
Prior Image-Constrained Reconstruction using Style-Based Generative Models

Varun A. Kelkar¹ Mark A. Anastasio¹

Abstract

Obtaining a useful estimate of an object from highly incomplete imaging measurements remains a holy grail of imaging science. Deep learning methods have shown promise in learning object priors or constraints to improve the conditioning of an ill-posed imaging inverse problem. In this study, a framework for estimating an object of interest that is semantically related to a known prior image, is proposed. An optimization problem is formulated in the disentangled latent space of a style-based generative model, and semantically meaningful constraints are imposed using the disentangled latent representation of the prior image. Stable recovery from incomplete measurements with the help of a prior image is theoretically analyzed. Numerical experiments demonstrating the superior performance of our approach as compared to related methods are presented.

1. Introduction

In recent years, generative models based on deep neural networks have risen to the forefront of machine learning research. By taking advantage of the approximately low dimensional structure of natural image data, generative models, such as generative adversarial networks (GANs) have been able to learn effective mappings from a simple, tractable, low dimensional distribution to a complex image data distribution, giving it the ability to generate highly realistic looking images, approximately from the distribution induced by the training data. Apart from image synthesis, generative models have found applications such as density estimation (Dinh et al., 2016), image restoration (Ledig et al., 2017), object detection (Prakash & Karam, 2019), video generation (Cai et al., 2018) and style transfer (Zhang et al., 2018) to name a few. Recent research in GANs has achieved state of the art performance in terms of visual

quality of images generated, invertibility and image representation, and a meaningful control of semantic features of the image (Karras et al., 2019; 2020).

Generative models have also found applications for image reconstruction in computed imaging systems, where a computational procedure is employed to form an estimate of an object of interest from imaging measurements. When the measurements are insufficient to uniquely determine the object of interest, this procedure amounts to solving an ill-posed inverse problem, and requires additional prior information about the object distribution. Approaches incorporating sparsity-based priors have been successful in obtaining accurate reconstructions from incomplete measurements (Candès & Wakin, 2008; Lustig et al., 2008). Seeking to better characterize the object distribution, generative models have been proposed as a prior for solving ill-posed inverse problems in imaging (Bora et al., 2017). Improvements in the image-synthesis performance, stability, and quality of generative neural networks have in-turn improved the performance of generative-model constrained reconstruction methods.

In this study, style-based generative models are investigated for constraining image reconstruction problems in which the solution is known to be close to a given prior image. This scenario, known as *prior image-constrained reconstruction*, is of particular significance when, for instance, the same object evolving over time is to be imaged multiple times. It finds applications in several scientific and medical imaging situations, such as monitoring tumor progression or perfusion (Chen et al., 2008), multi-contrast magnetic resonance imaging (MRI) (Weizman et al., 2016), or sequential radar imaging (Becquaert et al., 2019). Traditionally, this problem has been solved by assuming that the difference between the object of interest and the prior image is sparse in some domain (Chen et al., 2008; Mota et al., 2017), and solving an optimization problem penalizing this difference. However, although many natural images are compressible in some domain, their differences may not be so. Hence, penalizing the ℓ_1 norm of the difference between the object and the prior image in a linear transform domain may not be the best strategy to ensure that the ground truth is related to the prior image in a meaningful way. Style-based generative models have been known to be able to control individual semantic features, or styles, in an image by varying the disentangled

¹University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA. Correspondence to: Varun A. Kelkar <vak2@illinois.edu>, Mark A. Anastasio <maa@illinois.edu>.

latent representation of the image at different scales. In this work, the inverse problem is formulated as an optimization problem in the disentangled latent space of a style-based generative model. The disentangled latent representation of the prior image is computed and prior image-based regularization is imposed by constraining the estimate to have certain styles equal to the corresponding styles of the prior image.

Related work. Generative model-constrained reconstruction has been an active area of research, after being first proposed by (Bora et al., 2017). Several approaches have tried to reduce the representation error arising from the use of GANs as priors (Asim et al., 2020; Hussein et al., 2020). Other works have examined denoising and image reconstruction when the measurement noise is distributed in a complex way (Whang et al., 2020). Recently, StyleGAN has been used for image superresolution (Menon et al., 2020), and image reconstruction in general (Marinescu et al., 2020). Individual control over image semantics has been achieved (Karras et al., 2019; Tewari et al., 2020) by controlling the disentangled latent representation of the image. Since the disentangled latent representation of an image is crucial, several studies have focused on StyleGAN inversion (Wulff & Torralba, 2020; Abdal et al., 2019). Regularization by way of imposing constraints based on the latent space structure has also been explored (Kelkar et al., 2021). Prior image-constrained compressed sensing has been studied previously from the theoretical (Mota et al., 2017) point of view, to applications (Chen et al., 2008; Becquaert et al., 2019). Some studies have used adaptive weights on the prior image and the image estimate to better model the differences between the ground truth and the prior image (Weizman et al., 2016).

This study is organized as follows. Section 2 describes the background of compressed sensing, generative model-constrained image reconstruction, prior image-constrained reconstruction and style-based generative models. Section 3 describes the proposed approach motivated by StyleGANs. Section 4 describes a theoretical analysis of the problem at hand. Section 5 describes the setup of the proposed numerical studies, with the results being described in Section 6. Finally, discussion and conclusion is presented in Section 7.

2. Background

Several digital imaging systems can be approximately modeled by a linear imaging model, described as (Barrett & Myers, 2013)

$$\mathbf{g} = H\mathbf{f} + \mathbf{n}, \quad (1)$$

where $\mathbf{f} \in \mathbb{E}^n$ is a vector that approximates the object to-be-imaged, $\mathbf{g} \in \mathbb{E}^m$ corresponds to the imaging measurements, and $\mathbf{n} \in \mathbb{E}^m$ represents the measurement noise. Here, \mathbb{E}^l

corresponds to an $l \in \mathbb{N}$ dimensional Euclidean space. $H \in \mathbb{E}^{m \times n}$ corresponds to the linear operator that approximates the underlying physical model of the imaging system. Often, the measurements are incomplete ($m < n$) and, as such, \mathbf{f} cannot be uniquely recovered from \mathbf{g} . In this case, in order to obtain a useful estimate of the true object, prior knowledge about \mathbf{f} is needed to constrain the domain of H .

2.1. Compressed sensing

In recent decades, *compressed sensing* has emerged as a leading framework to solve such underdetermined systems of equations. It achieves this by constraining \mathbf{f} to a set of vectors that are sparse in some domain, and evoking certain conditions on H . Specifically, stable recovery of k -sparse objects can be guaranteed if H satisfies the *Restricted Isometry Property* (RIP) over the set of $2k$ -sparse signals (Candès et al., 2006; Candès & Wakin, 2008).

Definition 2.1 (Restricted Isometry). Let S_k be the set of all k -sparse vectors in \mathbb{R}^n . H is said to satisfy the RIP over S_k if $\exists \delta_k \in (0, 1)$ that satisfies

$$(1 - \delta_k) \|\mathbf{f}\|_2^2 \leq \|H\mathbf{f}\|_2^2 \leq (1 + \delta_k) \|\mathbf{f}\|_2^2, \quad (2)$$

for all $\mathbf{f} \in S_k$, and δ_k is not too close to 1 in a way prescribed in (Candès et al., 2006).

2.2. Prior image-constrained reconstruction

Prior image-constrained reconstruction is a scenario where the true object $\tilde{\mathbf{f}}$ is related, or close to, a previously known *prior image* $\mathbf{f}^{(PI)}$ (Chen et al., 2008; Mota et al., 2017). In traditional approaches to this problem, this similarity is imposed in the following way. In addition to the conventional sparsity constraint with respect to a transformation Φ , it is assumed that the difference of $\tilde{\mathbf{f}}$ and $\mathbf{f}^{(PI)}$ is sparse with respect to a transform Ψ , such as the wavelet transform or the 2D difference operator. Solving the inverse problem can then be cast as obtaining the solution to versions of the following optimization problem (Chen et al., 2008; Mota et al., 2017; Weizman et al., 2016):

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \|\mathbf{g} - H\mathbf{f}\|_2^2 + \lambda \left(\alpha \|\Psi(\mathbf{f} - \mathbf{f}^{(PI)})\|_1 + (1 - \alpha) \|\Phi\mathbf{f}\|_1 \right). \quad (3)$$

It is important to note that the assumption of sparse differences between the prior image and the ground truth may not always be valid. This is because in real life, images are better modeled as being compressible in a transform domain and are structured (Adcock et al., 2017). A difference of two such images may not only be denser than the original images, but also structured in a more complicated manner. Hence, it is important to develop new ways to quantify the similarity between the sought-after and the prior images.

2.3. Generative model-constrained reconstruction

Generative model-constrained reconstruction is a framework for image reconstruction, where the domain of H is constrained with the help of a generative model trained to approximate the distribution of objects (Bora et al., 2017; Asim et al., 2020; Marinescu et al., 2020). Let $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ be a generative model, typically parametrized by a deep neural network with parameters θ . G is trained on a dataset of images, such that under G , a sample $\mathbf{z} \in \mathbb{R}^k$ from a tractable distribution such as $\mathcal{N}(0, I_k)$ maps to a sample $G(\mathbf{z})$ that approximately comes from the distribution of the training dataset images. Here, \mathbf{z} is called the latent representation of $G(\mathbf{z})$. Since many real-life image datasets are approximately low dimensional, popular generative models, such as GANs often have a low dimensional domain with dimensionality $k \ll n$. Taking advantage of this fact, Bora *et al.* proposed a way to guarantee stable reconstruction of an object in the range of a generative model G having a Lipschitz constant \bar{L} from $O(k \log(\bar{L}r/\delta))$ measurements, if the object’s latent representation has an ℓ_2 norm of at most r , and if H satisfies the following *set-restricted eigenvalue condition* (S-REC):

Definition 2.2 (Set-restricted eigenvalue condition). Let $S \subseteq \mathbb{R}^n$. A matrix $H \in \mathbb{R}^{m \times n}$ satisfies the set-restricted eigenvalue condition S-REC(S, γ, δ) for some constants $\gamma > 0$ and $\delta \geq 0$, if for any $\mathbf{f}_1, \mathbf{f}_2 \in S$,

$$\|H(\mathbf{f}_1 - \mathbf{f}_2)\|_2 \geq \gamma \|\mathbf{f}_1 - \mathbf{f}_2\|_2 - \delta. \quad (4)$$

Intuitively, this property stipulates that two objects \mathbf{f}_1 and \mathbf{f}_2 in the range $\mathcal{R}(G)$ of G may give rise to measurements under H that are close, only if they themselves are close. Certain sensing matrices, such as iid Gaussian sensing matrices with an appropriate column length have been shown to satisfy the S-REC (Bora et al., 2017). The guarantees of stable recovery are applicable to the solution of the following constrained optimization problem:

$$\begin{aligned} \hat{\mathbf{z}} &= \arg \min_{\mathbf{z}, \|\mathbf{z}\| \leq r} \|\mathbf{g} - HG(\mathbf{z}; \theta)\|_2^2, \\ \hat{\mathbf{f}} &\equiv G(\hat{\mathbf{z}}; \theta), \end{aligned} \quad (5)$$

where $\mathbf{g} = H\tilde{\mathbf{f}} + \mathbf{n}$ is the measurement corresponding to the unknown true object $\tilde{\mathbf{f}}$. Since the above objective is non-convex, standard gradient descent-based methods are not guaranteed to converge to the optimal solution. However, it is observed that in practice, gradient-based methods give estimates of $\hat{\mathbf{z}}$ that are close to the optimum, at least in the case when $\tilde{\mathbf{f}} \in \mathcal{R}(G)$. Although Bora, *et al.* show numerical studies using a deep convolutional GAN (DCGAN), an optimization problem similar to Eq. (5) can also be formulated for recent advanced GAN architectures, such as StyleGAN, with demonstrably improved empirical performance (Marinescu et al., 2020).

2.4. Style-based Generative Adversarial Networks

StyleGAN and its successor, StyleGAN2, are well known for producing highly realistic samples from a real-life natural image distribution. They are characterized by an architecture, that consists of two sub-networks - (1) a *mapping network* $g_{\text{mapping}} : \mathbb{R}^k \rightarrow \mathbb{R}^k$, and an L -layer *synthesis network* $G : \mathbb{R}^{Lk} \rightarrow \mathbb{R}^n$. In the conventional image generation mode, the mapping network maps a sample $\mathbf{z} \in \mathcal{Z} \equiv \mathbb{R}^k$ from an iid standard normal distribution to a vector $\mathbf{u} \in \mathcal{W} \equiv \mathbb{R}^k$. The input \mathbf{w} to the L -layer synthesis network G is formed by stacking L copies of \mathbf{u} to form a $K = kL$ dimensional vector $\mathbf{w} \in \mathcal{W}^+ \equiv \mathbb{R}^K$. The i th copy of \mathbf{u} represents the input to the i th layer of G , which controls the i th level of detail in the generated image. In addition to these, G also takes as input a collection of latent-noise vectors η that control minor stochastic variations of the generated image at different resolutions. The ability of a StyleGAN to control features of the generated image at different scales comes in part due to this architecture, and in part, due to the style-mixing regularization during training (Karras et al., 2019). The latter loosely corresponds to evaluating the training loss using images generated by a “mixed” \mathbf{w} vector, formed by stacking the \mathcal{W} -space outputs of different realizations of \mathbf{z} . In addition to these basic characteristics, StyleGAN2 introduces path-length regularization, which aids in better conditioning of G and reducing the representation error (Karras et al., 2020). For conventional image generation, a sampled \mathbf{w} vector is degenerate, containing L copies of \mathbf{u} , and hence lies in a k -dimensional subspace of \mathcal{W}^+ . However, studies have shown that from the point-of-view of projecting an image to the range of G , utilizing the entire \mathcal{W}^+ space has benefits in terms of lower representation error (Wulff & Torralba, 2020).

3. Approach

3.1. Prior image-constrained reconstruction using StyleGANs (PICGM)

StyleGAN and StyleGAN2 are able to vary certain styles of an image while keeping certain other styles fixed. For example, for the StyleGAN trained on a dataset of faces, it is possible to vary the hairstyle and hair color while keeping the general structure of the face the same. For a medical imaging dataset such as a dataset of multi-contrast brain MRI images, it is possible to control the contrast, or exact placement of the folds, ventricles, and other fine scale features while keeping the general structure of the image the same. Hence, comparing the sought-after and the prior images in the latent space of a StyleGAN is a natural approach to quantify the similarity between the two images.

Motivated by such style-mixing properties of the StyleGAN, one possible way to formulate the prior image-constrained

inverse problem is as follows. Let $G : \mathbb{R}^K \rightarrow \mathbb{R}^n$ denote the synthesis network of a trained StyleGAN2. Note that the domain of G is taken to be the extended space \mathcal{W}^+ . Let $\mathbf{f}^{(\text{PI})} = G(\mathbf{w}^{(\text{PI})})$ denote a known prior image in the range of G . Then, the proposed measurement model can be written as

$$\mathbf{g} = H\tilde{\mathbf{f}} + \mathbf{n}, \quad \tilde{\mathbf{f}} \in \{G(\mathbf{w}) \text{ s.t. } \mathbf{w}_{1:p_1} = \mathbf{w}_{1:p_1}^{(\text{PI})}, \mathbf{w}_{p_2:K} = \mathbf{w}_{p_2:K}^{(\text{PI})}\}, \quad (6)$$

where p_1, p_2 are multiples of k , $1 \leq p_1 < p_2 \leq K$, $\tilde{\mathbf{f}}$ is the sought-after ground truth and $\mathbf{w}_{u:v}$ denotes the section of vector \mathbf{w} from indices u through v .

Even assuming that $\mathbf{f}^{(\text{PI})}$ and $\tilde{\mathbf{f}}$ are in-distribution images, there are some practical concerns about the measurement model described above. (1) In practice, $\mathbf{f}^{(\text{PI})}$ may not lie in $\mathcal{R}(G)$ since $\mathcal{R}(G)$ is a K -dimensional manifold in \mathbb{R}^n . (2) If $\mathbf{f}^{(\text{PI})} \in \mathcal{R}(G)$, its disentangled latent representation $\mathbf{w}^{(\text{PI})}$ may lie in an unstable region of \mathcal{W}^+ and the style mixing properties of G may not apply to it. (3) If $\mathbf{f}^{(\text{PI})} \in \mathcal{R}(G)$ and the style mixing performance using $\mathbf{w}^{(\text{PI})}$ is consistent with that of images drawn from G , $\tilde{\mathbf{f}}$, i.e. the sought-after object, may still not lie in the set $\{G(\mathbf{w}) \text{ s.t. } \mathbf{w}_{1:p_1} = \mathbf{w}_{1:p_1}^{(\text{PI})}, \mathbf{w}_{p_2:K} = \mathbf{w}_{p_2:K}^{(\text{PI})}\}$ for any p_1, p_2 . The following solutions are proposed in order to alleviate the aforementioned concerns. Concern (1) essentially refers to the representation error when approximating $\mathbf{f}^{(\text{PI})}$. It was observed that for in-distribution images, if a latent representation $\mathbf{w}^{(\text{PI})}$ in the extended \mathcal{W}^+ space is sought, even if the latent-noise vectors η are not optimized over, a close approximation to $\mathbf{f}^{(\text{PI})}$ can be obtained by gradient-descent based optimization, except for minor stochastic detail represented by η .

In an attempt to resolve concern (2), (Wulff & Torralba, 2020) observed that a transformed version of \mathbf{w} , given by $\mathbf{v} = \text{LReLU}_\alpha(\mathbf{w})$ approximately follows a multivariate gaussian distribution with a mean $\bar{\mathbf{v}} \in \mathbb{R}^K$ and covariance $\Sigma \in \mathbb{R}^{K \times K}$. Here, $\text{LReLU}_\alpha(\cdot)$ denotes the leaky-ReLU non-linear activation (Maas et al., 2013), defined as

$$\text{LReLU}_\alpha(\mathbf{x})_i = \begin{cases} x_i, & x_i \geq 0, \\ \alpha x_i, & x_i < 0. \end{cases} \quad (7)$$

The value of α is the reciprocal of the scaling value for negative numbers included in the last leakyReLU layer in the mapping network g_{mapping} . This means that it is possible to regularize the inversion of G with the help of a Gaussian prior on \mathbf{v} (Wulff & Torralba, 2020). The inversion process can then be formulated in terms of the following optimization problem (Wulff & Torralba, 2020)

$$\begin{aligned} \mathbf{w}^{(\text{PI})} &= \arg \min_{\mathbf{w}} \left\| \mathbf{f}^{(\text{PI})} - G(\mathbf{w}) \right\|_2^2 + \lambda \|\mathbf{v} - \bar{\mathbf{v}}\|_\Sigma^2, \\ \text{s.t. } \mathbf{v} &= \text{LReLU}_\alpha(\mathbf{w}), \end{aligned} \quad (8)$$

Algorithm 1 Projected Adam algorithm for minimizing the objective in Eq. (9).

Input: Measurements \mathbf{g} , prior image latent $\mathbf{w}^{(\text{PI})}$, Regularization parameters p_1, p_2, λ , maximum iterations n_{iter} . $\mathcal{L}(\mathbf{w}; \lambda)$: Objective function from Eq. (9).

Initialize Adam optimizer parameters $(\alpha, \beta_1, \beta_2)$. (Default parameters were used).

Initialize iteration number $t \leftarrow 0$.

Initialize $\mathbf{w}^{[0]} \leftarrow \mathbf{w}^{(\text{PI})}$.

while $\mathbf{w}^{[t]}$ **not** converged **do**

Adam update (Kingma & Ba, 2014):

$$\mathbf{w}^{[t]} \leftarrow \text{ADAM}_{\alpha, \beta_1, \beta_2}(\mathcal{L}(\mathbf{w}^{[t]}; \lambda))$$

Projection step:

$$\begin{aligned} \mathbf{w}_{1:p_1}^{[t]} &\leftarrow \mathbf{w}_{1:p_1}^{(\text{PI})} \\ \mathbf{w}_{p_2:K}^{[t]} &\leftarrow \mathbf{w}_{p_2:K}^{(\text{PI})} \end{aligned}$$

$t \leftarrow t + 1$

end while

where $\|\mathbf{x}\|_\Sigma^2 = \mathbf{x}^\top \Sigma^{-1} \mathbf{x}$ is used to impose a prior on \mathbf{v} corresponding to a Gaussian distribution with mean $\bar{\mathbf{v}}$ and covariance Σ . As observed in (Wulff & Torralba, 2020), $\mathcal{R}(G)$ -projected estimates of $\mathbf{f}^{(\text{PI})}$ obtained in this way inherit the style-mixing and stability properties of samples from G . The tradeoff between an accurate representation of $\mathbf{f}^{(\text{PI})}$ and the style-mixing properties is governed by the regularization parameter λ .

Concern (3) can be addressed by arguing that since $\tilde{\mathbf{f}}$ is an in-distribution image, it has minimal representation error when optimizing over \mathcal{W}^+ . Also, p_1, p_2 can be treated as tunable regularization parameters, which manage the trade-off between imposition of the prior from $\mathbf{f}^{(\text{PI})}$, and consistency with the measurements \mathbf{g} .

Taking into account the above arguments, the inverse problem associated with Eq. (6) is formulated as the following optimization problem:

$$\begin{aligned} \hat{\mathbf{w}} &= \arg \min_{\mathbf{w}} \|\mathbf{g} - HG(\mathbf{w})\|_2^2 + \lambda \|\mathbf{v} - \bar{\mathbf{v}}\|_\Sigma^2, \\ \text{s.t. } \mathbf{w}_{1:p_1} &= \mathbf{w}_{1:p_1}^{(\text{PI})}, \quad \mathbf{w}_{p_2:K} = \mathbf{w}_{p_2:K}^{(\text{PI})}, \\ \mathbf{v} &= \text{LReLU}_\alpha(\mathbf{w}). \end{aligned} \quad (9)$$

Although the above problem is non-convex, similar to previous works (Bora et al., 2017; Asim et al., 2020), useful estimates $\hat{\mathbf{f}}$ can be obtained by iterative gradient descent-based optimization. Algorithm 1 shows the projected-Adam algorithm used for this purpose (Kingma & Ba, 2014).

3.2. Compressed sensing using StyleGAN2 (CSGM)

Analogous to Eq. (9), CSGM using StyleGAN2 can be formulated in the following manner:

$$\begin{aligned} \hat{\mathbf{w}} &= \arg \min_{\mathbf{w}} \|\mathbf{g} - HG(\mathbf{w})\|_2^2 + \lambda \|\mathbf{v} - \bar{\mathbf{v}}\|_{\Sigma}^2, \\ \text{s.t. } \mathbf{v} &= \text{LReL}_{\alpha}(\mathbf{w}). \end{aligned} \quad (10)$$

The Adam algorithm was used for obtaining approximate solutions to the above problem (Kingma & Ba, 2014).

4. Theoretical Analysis

The theoretical analysis presented here is motivated by the results presented in (Bora et al., 2017). There, the authors provide a stable recovery guarantee in terms of the Lipschitz constant of the generative network. However, in practice, Lipschitz constants can be difficult to estimate, or a generative network may not be Lipschitz stable. Hence, theoretical analysis in terms of the Jacobians of the generative network is presented here, in order to derive a limited, non-uniform guarantee for the stable recovery of typical in-distribution objects that lie in the range of the generative network. In order to do so, we utilize properties of StyleGAN2.

Let $p_{\mathbf{w}}$ denote the distribution of the extended latent space vector $\mathbf{w} = [\mathbf{u}_1^{\top} \ \mathbf{u}_2^{\top} \ \dots \ \mathbf{u}_L^{\top}]^{\top} \in \mathcal{W}^+$, where $\mathbf{u}_i \sim g_{\text{mapping}}(\mathbf{z}_i)$, $\mathbf{z}_i \sim \mathcal{N}(0, I_k)$, \mathbf{z}_i 's are independently distributed. The following assumptions are made based on the StyleGAN2 properties:

1) *Path length regularity:*

$$\mathbb{E}_{\mathbf{w} \sim p_{\mathbf{w}}} (\|J(\mathbf{w})\|_F - a)^2 < b, \quad (\text{AS1})$$

$J(\mathbf{w})$ denotes the Jacobian of G evaluated at \mathbf{w} , $b > 0$ and $a = \mathbb{E}_{\mathbf{w}} \|J(\mathbf{w})\|_F$ are global constants. This assumption is inspired by the path-length regularization used in (Karras et al., 2020).

2) *Approximate local linearity:*

$$\mathbb{E}_{\mathbf{w} \sim p_{\mathbf{w}}} \max_{\|\mathbf{w}' - \mathbf{w}\| \leq \epsilon} \mathcal{L}(\mathbf{w}', \mathbf{w}) \leq \beta^2(\epsilon), \quad (\text{AS2})$$

where

$$\mathcal{L}(\mathbf{w}', \mathbf{w}) = \|G(\mathbf{w}') - G(\mathbf{w}) - J(\mathbf{w})(\mathbf{w}' - \mathbf{w})\|_2^2,$$

and β is a positive function with $\beta(0) = 0$. This property essentially measures how close to linear G behaves in an ϵ -neighborhood around a point \mathbf{w} .

As described in (Karras et al., 2020), estimates of a, b can be computed via an empirical estimation of $\mathbb{E}_{\mathbf{w}} \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I_n)} \|J(\mathbf{w})^{\top} \mathbf{y}\|_2$. Estimating $\beta(\epsilon)$ for a given

value of ϵ is however not tractable. Nevertheless, an approximate estimation can be obtained by first computing the Jacobian at a point $\mathbf{w} \sim p_{\mathbf{w}}$, and then iteratively maximizing $\mathcal{L}(\mathbf{w}', \mathbf{w})$ using a projected gradient ascent-type algorithm followed by an empirical mean of the maxima for several such $\mathbf{w} \sim p_{\mathbf{w}}$. Empirically estimated values of a, b and β are presented in the supplementary file.

Using these two assumptions, the following results are derived.

Notation 4.1. Let $p_{\mathbf{v}}$ denote the distribution of $\mathbf{v} = \text{LReL}_{\alpha}(\mathbf{w})$, and let $\bar{\mathbf{v}}, \Sigma$ be its mean and covariance matrix respectively. Let $\mathbf{w}^{(\text{PI})}$ be a sample from $p_{\mathbf{w}}$, and $1 \leq p_1 < p_2 \leq K$. Assume that p_1 and p_2 are multiples of k . Let

$$\begin{aligned} B_{\mathbf{w}}^{p_1, p_2}(r) &:= \left\{ \mathbf{w} \text{ s.t. } \|\text{LReL}_{\alpha}(\mathbf{w}) - \bar{\mathbf{v}}\|_{\Sigma} \leq r, \right. \\ &\quad \left. \mathbf{w}_{1:p_1} = \mathbf{w}_{1:p_1}^{(\text{PI})}, \mathbf{w}_{p_2:K} = \mathbf{w}_{p_2:K}^{(\text{PI})} \right\}, \end{aligned}$$

By Markov's inequality and concentration of norm, we have the following (Vershynin, 2018):

Lemma 4.1. If \mathbf{w} is a sample from $p_{\mathbf{w}}$, then it satisfies the following three properties with probability at least $1 - O(1/K)$:

$$\|J(\mathbf{w})\|_F \leq \sqrt{K}a, \quad (\text{P1})$$

$$\max_{\|\mathbf{w}' - \mathbf{w}\| \leq \epsilon} \mathcal{L}(\mathbf{w}', \mathbf{w}) \leq \sqrt{K}\beta(\epsilon) \quad (\text{P2})$$

$$\left\| \Sigma^{-1/2} \text{LReL}_{\alpha}(\mathbf{w}) \right\|_2 \leq \sqrt{K}(1 + o(1)) \quad (\text{P3})$$

Lemma 4.2 (Set-restricted eigenvalue condition). Let $\mathbf{w}^{(\text{PI})}$ be a sample from $p_{\mathbf{w}}$. Assume that $p_{\mathbf{v}}$ is a Gaussian distribution with covariance matrix Σ . Let $\tilde{B}_{\mathbf{w}}^{p_1, p_2}(r)$ be the set of all points in $B_{\mathbf{w}}^{p_1, p_2}(r)$ satisfying properties P1 and P2. Let $\tau < 1$, $\delta > 0$. Let $H \in \mathbb{R}^{m \times n}$ be an matrix with elements $h_{ij} \sim \mathcal{N}(0, 1/m)$. For all $\delta' < \delta$, let $\beta(\delta'/a)$ go polynomially as δ'/a , with $\beta(0) = 0$. If

$$m = \Omega \left(\frac{p_2 - p_1}{\tau^2} \log \frac{ar \|\Sigma\|_F}{\delta} \right), \quad (11)$$

then H satisfies the $S\text{-REC}(G(\tilde{B}_{\mathbf{w}}^{p_1, p_2}(r)), 1 - \tau, \delta + \sqrt{K}\beta(\delta/a))$ with probability $1 - e^{-\Omega(\alpha^2 m)}$.

Based on this, the following reconstruction guarantee is arrived at.

Theorem 4.1. Let $H \in \mathbb{R}^{m \times n}$ satisfy $S\text{-REC}(G(\tilde{B}_{\mathbf{w}}^{p_1, p_2}(r)), \gamma, \delta + \sqrt{K}\beta(\delta/a))$. Let \mathbf{n} be the measurement noise. Let $\mathbf{w}, \mathbf{w}^{(\text{PI})} \sim p_{\mathbf{w}}$. Let $\mathbf{f}^{(\text{PI})} = G(\mathbf{w}^{(\text{PI})})$ be the known prior image. Let

$$\tilde{\mathbf{w}} = [\mathbf{w}_{1:p_1}^{(\text{PI})\top} \quad \mathbf{w}_{p_1:p_2}^{\top} \quad \mathbf{w}_{p_2:K}^{(\text{PI})\top}]^{\top}.$$

Let $\tilde{\mathbf{f}} = G(\tilde{\mathbf{w}})$ represent the object to-be-imaged. Let $\mathbf{g} = H\tilde{\mathbf{f}} + \mathbf{n}$ be the imaging measurements. Let

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f} \in G(\tilde{B}_{\tilde{\mathbf{w}}}^{p_1, p_2}(r))} \|\mathbf{g} - H\mathbf{f}\|_2^2. \quad (12)$$

Then,

$$\|\hat{\mathbf{f}} - \tilde{\mathbf{f}}\| \leq \frac{1}{\gamma}(2\|\mathbf{n}\| + \delta + \sqrt{K}\beta(\delta/a)) \quad (13)$$

with probability $1 - O(1/K)$.

Note that instead of optimizing over $\mathbf{f} \in G(\tilde{B}_{\tilde{\mathbf{w}}}^{p_1, p_2}(r))$, Eq. (9) proposes to solve a relaxed version, over $G(B_{\tilde{\mathbf{w}}}^{p_1, p_2}(r))$. In addition to this, as mentioned previously, Eq. (9) is non-convex and convergence is not guaranteed. Due to this, the proposed approach may not always yield an optimum to Eq. (12). However, whether or not the solution lies in $G(\tilde{B}_{\tilde{\mathbf{w}}}^{p_1, p_2}(r))$ can always be checked by checking if P1, P2, P3 are satisfied. For estimating in-distribution objects in the range of G , it was observed empirically that the conditions P1, P2, P3 are satisfied by the estimated $\hat{\mathbf{w}}$.




Figure 1: Ground truth, prior image and image estimated from Gaussian measurements with $n/m = 10$ and $n/m = 50$ using the proposed approach in the inverse crime case.




Figure 2: Ensemble RMSE and SSIM values for in the inverse-crime setting, for various subsampling ratios. The error bars show a span of one standard deviation.

5. Numerical Studies

The numerical studies were split into three parts - (1) inverse-crime study, where the object was directly sampled from the StyleGAN2 and measurements were simulated using a Gaussian forward model, (2) face image study, where real face images were used to simulate noisy measurements using a Gaussian forward model, and (3) MR image study, where real brain MR images were used to simulate stylized undersampled MRI measurements with noise. Additionally, the robustness of the proposed approach to misalignment between the ground truth and the prior image was numerically assessed in the MR image study.

1) Dataset and forward model:

For the inverse-crime study, StyleGAN2 was trained on a composite brain MR image dataset consisting of a total of 200676 T1 and T2 weighted images of size 256×256 from the fastMRI initiative database (Zbontar et al., 2018) and 866 images from the brain tumor progression dataset (Schmainda & Prah, 2018). These images were resized to 256×256 . For the face image study, a StyleGAN2 with an output image size of $128 \times 128 \times 3$ was trained on images from the Flickr-Faces-HQ (FFHQ) dataset (Karras et al., 2019). For the MR image study, StyleGAN2 was trained on a composite brain MR image dataset consisting of 164741 T1 and T2 weighted images from the fastMRI database, 686 images from the Brain tumor progression dataset, 2206 T1 and T2 weighted images from the TCIA-GBM dataset (Scarpace et al., 2016), and 36978 T2 weighted images from the OASIS-3 dataset (LaMontagne et al., 2019).

For evaluating reconstruction performance in the inverse crime studies, iid images were sampled from the generative model, where each style vector \mathbf{u}_i was sampled independently from the mapping network, and the composite vector $\mathbf{w} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_L]$ was used to sample the image. This was treated as the prior image. Four of the style vectors, $\mathbf{u}_6 \dots \mathbf{u}_9$ were then replaced by new styles $\mathbf{u}'_6 \dots \mathbf{u}'_9$, which were each again sampled independently using the mapping network. This was treated as the ground truth object to-be-recovered. Reconstruction performance was evaluated on a dataset of 50 such images. An iid Gaussian matrix $H \in \mathbb{R}^{m \times n}$ was used as the forward model and real-valued iid Gaussian noise with signal-to-noise ratio (SNR) 20 dB was added to the measurements. The reconstruction performance was evaluated for undersampling ratios $n/m = 5, 10, 20$ and 50.

For evaluating reconstruction performance in the face image study, 17 pairs of face images corresponding to differences in hairstyles, hair color and facial expressions were obtained from the stock-image hosting service Shutterstock (Shutterstock, 2021). Images were manually inspected to ensure that one of the images in the pair was not a digitally altered

version of the other. The images were manually cropped and resized to $128 \times 128 \times 3$. In the numerical experiments, one of the images in the pair was used as the prior image. An iid Gaussian matrix with $n/m = 50$ was used to generate the measurements from the second image, which was treated as the ground truth. Real-valued iid Gaussian noise with 20 dB SNR was added to the measurements.

The brain tumor progression dataset contains pairs of brain MR images of test subjects, separated by time on the scale of a few months. The second image in the pair shows progression of the brain tumor with respect to the first image. For evaluating the reconstruction performance in the case of the MR image study, a dataset of 22 held-out image pairs from the tumor progression dataset were used. The first image was used as the prior image. Imaging measurements were simulated from the second image, which was treated as the ground truth. A Fourier undersampling forward model, described as $H = \mathbf{m} \odot \mathcal{F} \in \mathbb{R}^{m \times n}$ was used to simulate MRI measurements. Here, \mathbf{m} corresponds to a binary mask, and \mathcal{F} corresponds to the 2D discrete Fourier transform. This forward operator fully samples a fraction of the lower frequencies and randomly subsamples a fraction of high frequencies of the image. Five different undersampling ratios are considered: $n/m = 2, 4, 6, 8$ and 12 . Complex iid Gaussian noise with 20 dB SNR was added to the measurements.

2) *Generative network training details:* The StyleGAN2 architecture proposed in (Karras et al., 2020) was used. For an image size of $2^i \times 2^i$, it contains $2(i - 1)$ layers split across $i - 1$ resolution levels. The default latent space dimensionality of 512 was maintained. The networks were trained using Tensorflow 1.14/Python (Abadi et al., 2015) on an Intel Xeon E5-2620v4 CPU @ 2.1 GHz and four Nvidia TITAN X graphics processing units (GPUs).

3) *Baselines:* For obtaining estimates of the true object from \mathbf{g} , the performance of the following reconstruction methods were qualitatively and quantitatively compared – (1) Penalized least squares with TV regularization (PLS-TV), (2) compressed sensing using StyleGAN2 (CSGM) mentioned in Eq. (10), (3) prior image-constrained compressed sensing (PICCS) mentioned in Eq. (3), and (4) the proposed method (PICGM) introduced in Eq. (9). Note that the first two methods described do not utilize information from the prior image, while the last two do. For PICCS, discrete difference operator (corresponding to TV semi-norm) was used as the sparsifying transform Φ , and for the transform Ψ , a 2D Haar wavelet transform of level 7 (i.e. equal to the number of resolution levels in StyleGAN2) was utilized. For evaluating PICGM, the latent representation $\mathbf{w}^{(PI)}$ of the prior image was computed with the help of the trained StyleGAN using the procedure introduced in Eq. (8). The regularization parameters for all the methods were tuned using either a line search or a grid search depending upon

the number of regularization parameters, and the setting giving the lowest ensemble root mean-squared error (RMSE) was chosen. The image estimates obtained by use of each of the four methods were quantitatively evaluated using the RMSE and structural similarity (SSIM). All algorithms were implemented using Tensorflow 1.14/Python.¹

4) *Testing the robustness to misalignments:* The robustness of PICGM as compared to PICCS was numerically analyzed for the MR image study, by considering configurations where the ground truth is misaligned with the prior image. Five such configurations were considered. For each configuration C_i , the test ground truth images were each first rotated by $2i$ degrees clockwise or anticlockwise randomly with equal probability, and then translated by $2i$ pixels in a uniformly random direction. Estimates of the new ground truths were obtained with the help of PICCS and PICGM from 8 fold undersampled Fourier measurements, and the ensemble RMSE and SSIM performances of the two algorithms were compared.

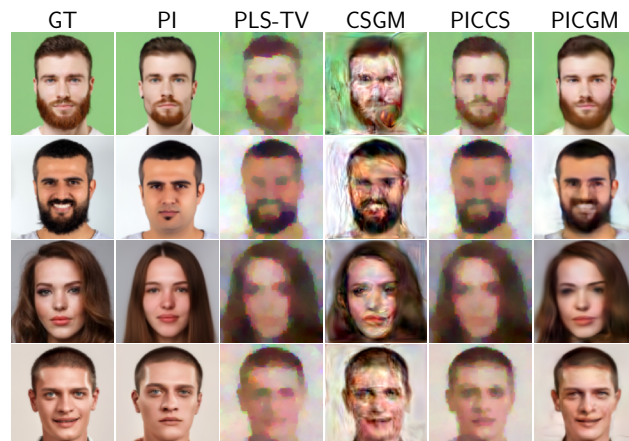


Figure 3: Ground truths (GT), prior images (PI) and image estimated from Gaussian measurements with $n/m = 50$ using the proposed approach in the face image study (zoom in for clarity).

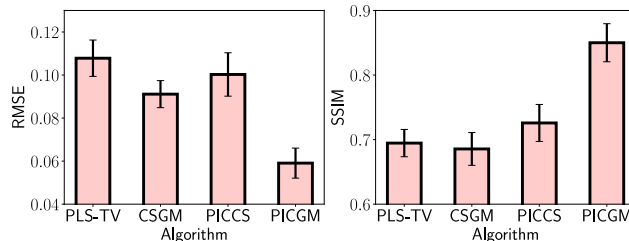


Figure 4: Ensemble RMSE and SSIM values for in the face image study for $n/m = 50$. The error bars span one standard deviation.

¹The Tensorflow/python implementation of the reconstruction methods can be found at <https://github.com/comp-imaging-sci/pic-recon>

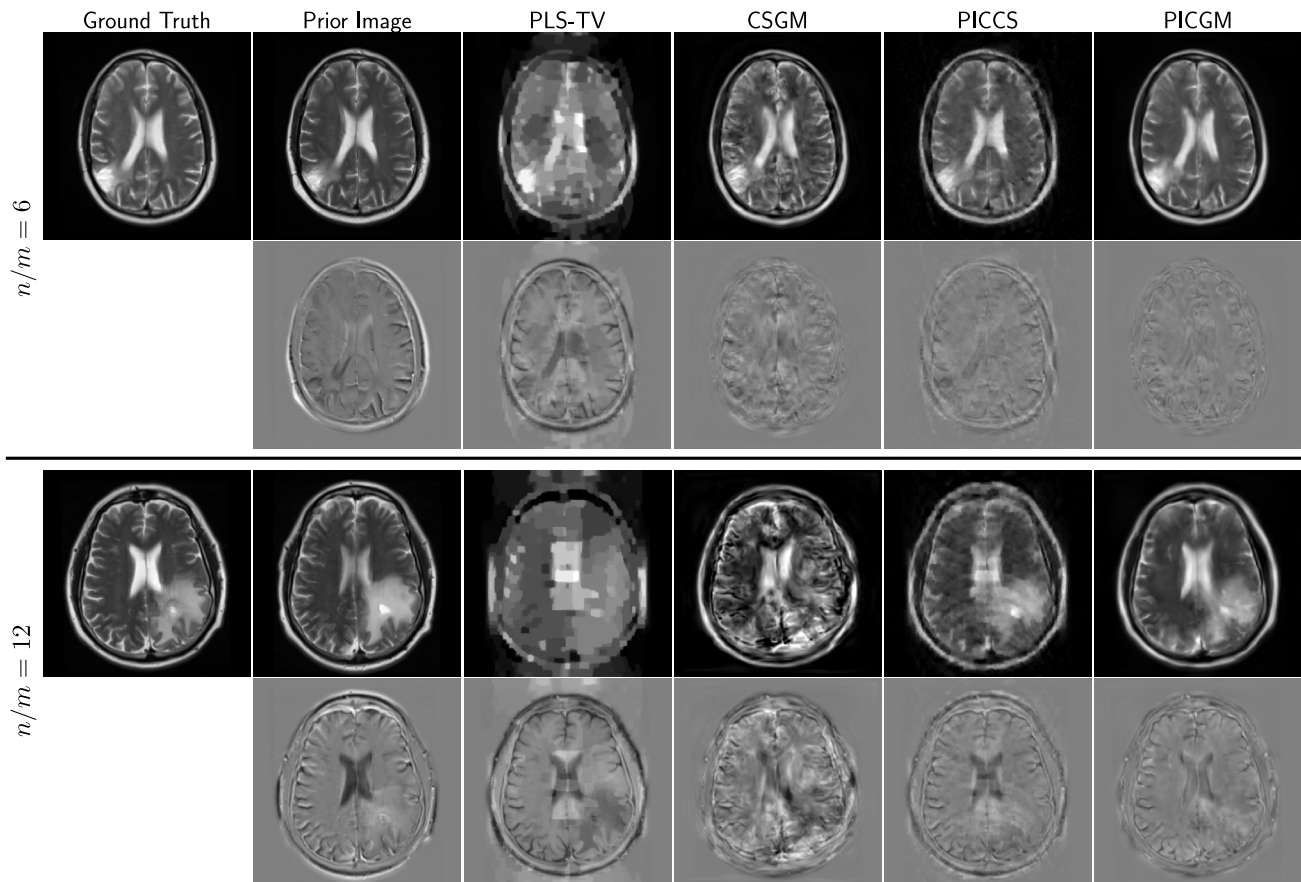


Figure 5: Ground truth, prior image, and images reconstructed from simulated MRI measurements with $n/m = 6$ and $n/m = 12$ along with difference images for the MR image study




Figure 6: Ensemble RMSE and SSIM values for in the MR image study. The error bars show a span of one standard deviation.

6. Results

Figure 1 shows some of the reconstructed images for the inverse crime case, with $n/m = 10$ and $n/m = 50$, respectively. The RMSE and SSIM values over the ensemble of images are shown in Fig. 2. It can be seen that the proposed method performs well in terms of RMSE and SSIM even in the severe undersampling case, such as $n/m = 50$.

Figure 3 shows some of the reconstructed images from the face image study. The ensemble RMSE and SSIM values are shown in Fig. 4. It can be seen that the proposed algorithm outperforms the other algorithms considered in terms of perceptual quality, as well as RMSE and SSIM. This is because it is able to better capture the semantic differences between the sought-after and the prior image.

Figure 5 shows the images reconstructed from simulated MRI measurements with undersampling ratios $n/m = 6$ and $n/m = 12$. Although PLS-TV and CSGM have difficulty in properly recovering the image, CSGM has noticeably better defined boundaries. Also, note that the ground truth and the prior image are visually similar in certain regions, and different in others. It can be seen that visually, the performance of the proposed method is the best. Ensemble RMSE and SSIM values for the MR image study are shown in Fig. 6. The results of the robustness study as shown in Fig. 7 indicate that while the proposed approach is robust to misalignments of the ground truth with respect to the prior image in the form of translations and rotations, the PICCS

method breaks down. An example for $i = 4$ is shown in Fig. 8. Additional images for the various studies are included in the supplementary file. Algorithm 1 takes around 5 minutes to converge to an optimal solution on a single Nvidia 1080 TX GPU.




Figure 7: Ensemble RMSE values for the MRI image study, with misaligned ground truth.




Figure 8: Ground truth (GT), prior image (PI) and image estimates from 8 fold undersampled measurements for misalignment configuration C_4

7. Discussion and Conclusion

Firstly, as seen in the numerical studies presented here, the object to-be-imaged and the prior image can be similar in some ways, while differing in others, due to factors such as evolution of the object over time, changes in the exact configuration of the imaging instrument, the depth at which the tomogram slice is recorded, changes in the defining features due to the tumor, changes in the contrast mechanism used, and small changes in pose and orientation. Such changes cannot always be effectively modeled by use of the sparse-differences model assumed in the case of PICCS. PICGM is able to outperform the other methods examined due to the fact that the StyleGAN gives control over individual semantic features in an image via the disentangled latent space.

An important point that needs to be addressed is the choice of p_1, p_2 which determines which styles are kept fixed to those of the prior image and which styles are optimized over. Since we require p_1, p_2 to be multiples of k , the maximum choices for each of p_1 and p_2 are limited to L , the number of layers in the synthesis network. Moreover, in several realistic scenarios, the nature of the semantic differences between the ground truth and the prior image are known, and can be directly related to specific style vectors by performing style-mixing experiments on the trained StyleGAN. For example, small tumors correspond to the medium-scale features, and can be controlled by using the corresponding style vectors. When recovering an MRI image corresponding to one contrast mechanism with the help of a prior image from a different contrast mechanism, the style vectors corresponding to the variation in contrast are known beforehand. In our numerical studies, the values of p_1, p_2 as well as the regu-

larization parameter λ were kept fixed over the test dataset. Since the dataset may not be homogeneous with respect to the differences between the prior image and ground truth, further investigation is needed to devise an adaptive strategy for picking p_1, p_2 . Several techniques for automatically picking regularization parameters exist in the compressive sensing literature which could be adapted to the proposed method (Weizman et al., 2016; Mathew & Paul, 2020). In the presented numerical studies, the best parameters were selected by tuning over 4-5 (p_1, p_2) pairs. The styles can be grouped into coarse, medium and fine for choices of p_i between $0-2k$, $2k-8k$ and $8k-14k$ respectively. As long as p_i resided in the correct group, a $\sim 10\%$ median degradation in RMSE across the incorrect choices of (p_1, p_2) was encountered.

The theoretical analysis presented here potentially offers a way to analyze the stability of the reconstructed estimate. In particular, the presented analysis applies to non-uniform recovery of in-distribution objects, that are sufficiently regular, in the sense that the Jacobian of the generative network computed at their latent representation is well behaved. This means that small changes in the estimated \hat{w} due to the measurement noise do not correspond to huge changes in the estimated image, which is what is obtained via the presented recovery guarantees. Hence, if the proposed algorithm converges to a solution that does not satisfy the properties presented in Lemma 4.1, it likely has converged to a highly unstable and nonlinear region of the optimization landscape. This provides a certain degree of graceful failure, in the face of hallucinations that may occur in the image reconstruction (Bhadra et al., 2021). It must be noted that certain assumptions in the theoretical analysis, such as the validity of a useful S-REC for non-Gaussian sensing matrices, and obtaining near-optimal approximations to Eq. (9) are not completely realistic, but have been made in recent literature, and there has been progress towards addressing them (Gómez et al., 2019; Daras et al., 2021).

Lastly, a task-informed approach to evaluating image reconstruction algorithms is necessary in addition to traditional image quality metrics employed in this study.

Acknowledgements

The authors would like to thank Sayantan Bhadra and Weimin Zhou for their help. This work was supported in part by NIH Awards EB020604, EB023045, NS102213, EB028652, and NSF Award DMS1614305.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard,

- M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Abdal, R., Qin, Y., and Wonka, P. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4432–4441, 2019.
- Adcock, B., Hansen, A. C., Poon, C., and Roman, B. Breaking the coherence barrier: A new theory for compressed sensing. *Forum of Mathematics, Sigma*, 5:e4 1–84, 2017.
- Asim, M., Daniels, G., Leong, O., Ahmed, A., and Hand, P. Invertible generative models for inverse problems: mitigating representation error and dataset bias. In *Proceedings of the International Conference on Machine Learning*, pp. 4577–4587, 2020.
- Barrett, H. H. and Myers, K. J. *Foundations of image science*. John Wiley & Sons, 2013.
- Becquaert, M., Cristofani, E., Lauwens, B., Vandewal, M., Stiens, J. H., and Deligiannis, N. Online sequential compressed sensing with multiple information for through-the-wall radar imaging. *IEEE Sensors Journal*, 19(11): 4138–4148, 2019.
- Bhadra, S., Kelkar, V. A., Brooks, F. J., and Anastasio, M. A. On hallucinations in tomographic image reconstruction. *IEEE Transactions on Medical Imaging*, pp. 1–1, 2021. doi: 10.1109/TMI.2021.3077857.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 537–546. JMLR. org, 2017.
- Cai, H., Bai, C., Tai, Y.-W., and Tang, C.-K. Deep video generation, prediction and completion of human action sequences. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 366–382, 2018.
- Candès, E. J. and Wakin, M. B. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2): 21–30, 2008.
- Candès, E. J., Romberg, J. K., and Tao, T. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006. doi: <https://doi.org/10.1002/cpa.20124>.
- Chen, G.-H., Tang, J., and Leng, S. Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets. *Medical physics*, 35(2):660–663, 2008.
- Daras, G., Dean, J., Jalal, A., and Dimakis, A. G. Intermediate layer optimization for inverse problems using deep generative models. *arXiv preprint arXiv:2102.07364*, 2021.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Gómez, F. L., Eftekhari, A., and Cevher, V. Fast and provable admm for learning with generative priors. *arXiv preprint arXiv:1907.03343*, 2019.
- Hussein, S. A., Tirer, T., and Giryes, R. Image-adaptive gan based reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3121–3129, 2020.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8110–8119, 2020.
- Kelkar, V. A., Bhadra, S., and Anastasio, M. A. Compressible latent-space invertible networks for generative model-constrained image reconstruction. *IEEE Transactions on Computational Imaging*, 7:209–223, 2021. doi: 10.1109/TCI.2021.3049648.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- LaMontagne, P. J., Benzinger, T. L., Morris, J. C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A., et al. Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *MedRxiv*, 2019.
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690, 2017.

- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. Compressed sensing MRI. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, pp. 3. Citeseer, 2013.
- Marinescu, R. V., Moyer, D., and Golland, P. Bayesian image reconstruction using deep generative models. *arXiv preprint arXiv:2012.04567*, 2020.
- Mathew, R. S. and Paul, J. S. Automated regularization parameter selection using continuation based proximal method for compressed sensing mri. *IEEE Transactions on Computational Imaging*, 6:1309–1319, 2020.
- Menon, S., Damian, A., Hu, S., Ravi, N., and Rudin, C. PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2437–2445, 2020.
- Mota, J. F., Deligiannis, N., and Rodrigues, M. R. Compressed sensing with prior information: Strategies, geometry, and bounds. *IEEE Transactions on Information Theory*, 63(7):4472–4496, 2017.
- Prakash, C. D. and Karam, L. J. It gan do better: Gan-based detection of objects on images with varying quality. *arXiv preprint arXiv:1912.01707*, 2019.
- Scarpace, L., Mikkelsen, T., Cha, S., Rao, S., Tekchandani, S., Gutman, D., Saltz, J. H., Erickson, B. J., Pedano, N., Flanders, A. E., Barnholtz-Sloan, J., Ostrom, Q., Barboriak, D., and Pierce, L. J. Radiology Data from The Cancer Genome Atlas Glioblastoma Multiforme [TCGA-GBM] collection. *The Cancer Imaging Archive*, 2016.
- Schmainda, K. and Prah, M. Data from Brain-Tumor-Progression. *The Cancer Imaging Archive*, 2018. URL <http://doi.org/10.7937/K9/TCIA.2018.15quzvn>.
- Shutterstock. Standard License., 2021. <https://www.shutterstock.com/home>.
- Tewari, A., Elgharib, M., Bharaj, G., Bernard, F., Seidel, H.-P., Pérez, P., Zollhofer, M., and Theobalt, C. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6142–6151, 2020.
- Vershynin, R. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Weizman, L., Eldar, Y. C., and Ben Bashat, D. Reference-based mri. *Medical physics*, 43(10):5357–5369, 2016.
- Whang, J., Lei, Q., and Dimakis, A. G. Compressed sensing with invertible generative models and dependent noise. *arXiv preprint arXiv:2003.08089*, 2020.
- Wulff, J. and Torralba, A. Improving inversion and generation diversity in stylegan using a gaussianized latent space. *arXiv preprint arXiv:2009.06529*, 2020.
- Zbontar, J., Knoll, F., Sriram, A., Muckley, M. J., Bruno, M., Defazio, A., Parente, M., Geras, K. J., Katsnelson, J., Chandarana, H., et al. fastMRI: An open dataset and benchmarks for accelerated MRI. *arXiv preprint arXiv:1811.08839*, 2018.
- Zhang, Y., Zhang, Y., and Cai, W. Separating style and content for generalized style transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8447–8455, 2018.