
Modeling Sparse Deviations for Compressed Sensing using Generative Models

Manik Dhar¹ Aditya Grover¹ Stefano Ermon¹

Abstract

In compressed sensing, a small number of linear measurements can be used to reconstruct an unknown signal. Existing approaches leverage assumptions on the structure of these signals, such as sparsity or the availability of a generative model. A domain-specific generative model can provide a stronger prior and thus allow for recovery with far fewer measurements. However, unlike sparsity-based approaches, existing methods based on generative models guarantee exact recovery only over their support, which is typically only a small subset of the space on which the signals are defined. We propose Sparse-Gen, a framework that allows for *sparse* deviations from the support set, thereby achieving the best of both worlds by using a domain specific prior and allowing reconstruction over the full space of signals. Theoretically, our framework provides a new class of signals that can be acquired using compressed sensing, reducing classic sparse vector recovery to a special case and avoiding the restrictive support due to a generative model prior. Empirically, we observe consistent improvements in reconstruction accuracy over competing approaches, especially in the more practical setting of *transfer* compressed sensing where a generative model for a data-rich, source domain aids sensing on a data-scarce, target domain.

1. Introduction

In many real-world domains, data acquisition is costly. For instance, magnetic resonance imaging (MRI) requires scan times proportional to the number of measurements, which can be significant for patients (Lustig et al., 2008). Geophysical applications like oil drilling require expensive simulation of seismic waves (Qaisar et al., 2013). Such appli-

¹Computer Science Department, Stanford University, CA, USA. Correspondence to: Manik Dhar <dmanik@cs.stanford.edu>, Aditya Grover <adityag@cs.stanford.edu>.

cations, among many others, can benefit significantly from *compressed sensing* techniques to acquire signals efficiently (Candès & Tao, 2005; Donoho, 2006; Candès et al., 2006).

In compressed sensing, we wish to acquire an n -dimensional signal $x \in \mathbb{R}^n$ using only $m \ll n$ measurements linear in x . The measurements could potentially be noisy, but even in the absence of any noise we need to impose additional structure on the signal to guarantee unique recovery. Classical results on compressed sensing impose structure by assuming the underlying signal to be approximately l -sparse in some known basis, *i.e.*, the l -largest entries dominate the rest. For instance, images and audio signals are typically sparse in the wavelet and Fourier basis respectively (Mallat, 2008). If the matrix of linear vectors relating the signal and measurements satisfies certain mild conditions, then one can provably recover x with only $m = O(l \log \frac{n}{l})$ measurements using LASSO (Tibshirani, 1996; Candès & Tao, 2005; Donoho, 2006; Candès et al., 2006; Bickel et al., 2009).

Alternatively, structural assumptions on the signals being sensed can be *learned* from data, *e.g.*, using a dataset of typical signals (Baraniuk et al., 2010; Peyre, 2010; Chen et al., 2010; Yu & Sapiro, 2011). Particularly relevant to this work, Bora et al. (2017) proposed an approach where structure is provided by a deep generative model learned from data. Specifically, the underlying signal x being sensed is assumed to be close to the range of a deterministic function expressed by a pretrained, latent variable model $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ such that $x \approx G(z)$ where $z \in \mathbb{R}^k$ denote the latent variables. Consequently, the signal x is recovered by optimizing for a latent vector z that minimizes the ℓ_2 distance between the measurements corresponding to $G(z)$ and the actual ones. Even though the objective being optimized in this case is non-convex, empirical results suggest that the reconstruction error decreases much faster than LASSO-based recovery as we increase the number of measurements.

A limitation of the above approach is that the recovered signal is constrained to be in the range of the generator function G . Hence, if the true signal being sensed is not in the range of G , the algorithm cannot drive the reconstruction error to zero even when $m \geq n$ (even if we ignore error due to measurement noise and non-convex optimization). This is also observed empirically, as the reconstruction error of generative model-based recovery saturates as we keep

increasing the number of measurements m . On the other hand, LASSO-based recovery continues to shrink the error with increasing number of measurements, eventually outperforming the generative model-based recovery.

To overcome this limitation, we propose a framework that allows recovery of signals with *sparse deviations* from the set defined by the range of the generator function. The recovered signals have the general form of $G(\hat{z}) + \hat{v}$, where $\hat{v} \in \mathbb{R}^n$ is a sparse vector. This allows the recovery algorithm to consider signals away from the range of the generator function. Similar to LASSO, we relax the hardness in optimizing for sparse vectors by minimizing the ℓ_1 norm of the deviations. Unlike LASSO-based recovery, we can exploit the rich structure imposed by a (deep) generative model (at the expense of solving a hard optimization problem if G is non-convex). In fact, we show that LASSO-based recovery is a special case of our framework if the generator function G maps all z to the origin. Unlike generative model-based recovery, the signals recovered by our algorithm are not constrained to be in the range of the generator function.

Our proposed algorithm, referred to as Sparse-Gen, has desirable theoretical properties and empirical performance. Theoretically, we derive upper bounds on the reconstruction error for an optimal decoder with respect to the proposed model and show that this error vanishes with $m = n$ measurements. We confirm our theory empirically, wherein we find that recovery using Sparse-Gen with variational autoencoders (Kingma & Welling, 2014) as the underlying generative model outperforms both LASSO-based and generative model-based recovery in terms of the reconstruction errors for the same number of measurements for MNIST and Omniglot datasets. Additionally, we observe significant improvements in the more practical and novel task of *transfer* compressed sensing where a generative model on a data-rich, source domain provides a prior for sensing a data-scarce, target domain.

2. Preliminaries

In this section, we review the necessary background and prior work in modeling domain specific structure in compressed sensing. We are interested in solving the following system of equations,

$$y = Ax \quad (1)$$

where $x \in \mathbb{R}^n$ is the signal of interest being sensed through measurements $y \in \mathbb{R}^m$, and $A \in \mathbb{R}^{m \times n}$ is a measurement matrix. For efficient acquisition of signals, we will design measurement matrices such that $m \ll n$. However, the system is under-determined whenever $\text{rank}(A) < n$. Hence, unique recovery requires additional assumptions on x . We now discuss two ways to model the structure of x .

Sparsity. Sparsity in a well-chosen basis is natural in many domains. For instance, natural images are sparse in the wavelet basis whereas audio signals exhibit sparsity in the Fourier basis (Mallat, 2008). Hence, it is natural to assume the domain of signals x we are interested in recovering is

$$S_l(0) = \{x : \|x - 0\|_0 \leq l\}. \quad (2)$$

This is the set of l -sparse vectors with the ℓ_0 distance measured from the origin. Such assumptions dominate the prior literature in compressed sensing and can be further relaxed to recover *approximately* sparse signals (Candès & Tao, 2005; Donoho, 2006; Candès et al., 2006).

Latent variable generative models. A latent variable model specifies a joint distribution $P_\theta(x, z)$ over the observed data x (e.g., images) and a set of latent variables $z \in \mathbb{R}^k$ (e.g., features). Given a training set of signals $\{x_1, \dots, x_M\}$, we can learn the parameters θ of such a model, e.g., via maximum likelihood. When $P_\theta(x, z)$ is parameterized using deep neural networks, such *generative* models can effectively model complex, high-dimensional signal distributions for modalities such as images and audio (Kingma & Welling, 2014; Goodfellow et al., 2014).

Given a pretrained latent variable generative model with parameters θ , we can associate a *generative model function* $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ mapping a latent vector z to the mean of the conditional distribution $P_\theta(x|z)$. Thereafter, the space of signals that can be recovered with such a model is given by the range of the generator function,

$$S_G = \{G(z) : z \in \mathbb{R}^k\}. \quad (3)$$

Note that the set is defined with respect to the latent vectors z , and we omit the dependence of G on the parameters θ (which are fixed for a pretrained model) for brevity.

2.1. Recovery algorithms

Signal recovery in compressed sensing algorithm typically involves solving an optimization problem consistent with the modeling assumptions on the domain of the signals being sensed.

Sparse vector recovery using LASSO. Under the assumptions of sparsity, the signal x can be recovered by solving an ℓ_0 minimization problem (Candès & Tao, 2005; Donoho, 2006; Candès et al., 2006).

$$\begin{aligned} & \min_x \|x\|_0 \\ & \text{s.t. } Ax = y. \end{aligned} \quad (4)$$

The objective above is however NP-hard to optimize, and hence, it is standard to consider a convex relaxation,

$$\begin{aligned} & \min_x \|x\|_1 \\ & \text{s.t. } Ax = y. \end{aligned} \quad (5)$$

In practice, it is common to solve the Lagrangian of the above problem. We refer to this method as LASSO-based recovery due to similarities of the objective in Eq. (5) to the LASSO regularization used broadly in machine learning (Tibshirani, 1996). LASSO-based recovery is the predominant technique for recovering sparse signals since it involves solving a tractable convex optimization problem.

In order to guarantee unique recovery to the underdetermined system in Eq. (1), the measurement matrix A is designed to satisfy the Restricted Isometry Property (RIP) or the Restricted Eigenvalue Condition (REC) for l -sparse matrices with high probability (Candès & Tao, 2005; Bickel et al., 2009). We define these conditions below.

Definition 1. Let $S_l(0) \subset \mathbb{R}^n$ be the set of l -sparse vectors. For some parameter $\alpha \in (0, 1)$, a matrix $A \in \mathbb{R}^{m \times n}$ is said to satisfy RIP(l, α) if $\forall x \in S_l(0)$,

$$(1 - \alpha)\|x\|_2 \leq \|Ax\|_2 \leq (1 + \alpha)\|x\|_2.$$

Definition 2. Let $S_l(0) \subset \mathbb{R}^n$ be the set of l -sparse vectors. For some parameter $\gamma > 0$, a matrix $A \in \mathbb{R}^{m \times n}$ is said to satisfy REC(l, γ) if $\forall x \in S_l(0)$,

$$\|Ax\|_2 \geq \gamma\|x\|_2.$$

Intuitively, RIP implies that A approximately preserves Euclidean norms for sparse vectors and REC implies that sparse vectors are far from the nullspace of A . Many classes of matrices satisfy these conditions with high probability, including random Gaussian and Bernoulli matrices where every entry of the matrix is sampled from a standard normal and uniform Bernoulli distribution respectively (Baraniuk et al., 2008).

Generative model vector recovery using gradient descent. If the signals being sensed are assumed to lie close to the range S_G of a generative model function G as defined in Eq. (3), then we can recover the best approximation to the true signal by ℓ_2 -minimization over z ,

$$\min_z \|AG(z) - y\|_2^2. \quad (6)$$

The function G is typically expressed as a deep neural network which makes the overall objective non-convex, but differentiable almost everywhere w.r.t z . In practice, good reconstructions can be recovered by gradient-based optimization methods. We refer to this method proposed by Bora et al. (2017) as *generative model-based recovery*.

To guarantee unique recovery, generative model-based recovery makes two key assumptions. First, the generator function G is assumed to be L -Lipschitz, i.e., $\forall z_1, z_2 \in \mathbb{R}^k$,

$$\|G(z_1) - G(z_2)\|_2 \leq L\|z_1 - z_2\|_2.$$

Secondly, the measurement matrix A is designed to satisfy the Set-Restricted Eigenvalue Condition (S-REC) with high probability (Bora et al., 2017).

Definition 3. Let $S \subseteq \mathbb{R}^n$. For some parameters $\gamma > 0$, $\delta \geq 0$, a matrix $A \in \mathbb{R}^{m \times n}$ is said to satisfy the S-REC(S, γ, δ) if $\forall x_1, x_2 \in S$,

$$\|A(x_1 - x_2)\|_2 \geq \gamma\|x_1 - x_2\|_2 - \delta.$$

S-REC generalizes REC to an arbitrary set of vectors S as opposed to just considering the set of approximately sparse vectors $S_l(0)$ and allowing an additional slack term δ . In particular, S is chosen to be the range of the generator function G for generative model-based recovery.

3. The Sparse-Gen framework

The modeling assumptions based on sparsity and generative modeling discussed in the previous section can be limiting in many cases. On one hand, sparsity assumes a relatively weak prior over the signals being sensed. Empirically, we observe that the recovered signals x_L have large reconstruction error $\|x_L - x\|_2^2$ especially when the number of measurements m is small. On the other hand, generative models imposes a very strong, but rigid prior which works well when the number of measurements is small. However, the performance of the corresponding recovery methods saturates with increasing measurements since the recovered signal $x_G = G(z_G)$ is constrained to lie in the range of the generator function G . If $z_G \in \mathbb{R}^k$ is the optimum value returned by an optimization procedure for Eq. (6), then the reconstruction error $\|x_G - x\|_2^2$ is limited by the dimensionality of the latent space and the quality of the generator function.

To sidestep the above limitations, we consider a strictly more expressive class of signals by allowing sparse deviations from the range of a generator function. Formally, the domain of the recovered signals is given by,

$$S_{l,G} = \cup_{z \in \text{Dom}(G)} S_l(G(z)) \quad (7)$$

where $S_l(G(z))$ denotes the set of sparse vectors centered on $G(z)$ and z varies over the domain of G (typically \mathbb{R}^k). We refer to this modeling assumption and the consequent algorithmic framework for recovery as *Sparse-Gen*.

Based on this modeling assumption, we will recover signals of the form $G(z) + \nu$ for some $\nu \in \mathbb{R}^n$ that is preferably sparse. Specifically, we consider the optimization of a hybrid objective,

$$\begin{aligned} & \min_{z, \nu} \|\nu\|_0 \\ & \text{s.t. } A(G(z) + \nu) = y. \end{aligned} \quad (8)$$

In the above optimization problem the objective is non-convex and non-differentiable, while the constraint is non-convex (for general G), making the above optimization




Figure 1. LASSO vs. Generative model vs. Sparse-Gen recovery. Unlike LASSO, Sparse-Gen imposes a stronger prior on the signals being sensed (**shaded grey regions**). Unlike generative model, the recovered signals are not constrained to lie on the range of the generator function (**red points**). Aspects of visualization due to eigenvalue conditions not shown for simplicity.

problem hard to solve. To ease the optimization problem, we propose two modifications. First, we relax the ℓ_0 minimization to an ℓ_1 minimization similar to LASSO.

$$\begin{aligned} & \min_{z, \nu} \|\nu\|_1 \\ \text{s.t. } & A(G(z) + \nu) = y. \end{aligned} \quad (9)$$

Next, we square the non-convex constraint on both sides and consider the Lagrangian of the above problem to get the final unconstrained optimization problem for Sparse-Gen,

$$\min_{z, \nu} \|\nu\|_1 + \lambda \|A(G(z) + \nu) - y\|_2^2 \quad (10)$$

where λ is the Lagrange multiplier.

The above optimization problem is non-differentiable w.r.t. ν and non-convex w.r.t. z (if G is non-convex). In practice, it can be solved in practice using gradient descent (since the non-differentiability is only at a finite number of points) or using sequential convex programming (SCP). SCP is an effective heuristic for non-convex problems where the convex portions of the problem are solved using a standard convex optimization technique (Boyd & Vandenberghe, 2004). In the case of Eq. (10), the optimization w.r.t. ν (for fixed z) is a convex optimization problem whereas the non-convexity typically involves differentiable terms (w.r.t. z) if G is a deep neural network. Empirically, we find excellent recovery by standard first order gradient-based methods (Duchi et al., 2011; Tieleman & Hinton, 2012; Kingma & Ba, 2015).

Unlike LASSO-based recovery which recovers only sparse signals, Sparse-Gen can impose a stronger domain-specific prior using a generative model. If we fix the generator function to map all z to the origin, we recover LASSO-based recovery as a special case of Sparse-Gen. Additionally, Sparse-Gen is not constrained to recover signals over the range of G , as in the case of generative model-based recovery. In fact, it can recover signals with sparse deviations from the range of G . Note that the sparse deviations can be

defined in a basis different from the canonical basis. In such cases, we consider the following optimization problem,

$$\min_{z, \nu} \|B\nu\|_1 + \lambda \|A(G(z) + \nu) - y\|_2^2 \quad (11)$$

where B is a change of basis matrix that promotes sparsity of the vector $B\nu$. Figure 1 illustrates the differences in modeling assumptions between Sparse-Gen and other frameworks.

4. Theoretical Analysis

The proofs for all results in this section are given in the Appendix. Our analysis and experiments account for measurement noise ϵ in compressed sensing, *i.e.*,

$$y = Ax + \epsilon. \quad (12)$$

Let $\Delta : \mathbb{R}^m \rightarrow \mathbb{R}^n$ denote an arbitrary decoding function used to recover the true signal x from the measurements $y \in \mathbb{R}^m$. Our analysis will upper bound the ℓ_2 -error in recovery incurred by our proposed framework using mixed norm guarantees (in particular, ℓ_2/ℓ_1). To this end, we first state some key definitions. Define the least possible ℓ_1 error for recovering x under the Sparse-Gen modeling as,

$$\sigma_{S_{l,G}}(x) = \inf_{\hat{x} \in S_{l,G}} \|x - \hat{x}\|_1$$

where the optimal \hat{x} is the closest point to x in the allowed domain $S_{l,G}$. We now state the main lemma guiding the theoretical analysis.

Lemma 1. *Given a function $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ and measurement noise ϵ with $\|\epsilon\|_2 \leq \epsilon_{\max}$, let A be any matrix that satisfies $S\text{-REC}(S_{1.5l,G}, (1 - \alpha), \delta)$ and $\text{RIP}(2l, \alpha)$ for some $\alpha \in (0, 1)$, $l > 0$. Then, there exists a decoder $\Delta : \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that,*

$$\|x - \Delta(Ax + \epsilon)\|_2 \leq (2l)^{-1/2} C_0 \sigma_{l,G}(x) + C_1 \epsilon_{\max} + \delta'$$

for all $x \in \mathbb{R}^n$, where $C_0 = 2((1 + \alpha)(1 - \alpha)^{-1} + 1)$, $C_1 = 2(1 - \alpha)^{-1}$, and $\delta' = \delta(1 - \alpha)^{-1}$.

The above lemma shows that there exists a decoder such that the error in recovery can be upper bounded for measurement matrices satisfying S-REC and RIP. Note that Lemma 1 only guarantees the existence of such a decoder and does not prescribe an optimization algorithm for recovery. Apart from the errors due to the bounded measurement noise ϵ_{\max} and a scaled slack term appearing in the S-REC condition δ' , the major term in the upper bound corresponds to (up to constants) the minimum possible error incurred by the best possible recovery vector in $S_{l,G}$ given by $\sigma_{l,G}(x)$. Similar terms appear invariably in the compressed sensing literature and are directly related to the modeling assumptions regarding x (for example, Theorem 8.3 in Cohen et al. (2009)).

Our next lemma shows that random Gaussian matrices satisfy the S-REC (over the range of Lipschitz generative model functions) and RIP conditions with high probability for G with bounded domain, both of which together are sufficient conditions for Lemma 1 to hold.

Lemma 2. *Let $G : B^k(r) \rightarrow \mathbb{R}^n$ be an L -Lipschitz function where $B^k(r) = \{z \mid z \in \mathbb{R}^k, \|z\|_2 \leq r\}$ is the ℓ_2 -norm ball in \mathbb{R}^k . For $\alpha \in (0, 1)$, if*

$$m = O\left(\frac{1}{\alpha^2} \left(k \log\left(\frac{Lr}{\delta}\right) + l \log(n/l)\right)\right)$$

then a random matrix $A \in \mathbb{R}^{m \times n}$ with i.i.d. entries such that $A_{ij} \sim \mathcal{N}(0, \frac{1}{m})$ satisfies the S-REC($S_{1.5l,G}, 1 - \alpha, \delta$) and RIP($2l, \alpha$) with $1 - e^{-\Omega(\alpha^2 m)}$ probability.

Using Lemma 1 and Lemma 2, we can bound the error due to decoding with generative models and random Gaussian measurement matrices in the following result.

Theorem 1. *Let $G : B^k(r) \rightarrow \mathbb{R}^n$ be an L -Lipschitz function. For any $\alpha \in (0, 1)$, $l > 0$, let $A \in \mathbb{R}^{m \times n}$ be a random Gaussian matrix with*

$$m = O\left(\frac{1}{\alpha^2} \left(k \log\left(\frac{Lr}{\delta}\right) + l \log(n/l)\right)\right)$$

rows of i.i.d. entries scaled such that $A_{i,j} \sim N(0, 1/m)$. Let Δ be the decoder satisfying Lemma 1. Then, we have with $1 - e^{-\Omega(\alpha^2 m)}$ probability,

$$\|x - \Delta(Ax + \epsilon)\|_2 \leq (2l)^{-1/2} C_0 \sigma_{l,G}(x) + C_1 \epsilon_{\max} + \delta'$$

for all $x \in \mathbb{R}^n$, $\|\epsilon\|_2 \leq \epsilon_{\max}$, where $C_0, C_1, \gamma, \delta'$ are constants defined in Lemma 1.

From the above lemma, we see that the number of measurements needed to guarantee upper bounds on the reconstruction error of any signal with high probability depends on two terms. The first term includes dependence on the Lipschitz constant L of the generative model function G . A high Lipschitz constant makes recovery harder (by requiring a larger

number of measurements), but only contributes logarithmically. The second term, typical of results in sparse vector recovery, shows a logarithmic growth on the dimensionality n of the signals. Ignoring logarithmic dependences and constants, recovery using Sparse-Gen requires about $O(k + l)$ measurements for recovery. Note that Theorem 1 assumes access to an optimization oracle for decoding. In practice, we consider the solutions returned by gradient-based optimization methods to a non-convex objective defined in Eq. (11) that are not guaranteed to correspond to the optimal decoding in general.

Finally, we obtain tighter bounds for the special case when G is expressed using a neural network with only ReLU activations. These bounds do not rely explicitly on the Lipschitz constant L or require the domain of G to be bounded.

Theorem 2. *If $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ is a neural network of depth d with only ReLU activations and at most c nodes in each layer, then the guarantees of Theorem 1 hold for*

$$m = O\left(\frac{1}{\alpha^2} \left((k + l)d \log c + (k + l) \log(n/l)\right)\right).$$

Our theoretical analysis formalizes the key properties of recovering signals using Sparse-Gen. As shown in Lemma 1, there exists a decoder for recovery based on such modeling assumptions that extends recovery guarantees based on vanilla sparse vector recovery and generative model-based recovery. Such recovery requires measurement matrices that satisfy both the RIP and S-REC conditions over the set of vectors that deviate in sparse directions from the range of a generative model function. In Theorems 1-2, we observed that the number of measurements required to guarantee recovery with high probability grow almost linearly (with some logarithmic terms) with the latent space dimensionality k of the generative model and the permissible sparsity l for deviating from the range of the generative model.

5. Experimental Evaluation

We evaluated Sparse-Gen for compressed sensing of high-dimensional signals from the domain of benchmark image datasets. Specifically, we considered the MNIST dataset of handwritten digits (LeCun et al., 2010) and the OMNIGLOT dataset of handwritten characters (Lake et al., 2015). Both these datasets have the same data dimensionality (28×28), but significantly different characteristics. The MNIST dataset has fewer classes (10 digits from 0-9) as opposed to Omniglot which shows greater diversity (1623 characters across 50 alphabets). Additional experiments with generative adversarial networks on the CelebA dataset are reported in the Appendix.

Baselines. We considered methods based on sparse vector recovery using LASSO (Tibshirani, 1996; Candès & Tao,

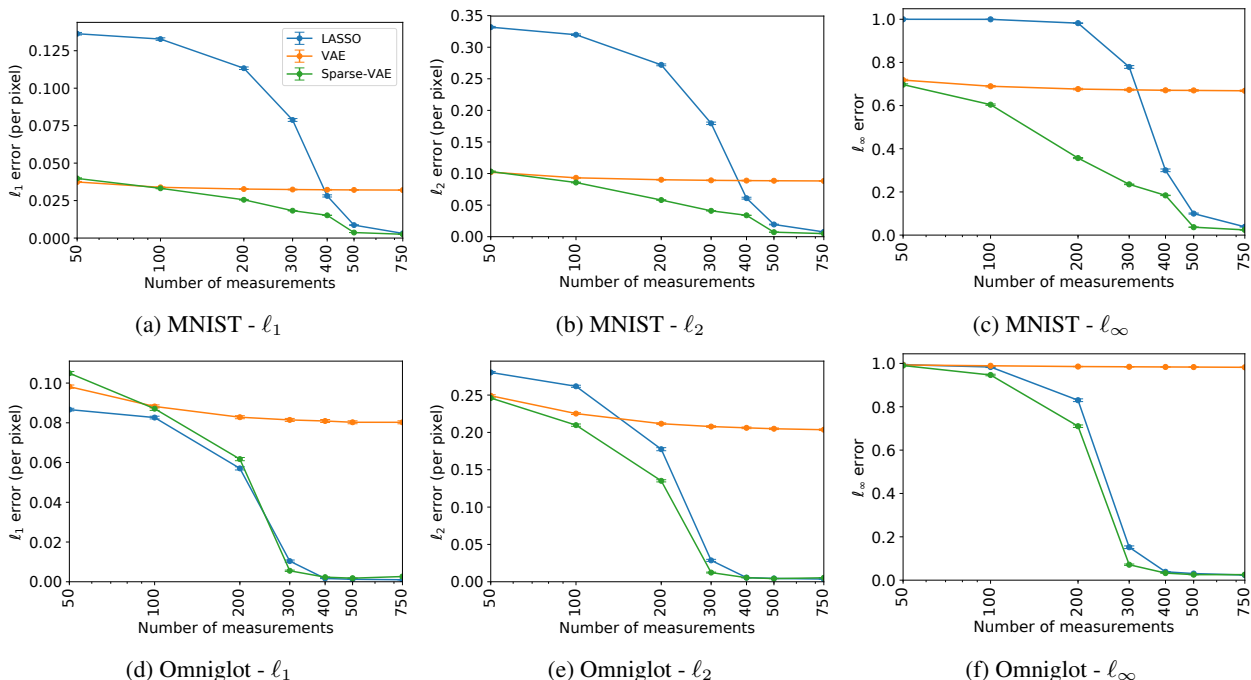


Figure 2. Reconstruction error using ℓ_1 (left), ℓ_2 (center), and ℓ_∞ (right) norms for MNIST (top) and Omniglot (bottom). The performance of Sparse-Gen is better or competitive with both generative model-based methods and sparse vector recovery methods.

2005) and generative model based recovery using variational autoencoders (VAE) (Kingma & Welling, 2014; Bora et al., 2017). For VAE training, we used the standard train/hold-out splits of both datasets. Compressed sensing experiments that we report were performed on the entire test set of images. The architecture and other hyperparameter details are given in the Appendix.

Experimental setup. For the held-out set of instances, we artificially generated measurements y through a random matrix $A \in \mathbb{R}^{m \times n}$ with entries sampled i.i.d. from a Gaussian with zero mean and standard deviation of $1/m$. Measurement noise is sampled from zero mean and diagonal scalar covariance matrix with entries as 0.01. For evaluation, we report the reconstruction error measured as $\|\hat{x} - x\|_p$ where \hat{x} is the recovered signal and p is a norm of interest, varying the number of measurements m from 50 to the highest value of 750. We report results for the $p = \{1, 2, \infty\}$ norms.

We evaluated sensing of both continuous signals (MNIST) with pixel values in range $[0, 1]$ and discrete signals (Omniglot) with binary pixel values $\{0, 1\}$. For all algorithms considered, recovery was performed by optimizing over a continuous space. In the case of sparse recovery methods (including Sparse-Gen) it is possible that unconstrained optimization returns signals outside the domain of interest, in which case they are projected to the required domain by simple clipping, *i.e.*, any signal less than zero is clipped to 0 and similarly any signal greater than one is clipped to 1.

Results and Discussion. The reconstruction errors for varying number of measurements are given in Figure 2. Consistent with the theory, the strong prior in generative model-based recovery methods outperforms the LASSO-based methods for sparse vector recovery. In the regime of low measurements, the performance of algorithms that can incorporate the generative model prior dominates over methods modeling sparsity using LASSO. The performance of plain generative model-based methods however saturates with increasing measurements, unlike Sparse-Gen and LASSO which continue to shrink the error. The trends are consistent for both MNIST and Omniglot, although we observe the relative magnitudes of errors in the case of Omniglot are much higher than that of MNIST. This is expected due to the increased diversity and variations of the structure of the signals being sensed in the case of Omniglot. We also observe the trends to be consistent across the various norms considered.

5.1. Transfer compressed sensing

One of the primary motivations for compressive sensing is to directly acquire the signals using few measurements. On the contrary, learning a deep generative model requires access to large amounts of training data. In several applications, getting the data for training a generative model might not be feasible. Hence, we test the generative model-based recovery on the novel task of *transfer* compressed sensing.

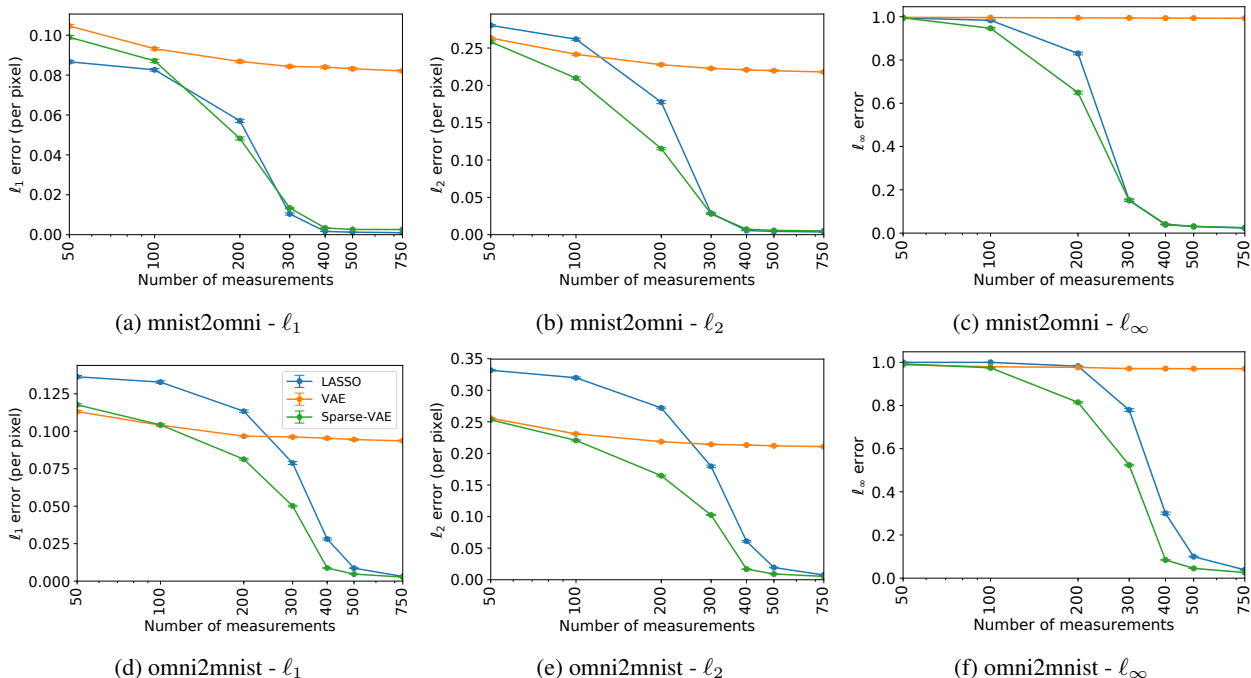


Figure 3. Reconstruction error in terms of ℓ_1 (left), ℓ_2 (center), and ℓ_∞ (right) norms for transfer compressed sensing from MNIST as source to Omniglot as target (top) and Omniglot as source and MNIST as target (bottom). Under the ℓ_2 and ℓ_∞ metric, Sparse-VAE beats both VAE and LASSO. Under the ℓ_1 metric, performance is slightly poorer until 100 measurements.

Experimental setup. We train the generative model on a source domain (assumed to be data-rich) and related to a data-hungry target domain we wish to sense. Given the matching dimensions of MNIST and Omniglot, we conduct experiments transferring from MNIST (source) to Omniglot (target) and vice versa.

Results and Discussion. The reconstruction errors for the norms considered are given in Figure 3. For both the source-target pairs, we observe that the Sparse-Gen consistently performs well. Vanilla generative model-based recovery shows hardly an improvements with increasing measurements. We can qualitatively see this phenomena for transferring from MNIST (source) to Omniglot (target) in Figure 4. With only $m = 100$ measurements, all models perform poorly and generative model based methods particularly continue to sense images similar to MNIST. On the other hand, there is a noticeable transition at $m = 200$ measurements for Sparse-VAE where it adapts better to the domain being sensed than plain generative model-based recovery and achieves lower reconstruction error.

6. Related Work

Since the introduction of compressed sensing over a decade ago, there has been a vast body of research studying various extensions and applications (Candès & Tao, 2005; Donoho,

2006; Candès et al., 2006). This work explores the effect of modeling different structural assumptions on signals in theory and practice.

Themes around sparsity in a well-chosen basis has driven much of the research in this direction. For instance, the paradigm of model-based compressed sensing accounts for the interdependencies between the dimensions of a sparse data signal (Baraniuk et al., 2010; Duarte & Eldar, 2011; Gilbert et al., 2017). Alternatively, adaptive selection of basis vectors from a dictionary that best capture the structure of the particular signal being sensed has also been explored (Peyre, 2010; Tang et al., 2013). Many of these methods have been extended to recovery of structured tensors (Zhang et al., 2013; 2014). In another prominent line of research involving Bayesian compressed sensing, the sparseness assumption is formalized by placing sparseness-promoting priors on the signals (Ji et al., 2008; He & Carin, 2009; Babacan et al., 2010; Baron et al., 2010).

Research exploring structure beyond sparsity is relatively scarce. Early works in this direction can be traced to Baraniuk & Wakin (2009) who proposed algorithms for recovering signals lying on a smooth manifold. The generative model-based recovery methods consider functions that do not necessarily define manifolds since the range of a generator function could intersect with itself. Yu & Sapiro (2011) coined the term statistical compressed sensing and proposed

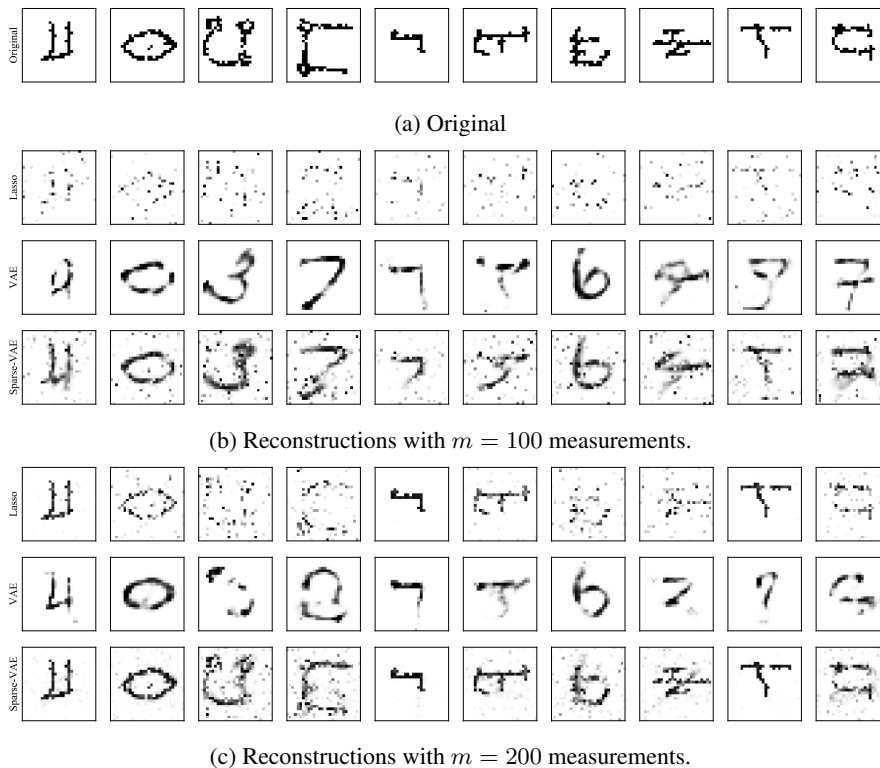


Figure 4. The image matrix above shows the recovery of Omniglot images (**top**) using a VAE trained over MNIST using 100 measurements (**middle**) and 200 measurements (**bottom**). We note that with 100 measurements recovery is very poor across the board. With 200 measurements, all three methods show significant improvement and Sparse-VAE shows greater sensitivity to the domain being sensed.

algorithms for efficient sensing of signals from a mixture of Gaussians. The recent work in deep generative model-based recovery differs in key theoretical aspects as well in the use of a more expressive family of models based on neural networks. A related recent work by Hand & Voroninski (2017) provides theoretical guarantees on the solution recovered for solving non-convex linear inverse problems with deep generative priors. Empirical advances based on well-designed deep neural network architectures that sacrifice many of the theoretical guarantees have been proposed for applications such as MRI (Mardani et al., 2017; 2018). Many recent methods propose to learn mappings of signals to measurements using neural networks, instead of restricting them to be linear, random matrices (Mousavi et al., 2015; Kulkarni et al., 2016; Chang et al., 2017; Lu et al., 2018).

Our proposed framework bridges the gap between algorithms that model structure using sparsity and enjoy good theoretical properties with advances in deep generative models, in particular their use for compressed sensing.

7. Conclusion and Future Work

The use of deep generative models as priors for compressed sensing presents a new outlook on algorithms for inexpen-

sive data acquisition. In this work, we showed that these priors can be used in conjunction with classical modeling assumptions based on sparsity. Our proposed framework, Sparse-Gen, generalizes both sparse vector recovery and recovery using generative models by allowing for sparse deviations from the range of a generative model function. The benefits of using such modeling assumptions are observed both theoretically and empirically.

In the future, we would like to design algorithms that can better model the structure *within* sparse deviations. Follow-up work in this direction can benefit from the vast body of prior work in structured sparse vector recovery (Duarte & Eldar, 2011). From a theoretical perspective, a better understanding of the non-convexity resulting from generative model-based recovery can lead to stronger guarantees and consequently better optimization algorithms for recovery. Finally, it would be interesting to extend Sparse-Gen for compressed sensing of other data modalities such as graphs for applications in network tomography and reconstruction (Xu et al., 2011). Real-world graph networks are typically sparse in the canonical basis and can be modeled effectively using deep generative models (Grover et al., 2018), which is consistent with the modeling assumptions of the Sparse-Gen framework.

Acknowledgements

We are thankful to Tri Dao, Jonathan Kuck, Daniel Levy, Aditi Raghunathan, and Yang Song for helpful comments on early drafts. This research was supported by Intel Corporation, TRI, a Hellman Faculty Fellowship, ONR, NSF (#1651565, #1522054, #1733686) and FLI (#2017-158687). AG is supported by a Microsoft Research PhD Fellowship.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- Andersen, M. S., Dahl, J., and Vandenberghe, L. Cvxopt: A python package for convex optimization, version 1.1. 6. Available at cvxopt.org, 54, 2013.
- Babacan, S. D., Molina, R., and Katsaggelos, A. K. Bayesian compressive sensing using Laplace priors. *IEEE Transactions on Image Processing*, 19(1):53–63, 2010.
- Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- Baraniuk, R. G. and Wakin, M. B. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009.
- Baraniuk, R. G., Cevher, V., Duarte, M. F., and Hegde, C. Model-based compressive sensing. *IEEE Transactions on Information Theory*, 56(4):1982–2001, 2010.
- Baron, D., Sarvotham, S., and Baraniuk, R. G. Bayesian compressive sensing via belief propagation. *IEEE Transactions on Signal Processing*, 58(1):269–280, 2010.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. Simultaneous analysis of lasso and Dantzig selector. *The Annals of Statistics*, pp. 1705–1732, 2009.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models. In *International Conference on Machine Learning*, 2017.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Candès, E. J. and Tao, T. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- Candès, E. J., Romberg, J., and Tao, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- Chang, J. R., Li, C.-L., Póczos, B., Kumar, B. V., and Sankaranarayanan, A. C. One network to solve them all: solving linear inverse problems using deep projection models. *arXiv preprint*, 2017.
- Chen, M., Silva, J., Paisley, J., Wang, C., Dunson, D., and Carin, L. Compressive sensing on manifolds using a nonparametric mixture of factor analyzers: Algorithm and performance bounds. *IEEE Transactions on Signal Processing*, 58(12):6140–6155, 2010.
- Cohen, A., Dahmen, W., and DeVore, R. Compressed sensing and best k -term approximation. *Journal of the American mathematical society*, 22(1):211–231, 2009.
- Donoho, D. L. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- Duarte, M. F. and Eldar, Y. C. Structured compressed sensing: From theory to applications. *IEEE Transactions on Signal Processing*, 59(9):4053–4085, 2011.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- Gilbert, A. C., Zhang, Y., Lee, K., Zhang, Y., and Lee, H. Towards understanding the invertibility of convolutional neural networks. *arXiv preprint arXiv:1705.08664*, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- Grover, A., Zweig, A., and Ermon, S. Graphite: Iterative generative modeling of graphs. *arXiv preprint arXiv:1803.10459*, 2018.
- Hand, P. and Voroninski, V. Global guarantees for enforcing deep generative priors by empirical risk. *arXiv preprint arXiv:1705.07576*, 2017.
- He, L. and Carin, L. Exploiting structure in wavelet-based bayesian compressive sensing. *IEEE Transactions on Signal Processing*, 57(9):3488–3497, 2009.
- Ji, S., Xue, Y., and Carin, L. Bayesian compressive sensing. *IEEE Transactions on Signal Processing*, 56(6):2346–2356, 2008.

- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Kingma, D. and Welling, M. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Kulkarni, K., Lohit, S., Turaga, P., Kerviche, R., and Ashok, A. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- LeCun, Y., Cortes, C., and Burges, C. J. MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist>, 2010.
- Lu, X., Dong, W., Wang, P., Shi, G., and Xie, X. Convcsnet: A convolutional compressive sensing framework based on deep learning. *arXiv preprint arXiv:1801.10342*, 2018.
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. Compressed sensing MRI. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- Mallat, S. *A wavelet tour of signal processing: the sparse way*. Academic press, 2008.
- Mardani, M., Gong, E., Cheng, J. Y., Vasanawala, S., Zaharchuk, G., Alley, M., Thakur, N., Han, S., Dally, W., Pauly, J. M., et al. Deep generative adversarial networks for compressed sensing automates MRI. *arXiv preprint arXiv:1706.00051*, 2017.
- Mardani, M., Monajemi, H., Pappas, V., Vasanawala, S., Donoho, D., and Pauly, J. Recurrent generative adversarial networks for proximal learning and automated compressive image recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Matouek, J. *Lectures on Discrete Geometry - 2002*. Springer Science and Business Media, 2002.
- Mousavi, A., Patel, A. B., and Baraniuk, R. G. A deep learning approach to structured signal recovery. In *Annual Allerton Conference on Communication, Control, and Computing*, 2015.
- Peyre, G. Best basis compressed sensing. *IEEE Transactions on Signal Processing*, 58(5):2613–2622, 2010.
- Qaisar, S., Bilal, R. M., Iqbal, W., Naureen, M., and Lee, S. Compressive sensing: From theory to applications, a survey. *Journal of Communications and networks*, 15(5): 443–456, 2013.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Tang, G., Bhaskar, B. N., Shah, P., and Recht, B. Compressed sensing off the grid. *IEEE Transactions on Information Theory*, 59(11):7465–7490, 2013.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- Tieleman, T. and Hinton, G. Lecture 6.5-RMSProp: Divide the gradient by a running average of its recent magnitude. *Coursera: Neural networks for machine learning*, 4(2): 26–31, 2012.
- Xu, W., Mallada, E., and Tang, A. Compressive sensing over graphs. In *IEEE INFOCOM*, 2011.
- Yu, G. and Sapiro, G. Statistical compressed sensing of Gaussian mixture models. *IEEE Transactions on Signal Processing*, 59(12):5842–5858, 2011.
- Zhang, X., Wang, D., Zhou, Z., and Ma, Y. Simultaneous rectification and alignment via robust recovery of low-rank tensors. In *Advances in Neural Information Processing Systems*, 2013.
- Zhang, X., Zhou, Z., Wang, D., and Ma, Y. Hybrid singular value thresholding for tensor completion. In *AAAI Conference on Artificial Intelligence*, 2014.