distribution test set images with a number $m = 2500$ ($\approx 5\%$) of measurements under the Glow prior, the PGGAN prior, the IA-PGGAN prior, and the Deep Decoder.



Figure 24. Compressive sensing visual comparisons — Recoveries on a sample of out-of-distribution (FFHQ) test set images with a number $m = 10000$ ($\approx 20\%$) of measurements under the Glow prior, the PGGAN prior, the IA-PGGAN prior, and the Deep Decoder.



Figure 25. Compressive sensing visual comparisons — Recoveries on a sample of out-of-distribution (FFHQ) test set images with a number $m = 2500$ ($\approx 5\%$) of measurements under the Glow prior, the PGGAN prior, the IA-PGGAN prior, and the Deep Decoder.

6.5. Inpainting

We present here qualitative results on image inpainting under the DCGAN prior and the Glow prior on the CelebA test set. Compared to DCGAN, the reconstructions from Glow are of noticeably higher visual quality.



Figure 26. Image inpainting results on a sample of the CelebA test set. Masked images are recovered using the DCGAN prior and the Glow prior. Recoveries under the DCGAN prior are skewed and blurred whereas the Glow prior leads to sharper and coherent inpainted images. For both Glow and DCGAN, we set $\gamma = 0$.

6.5.1. IMAGE INPAINTING ON OUT OF DISTRIBUTION IMAGES

We now perform image inpainting under the Glow prior and the DCGAN prior, each trained on CelebA. Figure 27 shows the visuals of out-of-distribution inpainting. As before, the DCGAN prior continues to suffer due to representation error and data bias while Glow achieves reasonable reconstructions on out-of-distribution images which are semantically similar to CelebA faces. As one deviates to other natural images such as houses, doors, and butterfly wings, the inpainting performance deteriorates. In compressive sensing, Glow performs much better on such arbitrarily out-of-distribution images as in this case, good recoveries only require the network only to assign a higher likelihood score to the true image compared to the all the candidate static images given by the null space of the measurement operator.

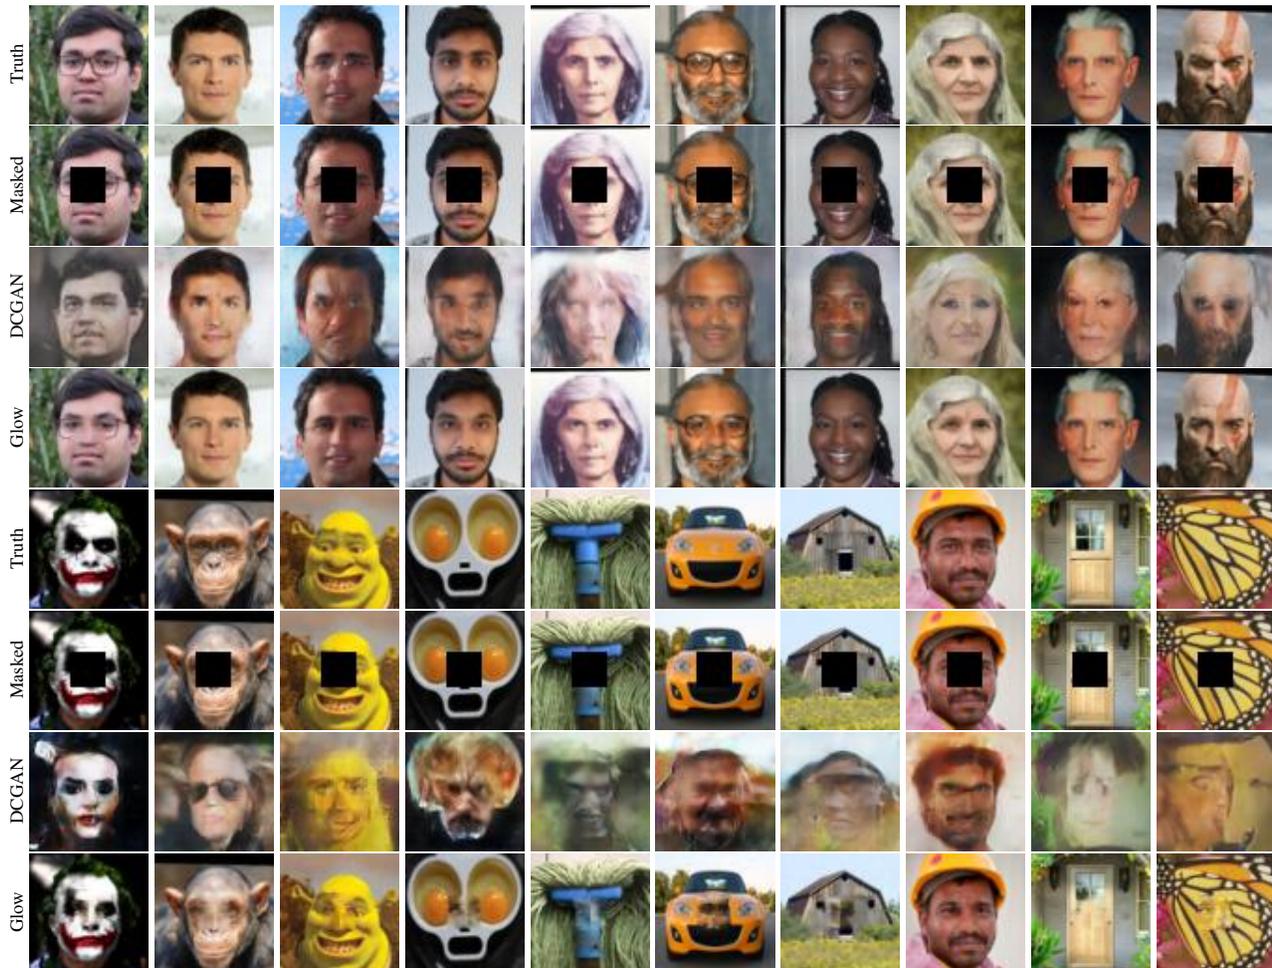


Figure 27. Image inpainting results on out-of-distribution images. Masked images are recovered under the DCGAN prior and the Glow prior. Recoveries under the DCGAN prior are skewed and blurred whereas Glow prior leads to sharper and coherent inpainted images. For both Glow and DCGAN, we set $\gamma = 0$.

6.6. Discussion

We provide evidence that random perturbations in image space induce larger changes in z than comparable natural perturbations in image space. Figure 28 shows a plot of the norm of the change in image space, averaged over 100 test images, as a function of the size of a perturbation in latent space. Natural directions are given by the interpolation between the latent representation of two test images. For the denoising problem, this difference in sensitivity indicates that the optimization algorithm might obtain a larger decrease in $\|z\|$ by an image modification that reduces unnatural image components than by a correspondingly large modification in a natural direction.

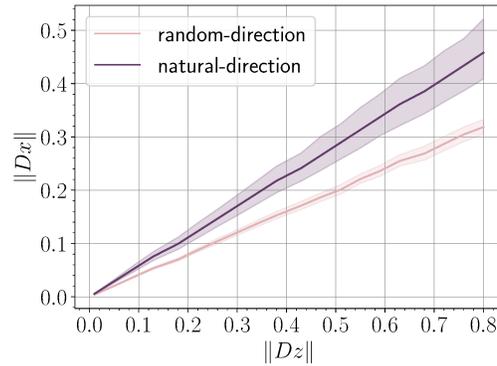


Figure 28. The magnitude of the change in image space as a function of the size of a perturbation in latent space. Solid lines are the mean behavior and shaded region depicts 95% confidence interval.

To further illustrate this point, we investigate in Figure 29 the decay of the singular values of the Glow model’s Jacobian for random points.

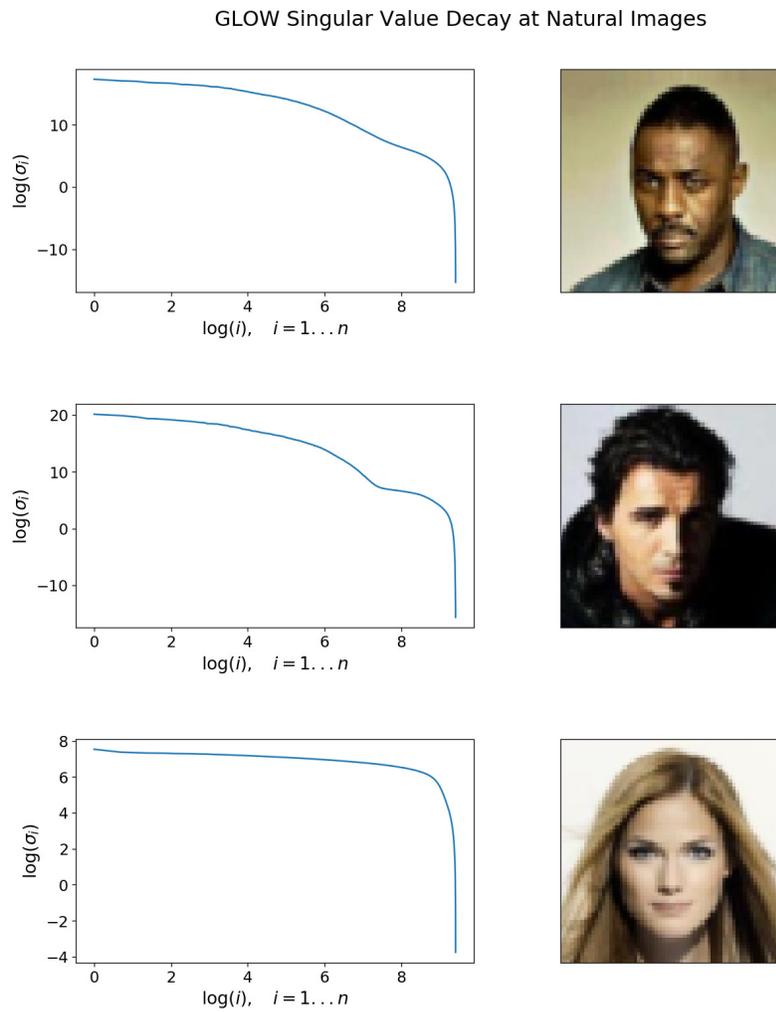


Figure 29. Log-log plots of the decay of the singular values of the trained Glow model's Jacobian for three random CelebA images.