
Learning a Compressed Sensing Measurement Matrix via Gradient Unrolling

Shanshan Wu¹ Alexandros G. Dimakis¹ Sujay Sanghavi¹ Felix X. Yu² Daniel Holtmann-Rice²
Dmitry Storcheus² Afshin Rostamizadeh² Sanjiv Kumar²

Abstract

Linear encoding of sparse vectors is widely popular, but is commonly data-independent – missing any possible extra (but a priori unknown) structure beyond sparsity. In this paper we present a new method to learn linear encoders that adapt to data, while still performing well with the widely used ℓ_1 decoder. The convex ℓ_1 decoder prevents gradient propagation as needed in standard gradient-based training. Our method is based on the insight that unrolling the convex decoder into T projected subgradient steps can address this issue. Our method can be seen as a data-driven way to learn a compressed sensing measurement matrix. We compare the empirical performance of 10 algorithms over 6 sparse datasets (3 synthetic and 3 real). Our experiments show that there is indeed additional structure beyond sparsity in the real datasets; our method is able to discover it and exploit it to create excellent reconstructions with fewer measurements (by a factor of 1.1-3x) compared to the previous state-of-the-art methods. We illustrate an application of our method in learning label embeddings for extreme multi-label classification, and empirically show that our method is able to match or outperform the precision scores of SLEEC, which is one of the state-of-the-art embedding-based approaches.

1. Introduction

Assume we have some unknown data vector $x \in \mathbb{R}^d$. We can observe only a few ($m < d$) linear equations of its entries and we would like to design these projections by creating a measurement matrix $A \in \mathbb{R}^{m \times d}$ such that the projections $y = Ax$ allow exact (or near-exact) recovery of the original vector $x \in \mathbb{R}^d$.

¹Department of Electrical and Computer Engineering, University of Texas at Austin, USA ²Google Research, New York, USA. Correspondence to: Shanshan Wu <shanshan@utexas.edu>.

If $d > m$, this is an ill-posed problem in general: we are observing m linear equations with d unknowns, so any vector x' on the subspace $Ax' = y$ satisfies our observed measurements. In this high-dimensional regime, the only hope is to make a structural assumption on x , so that unique reconstruction is possible. A natural approach is to assume that the data vector is sparse. The problem of designing measurement matrices and reconstruction algorithms that recover sparse vectors from linear observations is called Compressed Sensing (CS), Sparse Approximation or Sparse Recovery Theory (Donoho, 2006; Candès et al., 2006).

A natural way to recover is to search for the sparsest solution that satisfies the linear measurements:

$$\arg \min_{x' \in \mathbb{R}^d} \|x'\|_0 \quad \text{s.t. } Ax' = y. \quad (1)$$

Unfortunately this problem is NP-hard and for that reason the ℓ_0 norm is relaxed into an ℓ_1 -norm minimization¹

$$D(A, y) := \arg \min_{x' \in \mathbb{R}^d} \|x'\|_1 \quad \text{s.t. } Ax' = y. \quad (2)$$

Remarkably, if the measurement matrix A satisfies some properties, e.g. Restricted Isometry Property (RIP) (Candès, 2008) or the nullspace condition (NSP) (Rauhut, 2010), it is possible to prove that the ℓ_1 minimization in (2) produces the same output as the intractable ℓ_0 minimization in (1). However, it is NP-hard to test whether a matrix satisfies RIP (Bandeira et al., 2013).

In this paper we are interested in vectors that are not only sparse but have *additional structure* in their support. This extra structure may not be known or a-priori specified. We propose a data-driven algorithm that *learns a good linear measurement matrix A from data samples*. Our linear measurements are subsequently decoded with the ℓ_1 -minimization in (2) to estimate the unknown vector x .

Many real-world sparse datasets have additional structures beyond simple sparsity. For example, in a demographic dataset with (one-hot encoded) categorical features, a person's income level may be related to his/her education. Similarly, in a text dataset with bag-of-words representation, it is much more likely for two related words (e.g., football and

¹Other examples are greedy (Tropp & Gilbert, 2007), and iterative algorithms, e.g., CoSaMP (Needell & Tropp, 2009), IHT (Blumensath & Davies, 2009), and AMP (Donoho et al., 2009).

game) to appear in the same document than two unrelated words (e.g., football and biology). A third example is that in a genome dataset, certain groups of gene features may be correlated. In this paper, the goal is to *learn* a measurement matrix A to leverage such additional structure.

Our method is an autoencoder for sparse data, with a linear encoder (the measurement matrix) and a complex non-linear decoder that approximately solves an optimization problem. The latent code is the measurement $y \in \mathbb{R}^m$ which forms the bottleneck of the autoencoder that makes the representation interesting. A popular data-driven dimensionality reduction method is Principal Components Analysis (PCA) (see e.g., Hotelling 1933; Boutsidis et al. 2015; Wu et al. 2016; Li et al. 2017). PCA is also an autoencoder where both the encoder and decoder are linear and learned from samples. Given a data matrix $X \in \mathbb{R}^{n \times d}$ (each row is a sample), PCA projects each data point $x \in \mathbb{R}^d$ onto the subspace spanned by the top right-singular vectors of X . As is well-known, PCA provides the lowest mean-squared error when used with a linear decoder. However, when the data is sparse, non-linear recovery algorithms like (2) can yield significantly better recovery performance.

In this paper, we focus on learning a linear encoder for sparse data. Compared to non-linear embedding methods such as kernel PCA (Mika et al., 1999), a linear method enjoys two broad advantages: 1) it is easy to compute, as it only needs a matrix-vector multiplication; 2) it is easy to interpret, as every column of the encoding matrix can be viewed as a feature embedding. Interestingly, Arora et al. (2018) recently observed that the pre-trained word embeddings such as GloVe and word2vec (Mikolov et al., 2013; Pennington et al., 2014) formed a good measurement matrix for text data. Those measurement matrices, when used with ℓ_1 -minimization, need fewer measurements than the random matrices to achieve near-perfect recovery.

Given a sparse dataset that has additional (but unknown) structure, our goal is to learn a good measurement matrix A , when the recovery algorithm is the ℓ_1 -minimization in (2). More formally, given n sparse samples $x_1, x_2, \dots, x_n \in \mathbb{R}^d$, our problem of finding the best A can be formulated as

$$\min_{A \in \mathbb{R}^{m \times d}} f(A), \text{ where } f(A) := \sum_{i=1}^n \|x_i - D(A, Ax_i)\|_2^2.$$

Here $D(\cdot, \cdot)$ is the ℓ_1 decoder defined in (2). Unfortunately, there is no easy way to compute the gradient of $f(A)$ with respect to A , due to the optimization problem defining $D(\cdot, \cdot)$. Our main insight, which will be elaborated on in Section 3.1, is that **replacing the ℓ_1 -minimization with a T -step projected subgradient update of it**, results in gradients being (approximately) computable. In particular, consider the approximate function $\tilde{f}(A) : \mathbb{R}^{m \times d} \mapsto \mathbb{R}$ defined as

$$\begin{aligned} \tilde{f}(A) &:= \sum_{i=1}^n \|x_i - \hat{x}_i\|_2^2, \quad \text{where} \\ \hat{x}_i &= T\text{-step projected subgradient of} \\ &D(A, Ax_i), \text{ for } i = 1, \dots, n. \end{aligned} \quad (3)$$

As we will show, now it is *possible* to compute the gradients of $\tilde{f}(A)$ with respect to A . This idea is sometimes called *unrolling* and has appeared in various other applications as we discuss in the related work section. To the best of our knowledge, we are the first to use unrolling to learn a measurement matrix for compressed sensing.

Our contributions can be summarized as follows:

- We design a novel autoencoder, called ℓ_1 -AE, to learn an efficient and compressed representation for structured sparse vectors. Our autoencoder is easy to train. It has only two tuning hyper-parameters in the network architecture: the encoding dimension m and the network depth T . The architecture of ℓ_1 -AE is inspired by the projected subgradient method of solving the ℓ_1 decoder in (2). While the exact projected subgradient method requires computing the pseudoinverse, we circumvent this by observing that it is possible to replace the expensive pseudoinverse operation by a simple transpose (Lemma 1).
- The most surprising result in this paper is that we can learn a linear encoder using an unfolded T -step projected subgradient decoder and the learned measurement matrix *still* performs well for the original ℓ_1 -minimization decoder. We empirically compare 10 algorithms over 6 sparse datasets (3 synthetic and 3 real). As shown in Figure 2, using the measurement matrix learned from our autoencoder, we can compress the sparse vectors (in the test set) to a lower dimension (by a factor of 1.1-3x) than random Gaussian matrices while still being able to *perfectly* recover the original sparse vectors (see also Table 2). This demonstrates the superior ability of our autoencoder in learning and adapting to the additional structures in the given data.
- Although our autoencoder is specifically designed for ℓ_1 -minimization decoder, the learned measurement matrix also performs well (and can perform even better) with the model-based decoder (Baraniuk et al., 2010) (Figure 3). This further demonstrates the benefit of learning a measurement matrix from data. As a baseline algorithm, we slightly modify the original model-based CoSaMP algorithm by adding a positivity constraint without changing the theoretical guarantee (Appendix C), which could be of independent interest.
- Besides the application in compressed sensing, one interesting direction for future research is to use the proposed ℓ_1 -AE in other supervised tasks. We illustrate a potential application of ℓ_1 -AE in extreme multi-label learning.

We show that ℓ_1 -AE can be used to learn label embeddings for multi-label datasets. We compare the resulted method with one of the state-of-the-art embedding-based methods SLEEC (Bhatia et al., 2015) over two benchmark datasets. Our method is able to achieve better or comparable precision scores than SLEEC (see Table 3).

2. Related Work

We briefly review the relevant work, and highlight the differences compared to our paper.

Model-based compressed sensing (CS). Model-based CS (Baraniuk et al., 2010; Hegde et al., 2014) extends the conventional compressed sensing theory by considering more realistic structured models than simple sparsity. It requires to know the sparsity structure a priori, which is not always possible in practice. Our approach, by contrast, does not require a priori knowledge about the sparsity structure.

Learning-based measurement design. Most theoretical results in CS are based on random measurement matrices. There are a few approaches proposed to learn a measurement matrix from training data. One approach is to learn a near-isometric embedding that preserves pairwise distance (Hegde et al., 2015; Bah et al., 2013). This approach usually requires computing the pairwise difference vectors, and hence is computationally expensive if both the number of training samples and the dimension are large (which is the setting that we are interested in). Another approach restricts the form of the measurement matrix, e.g., matrices formed by a subset of rows of a given basis matrix. The learning problem then becomes selecting the best subset of rows (Baldassarre et al., 2016; Li & Cevher, 2016; Gözcü et al., 2018). In Figure 2, we compare our method with the learning-based subsampling method proposed in (Baldassarre et al., 2016), and show that our method needs fewer measurements to recover the original sparse vector.

Adaptive CS. In adaptive CS (Seeger & Nickisch, 2011; Arias-Castro et al., 2013; Malloy & Nowak, 2014), the new measurement is designed based on the previous measurements in order to maximize the gain of new information. This is in contrast to our setting, where we are given a set of training samples. Our goal is to learn a good measurement matrix to leverage additional structure in the given samples.

Dictionary learning. Dictionary learning (Aharon et al., 2006; Mairal et al., 2009) is the problem of learning an overcomplete set of basis vectors from data so that every datapoint (presumably dense) can be represented as a sparse linear combination of the basis vectors. By contrast, this paper focuses on learning a good measurement matrix for data that are already sparse in the canonical basis.

Sparse coding. The goal of sparse coding (Olshausen & Field, 1996; Donoho & Elad, 2003) is to find the sparsest representation of a dense vector, given a fixed family of basis vectors (aka a dictionary). Training a deep neural network to compute the sparse codes becomes popular recently (Gregor & LeCun, 2010; Sprechmann et al., 2015; Wang et al., 2016). Several recent papers (Kulkarni et al., 2016; Xin et al., 2016; Shi et al., 2017; Jin et al., 2017; Mardani et al., 2017; Mousavi et al., 2015; 2017; 2019; Mousavi & Baraniuk, 2017; He et al., 2017; Zhang & Ghanem, 2017; Lohit et al., 2018) proposes new convolutional architectures for image reconstruction from low-dimensional measurements. Some of the networks also have an image sensing component. Sparse coding is different from this paper, where our focus is on learning a good measurement/encoding matrix rather than learning a good recovery/decoding algorithm.

Autoencoders. An autoencoder is a popular neural network architecture for unsupervised learning. It has applications in dimensionality reduction (Hinton & Salakhutdinov, 2006), pre-training (Erhan et al., 2010), image compression and recovery (Lohit et al., 2018; Mousavi et al., 2015; 2017; 2019), denoising (Vincent et al., 2010), and generative modeling (Kingma & Welling, 2014). In this paper we design a novel autoencoder ℓ_1 -AE to learn a compressed sensing measurement matrix for the ℓ_1 -minimization decoder.

Unrolling. The idea of unfolding an iterative algorithm into a neural network is a natural way to make the algorithm trainable (Gregor & LeCun, 2010; Hershey et al., 2014; Sprechmann et al., 2015; Xin et al., 2016; Wang et al., 2016; Shi et al., 2017; Jin et al., 2017; Mardani et al., 2017; He et al., 2017; Zhang & Ghanem, 2017; Mardani et al., 2018; Kellman et al., 2019). The main difference between the previous papers and this paper is that, most previous papers seek a trained neural network as a replacement of the original optimization-based algorithm, while in this paper we design an autoencoder based on the unrolling idea, and after training, we show that the learned measurement matrix still performs well with the *original* ℓ_1 -minimization decoder.

Extreme multi-label learning (XML). The goal of XML is to learn a classifier to identify (for each datapoint) a subset of relevant labels from a extreme large label set. Different approaches have been proposed for XML, e.g., embedding-based (Bhatia et al., 2015; Yu et al., 2014; Mineiro & Karampatziakis, 2015; Tagami, 2017), tree-based (Prabhu & Varma, 2014; Jain et al., 2016; Jasinska et al., 2016; Prabhu et al., 2018a), one-vs-all (Prabhu et al., 2018b; Babbar & Schölkopf, 2017; Yen et al., 2016; 2017; Niculescu-Mizil & Abbasnejad, 2017; Hariharan et al., 2012), and deep learning (Jernite et al., 2017; Liu et al., 2017). Here we focus on the embedding-based approach. In Section 5 we show that the proposed autoencoder can be used to learn label embeddings for multi-label datasets.

3. Our Algorithm

Our goal is to learn a measurement matrix A from the given sparse vectors. This is done via training a novel autoencoder, called ℓ_1 -AE. In this section, we will describe the key idea behind the design of ℓ_1 -AE. In this paper we focus on the vectors that are sparse in the standard basis and also non-negative². This is a natural setting for many real-world datasets, e.g., categorical data and bag-of-words data.

3.1. Intuition

Our design is strongly motivated by the projected subgradient method used to solve an ℓ_1 -minimization problem. Consider the following ℓ_1 -minimization problem:

$$\min_{x' \in \mathbb{R}^d} \|x'\|_1 \quad \text{s.t. } Ax' = y, \quad (4)$$

where $A \in \mathbb{R}^{m \times d}$ and $y \in \mathbb{R}^m$. We assume that $m < d$ and that A has rank m . One approach³ to solving (4) is the projected subgradient method. The update is given by

$$x^{(t+1)} = \Pi(x^{(t)} - \alpha_t g^{(t)}), \quad \text{where } g^{(t)} = \text{sign}(x^{(t)}) \quad (5)$$

where Π denotes the (Euclidean) projection onto the convex set $\{x' : Ax' = y\}$, $g^{(t)}$ is the sign function, i.e., the subgradient of the objective function $\|\cdot\|_1$ at $x^{(t)}$, and α_t is the step size at the t -th iteration. Since A has full row rank, Π has a closed-form solution given by

$$\Pi(z) = \arg \min_h \|h - z\|_2^2 \quad \text{s.t. } Ah = y \quad (6)$$

$$= z + \arg \min_{h'} \|h' - z\|_2^2 \quad \text{s.t. } Ah' = y - Az \quad (7)$$

$$= z + A^\dagger(y - Az), \quad (8)$$

where $A^\dagger = A^T(AA^T)^{-1}$ is the Moore-Penrose inverse of matrix A . Substituting (8) into (5), and using the fact that $Ax^{(t)} = y$, we get the following update equation

$$x^{(t+1)} = x^{(t)} - \alpha_t(I - A^\dagger A)\text{sign}(x^{(t)}), \quad (9)$$

where I is the identity matrix. We use $x^{(1)} = A^\dagger y$ (which satisfies the constraint $Ax' = y$) as the starting point.

As mentioned in the Introduction, our main idea is to replace the solution of an ℓ_1 decoder given in (4) by a T -step projected subgradient update given in (9). One technical difficulty⁴ in simulating (9) using neural networks is back-propagating through the pseudoinverse A^\dagger . Fortunately, Lemma 1 shows that it is possible to replace A^\dagger by A^T .

Lemma 1. *For any vector $x \in \mathbb{R}^d$, and any matrix $A \in \mathbb{R}^{m \times d}$ ($m < d$) with rank m , there exists an $\tilde{A} \in \mathbb{R}^{m \times d}$*

²Extending our method to general cases is left for future work.

³Another approach is via linear programming.

⁴One approach is to replace A^\dagger by a trainable matrix $B \in \mathbb{R}^{d \times m}$. This approach performs worse than ours (see Figure 3).

with all singular values being ones, such that the following two ℓ_1 -minimization problems have the same solution:

$$P_1 : \quad \min_{x' \in \mathbb{R}^d} \|x'\|_1 \quad \text{s.t. } Ax' = Ax. \quad (10)$$

$$P_2 : \quad \min_{x' \in \mathbb{R}^d} \|x'\|_1 \quad \text{s.t. } \tilde{A}x' = \tilde{A}x. \quad (11)$$

Furthermore, the projected subgradient update of P_2 is

$$x^{(t+1)} = x^{(t)} - \alpha_t(I - \tilde{A}^T \tilde{A})\text{sign}(x^{(t)}), \quad x^{(1)} = \tilde{A}^T \tilde{A}x. \quad (12)$$

A natural choice for \tilde{A} is $U(AA^T)^{-1/2}A$, where $U \in \mathbb{R}^{m \times m}$ can be any unitary matrix.

Lemma 1 essentially says that: 1) Instead of searching over all matrices (of size m -by- d with rank m), it is enough to search over a subset of matrices \tilde{A} , whose singular values are all ones. This is because A and \tilde{A} has the same recovery performance for ℓ_1 -minimization (this is true as long as \tilde{A} and A have the same null space). 2) The key benefit of searching over matrices with singular values being all ones is that the corresponding projected subgradient update has a simpler form: the annoying pseudoinverse term A^\dagger in (9) is replaced by a simple matrix transpose A^T in (12).

As we will show in the next section, our decoder is designed to simulate (12) instead of (9): the only difference is that the pseudoinverse term A^\dagger is replaced by matrix transpose A^T . Ideally we should train our ℓ_1 -AE by enforcing the constraint that the matrices have singular values being ones. In practice, we do not enforce that constraint during training⁵. We empirically observe that the learned measurement matrix A is not far from the constraint set (see Appendix D.5).

3.2. Network Structure of ℓ_1 -AE

As shown in Figure 1, ℓ_1 -AE has a simple linear encoder and a non-linear decoder. When a data point $x \in \mathbb{R}^d$ comes, it is encoded as $y = Ax$, where $A \in \mathbb{R}^{m \times d}$ is the encoding matrix that will be learned from data. A decoder is then used to recover the original vector x from its embedding y .

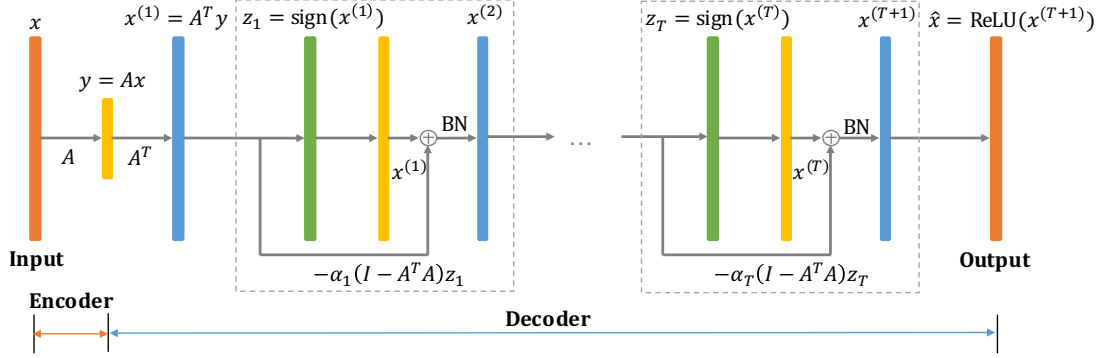
The decoder network of ℓ_1 -AE consists of T blocks connected in a feedforward manner: the output vector of the t -th block is the input vector to the $(t+1)$ -th block. The network structure inside each block is identical. Let $x^{(1)} = A^T y$. For $t \in \{1, 2, \dots, T\}$, if $x^{(t)} \in \mathbb{R}^d$ is the input to the t -th block, then its output vector $x^{(t+1)} \in \mathbb{R}^d$ is

$$x^{(t+1)} = x^{(t)} - \alpha_t(I - A^T A)\text{sign}(x^{(t)}), \quad (13)$$

where $\alpha_1, \alpha_2, \dots, \alpha_T \in \mathbb{R}$ are scalar variables to be learned from data. We empirically observe that regularizing α_t to have the following form⁶ $\alpha_t = \beta/t$ for $t \in \{1, 2, \dots, T\}$

⁵Efficient training with this manifold constraint (Meghwanshi et al., 2018) is an interesting direction for future work. We empirically found that adding a regularizer $\|I - AA^T\|_2$ to the loss function would degrade the performance.

⁶It is called *square summable but not summable* (Boyd, 2014).


 Figure 1. Network structure of the proposed autoencoder ℓ_1 -AE .

improves test accuracy. Here, $\beta \in \mathbb{R}$ is the only scalar variable to be learned from data. We also add a standard batch normalization (BN) layer (Ioffe & Szegedy, 2015) inside each block, because empirically it improves the test accuracy (see Figure 3). After T blocks, we use rectified linear units (ReLU) in the last layer⁷ to obtain the final output $\hat{x} \in \mathbb{R}^d$: $\hat{x} = \text{ReLU}(x^{(T+1)})$.

It is worth noting that the low-rank structure of the weight matrix $I - A^T A$ in (13) is essential for reducing the computational complexity. A fully-connected layer requires a weight matrix of size $d \times d$, which is intractable for large d .

Given n unlabeled training examples $\{x_i\}_{i=1}^n$, we will train an ℓ_1 -AE to minimize the average squared ℓ_2 distance between $x \in \mathbb{R}^d$ and $\hat{x} \in \mathbb{R}^d$:

$$\min_{A \in \mathbb{R}^{m \times d}, \beta \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \|x_i - \hat{x}_i\|_2^2. \quad (14)$$

4. Experiments

We implement ℓ_1 -AE in Tensorflow. Our code is available online: <https://github.com/wushanshan/L1AE>. In this section, we demonstrate that ℓ_1 -AE is able to learn a good measurement matrix A for the structured sparse datasets, when decoding is done by ℓ_1 -minimization⁸.

4.1. Datasets and Training

Synthetic datasets. As shown in Table 1, we generate three synthetic datasets⁹: two satisfy the block sparsity model¹⁰ (Baraniuk et al., 2010), and one follows the power-law structured sparsity (The i -th coordinate is nonzero with

⁷This makes sense as we focus on non-negative sparse vectors in this paper. Non-negativity is a natural setting for many real-world sparse datasets, e.g., categorical data and text data.

⁸We use Gurobi (a commercial optimization solver) to solve it.

⁹We also have experiments on a synthetic dataset with no extra structure (see Appendix D.6).

¹⁰A signal $x \in \mathbb{R}^d$ is called K -block sparse with block size J if it satisfies: 1) x can be reshaped into a matrix X of size $J \times N$, where $JN = d$; 2) every column of X acts as a group, i.e., the entire column is either zero or nonzero; 3) X has K nonzero columns, and hence x has sparsity KJ .

probability proportional to $1/i$). Each sample is generated as follows: 1) choose a random support set satisfying the sparsity model; 2) set the nonzeros to be uniform in $[0, 1]$.

Real datasets. Our first dataset is from Kaggle ‘‘Amazon Employee Access Challenge’’¹¹. Each training example contains 9 categorical features. We use one-hot encoding to transform each example. Our second dataset Wiki10-31K is a multi-label dataset downloaded from this repository (Bhatia et al., 2017). We only use the label vectors to train our autoencoder. Our third dataset is RCV1 (Lewis et al., 2004), a popular text dataset. We use scikit-learn to fetch the training set and randomly split it into train/validation/test sets.

Training. We use stochastic gradient descent to train the autoencoder. Before training, every sample is normalized to have unit ℓ_2 norm. The parameters are initialized as follows: $A \in \mathbb{R}^{m \times d}$ is a random Gaussian matrix with standard deviation $1/\sqrt{d}$; β is initialized as 1.0. Other hyperparameters are given in Appendix B. A single NVIDIA Quadro P5000 GPU is used in the experiments. We set the decoder depth $T = 10$ for most datasets¹². Training an ℓ_1 -AE can be done in several minutes for small-scale datasets and in around an hour for large-scale datasets.

4.2. Algorithms

We compare 10 algorithms in terms of their recovery performance. The results are shown in Figure 2. All the algorithms follow a two-step ‘‘encoding + decoding’’ process.

ℓ_1 -AE + ℓ_1 -min pos (our algorithm): After training an ℓ_1 -AE, we use the encoder matrix A as the measurement matrix. To decode, we use Gurobi (a commercial optimization solver) to solve the following ℓ_1 -minimization problem with positivity constraint (denoted as ‘‘ ℓ_1 -min pos’’):

$$\min_{x' \in \mathbb{R}^d} \|x'\|_1 \quad \text{s.t. } Ax' = y, x' \geq 0. \quad (15)$$

¹¹<https://www.kaggle.com/c/amazon-employee-access-challenge>

¹²Although the subgradient method (Boyd, 2014) has $1/\epsilon^2$ convergence, in practice, we found that a small value of T (e.g., $T = 10$ in Table 2) was good enough (Giryas et al., 2018).

Table 1. Summary of the datasets. The validation set is used for parameter tuning and early stopping.

Dataset	Dimension	Avg. no. of nonzeros	Train / Valid / Test Size	Description
Synthetic1	1000	10	6000 / 2000 / 2000	1-block sparse with block size 10
Synthetic2	1000	10	6000 / 2000 / 2000	2-block sparse with block size 5
Synthetic3	1000	10	6000 / 2000 / 2000	Power-law structured sparsity
Amazon	15626	9	19661 / 6554 / 6554	1-hot encoded categorical data
Wiki10-31K	30398	19	14146 / 3308 / 3308	Extreme multi-label data
RCV1	47236	76	13889 / 4630 / 4630	Text data with TF-IDF features

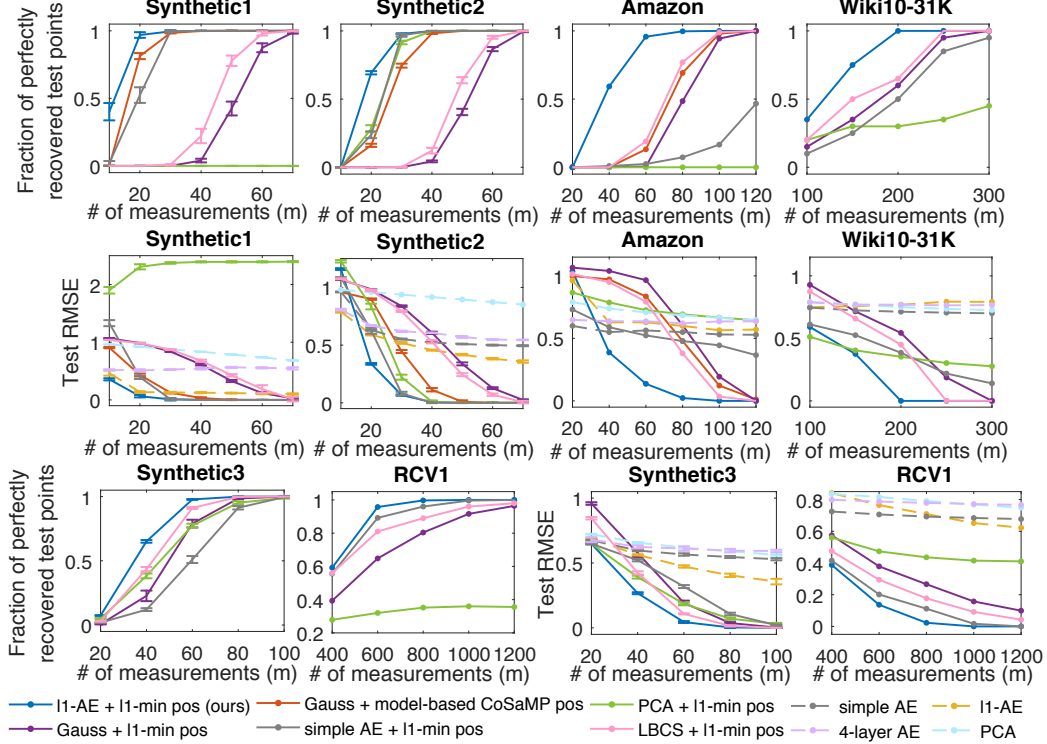


Figure 2. Best viewed in color. Recovery performance over the test set: fraction of exactly recovered data points (the 1st row and the left half of the 3rd row); reconstruction error (the 2nd row and the right half of the 3rd row). For the following four baselines, we do not plot the fraction of recovered points: simple AE, 4-layer AE, ℓ_1 -AE, and PCA. This is because they cannot produce a perfect reconstruction \hat{x} that satisfies $\|x - \hat{x}\|_2 \leq 10^{-10}$ (see also Table 2). For synthetic data, we plot the mean and standard deviation (indicated by the error bars) across 10 randomly generated datasets. Note that model-based CoSaMP decoder is not applicable for the Synthetic3, Wiki10-31K, and the RCV1 datasets. Although the model-based CoSaMP decoder has more information about the given data (such as block sparsity and one-hot sparsity) than the ℓ_1 -minimization decoder, our “ ℓ_1 -AE + ℓ_1 -min pos” still gives the best recovery performance.

Since we focus on non-negative sparse vectors in this paper, adding a positivity constraint improves the recovery performance¹³ (see Appendix D.4).

Gauss + ℓ_1 -min pos / model-based CoSaMP pos: A random Gaussian matrix $G \in \mathbb{R}^{m \times d}$ with i.i.d. $\mathcal{N}(0, 1/m)$ entries is used as the measurement matrix¹⁴. We experiment with two decoding methods: 1) Solve the optimization problem given in (15); 2) Use the model-based CoSaMP algorithm¹⁵ (Algorithm 1 in Baraniuk et al. 2010) with an

¹³The sufficient and necessary condition (Theorem 3.1 of Khajepnejad et al. 2011) for exact recovery using (15) is weaker than the nullspace property (NSP) (Rauhut, 2010) for (4).

¹⁴Additional experiments with random partial Fourier matrices (Haviv & Regev, 2017) can be found in Appendix D.2.

¹⁵Model-based method needs the explicit sparsity model, and hence is not applicable for RCV1, Wiki10-31K, and Synthetic3.

additional positivity constraint (see Appendix C).

PCA or PCA + ℓ_1 -min pos: We perform truncated singular value decomposition (SVD) on the training set. Let $A \in \mathbb{R}^{m \times d}$ be the top- m singular vectors. For PCA, every vector $x \in \mathbb{R}^d$ in the test set is estimated as $A^T A x$. We can also use “ ℓ_1 -min pos” (15) as the decoder.

Simple AE or Simple AE + ℓ_1 -min pos: We train a simple autoencoder: for an input vector $x \in \mathbb{R}^d$, the output is $\text{ReLU}(B^T A x) \in \mathbb{R}^d$, where both $B \in \mathbb{R}^{m \times d}$ and $A \in \mathbb{R}^{m \times d}$ are learned from data. We use the same loss function as our autoencoder. After training, we use the learned A matrix as the measurement matrix. Decoding is performed either by the learned decoder or solving “ ℓ_1 -min pos” (15).

LBCS + ℓ_1 -min pos / model-based CoSaMP pos: We implement the learning-based compressive subsampling

(LBCS) method in (Baldassarre et al., 2016). The idea is to select a subset (of size m) of coordinates (in the transformed space) that preserves the most energy. We use Gaussian matrix as the basis matrix and “ ℓ_1 -min pos” as the decoder¹⁶. Decoding with “model-based CoSaMP pos” is in Figure 3.

4-layer AE: We train a standard 4-layer autoencoder (we do not count the input layer), whose encoder network (and decoder network) consists of two feedforward layers with ReLU activation. The dimension of the 1st (and 3rd) layer is tuned based on the performance on the validation set.

4.3. Results and Analysis

The experimental results are shown in Figure 2. Two performance metrics are compared. The first one is the fraction of exactly recovered test samples. We say that a vector x is exactly recovered by an algorithm if it produces an \hat{x} that satisfies $\|x - \hat{x}\|_2 \leq 10^{-10}$. The second metric is the root mean-squared error (RMSE) over the test set¹⁷.

As shown in Figure 2, our algorithm “ ℓ_1 -AE + ℓ_1 -min pos” outperforms the rest baselines across all datasets. By learning a data-dependent linear encoding matrix, our method requires fewer measurements to achieve perfect recovery.

Learned decoder versus ℓ_1 -min decoder. We now compare two methods: ℓ_1 -AE and ℓ_1 -AE + ℓ_1 -min pos. They have a common encoder but different decoders. As shown in Figure 2, “ ℓ_1 -AE + ℓ_1 -min pos” almost always gives a smaller RMSE. In fact, as shown in Table 2, “ ℓ_1 -min pos” is able to achieve reconstruction errors in the order of $1e-15$, which is impossible for a neural network. The strength of optimization-based decoder over neural network-based decoder has been observed before, e.g., see Figure 1 in (Bora et al., 2017)¹⁸. Nevertheless, neural network-based decoder usually has an advantage that it can handle nonlinear encoders, for which the corresponding optimization problem may become non-convex and hard to solve exactly.

Model-based decoder versus learned encoder. It is interesting to see that our algorithm even outperforms model-based method (Baraniuk et al., 2010), even though the model-based decoder has more information about the given data than ℓ_1 -minimization decoder. For the Amazon dataset, compared to “Gauss + model-based CoSaMP pos”, our method reduces the number of measurements needed for exact recovery by about 2x. This is possible because the model-

based decoder only knows that the input vector comes from one-hot encoding, which is a *coarse* model for the underlying sparsity model. By contrast, ℓ_1 -AE learns a measurement matrix directly from the given training data.

A natural question to ask is whether the measurement matrix learned by ℓ_1 -AE can improve the recovery performance of the model-based decoding algorithm. As shown in Figure 3, the recovery performance of “ ℓ_1 -AE + model-based CoSaMP pos” is better than “Gauss + model-based CoSaMP pos”. This further demonstrates the benefit of learning a data-adaptive measurement matrix.

Variations of ℓ_1 -AE. We now examine how slightly varying the decoder structure would affect the performance. We make the following changes to the decoder structure: 1) remove the Batch Normalization layer; 2) remove the ReLU operation in the last layer; 3) change the nonlinearity from sign to tanh; 4) replace the A^T term in the decoder network by a matrix $B \in \mathbb{R}^{d \times m}$ that is learned from data; 5) use one-layer neural network as the decoder, i.e., set $T = 0$ in ℓ_1 -AE. Each change is applied in isolation. As shown in Figure 3, our design gives the best recovery performance among all the variations for the Amazon dataset¹⁹.

Decoder depth of ℓ_1 -AE. The decoder depth T is a tuning parameter of ℓ_1 -AE. Empirically, the performance of the learned matrix improves as T increases (see Table 2). On the other hand, the training time increases as T increases. The parameter T is tuned as follows: we start with $T = 5, 10, \dots, T_{\max}$, and stop if the validation performance improvement is smaller than a threshold or if T equals T_{\max} . The hyper-parameters used for training ℓ_1 -AE are given in Appendix B. We set $T = 5$ for Synthetic2 and Synthetic3, $T = 10$ for Synthetic1, RCV1, and Wiki10-31K, and $T = 60$ for Amazon dataset. The autoencoder is trained on a single GPU. Training an ℓ_1 -AE takes a few minutes for small datasets and around an hour for large datasets.

5. Extreme Multi-label Learning

We have proposed a novel autoencoder ℓ_1 -AE to learn a compressed sensing measurement matrix for high-dimensional sparse data. Besides the application in compressed sensing, one interesting direction for future research is to use ℓ_1 -AE in other machine learning tasks. Here we illustrate a potential application of ℓ_1 -AE in extreme multi-label learning (XML). For every input feature vector, the goal of XML is to predict a subset of relevant labels from a extremely large label set. As a result, the output label vector is high-dimensional, sparse, and non-negative (with 1’s denoting the relevant labels and 0’s otherwise).

¹⁹Besides these variations, we can design new autoencoders by unrolling other decoding algorithms such as ISTA (Gregor & LeCun, 2010). More discussions can be found in Appendix D.7.

¹⁶We tried four variations of LBCS: two different basis matrices (random Gaussian matrix and DCT matrix), two different decoders (ℓ_1 -minimization and linear decoder). The combination of Gaussian and ℓ_1 -minimization performs the best (see Appendix D.8).

¹⁷Note that in Figure 2, test RMSE has similar scale across all datasets, because the vectors are normalized to have unit ℓ_2 norm.

¹⁸As indicated by Figure 1 in (Bora et al., 2017), LASSO gives better reconstruction than GAN-based method when given enough Gaussian measurements.

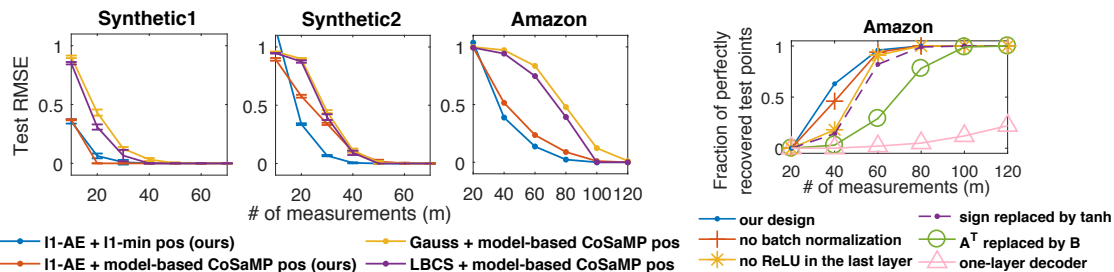


Figure 3. Best viewed in color. **Left three plots:** Although ℓ_1 -AE is designed for the ℓ_1 -min decoder, the matrix learned from ℓ_1 -AE also improves the recovery performance when the decoder is model-based CoSaMP. **Right plot:** Recovery performance of “ ℓ_1 -AE + ℓ_1 -min pos” on the Amazon test set when we slightly change the decoder structure. Each change is applied in isolation.

Table 2. **Left:** Comparison of test RMSE: ℓ_1 -AE versus ℓ_1 -AE + ℓ_1 -min pos. **Right:** Test RMSE of our method “ ℓ_1 -AE + ℓ_1 -min pos” on the synthetic1 dataset: the error decreases as the decoder depth T increases.

Dataset	Synthetic1			Amazon			# measurements	20	30
# measurements	10	30	50	40	80	120	$T = 10$	0.097	9.9e-3
ℓ_1 -AE	0.465	0.129	0.118	0.638	0.599	0.565	$T = 20$	0.063	1.2e-15
ℓ_1 -AE + ℓ_1 -min pos (ours)	0.357	9.9e-3	1.4e-15	0.387	0.023	2.8e-15	$T = 30$	9.6e-16	1.1e-15

Many approaches have been proposed for XML (Bhatia et al., 2017). Here we focus on the embedding-based approach, and one of the state-of-the-art embedding-based approaches²⁰ is SLEEC (Bhatia et al., 2015). Given n training samples (x_i, y_i) , $i = 1, \dots, n$, where $x_i \in \mathbb{R}^p$, $y_i \in \mathbb{R}^d$, we use $X \in \mathbb{R}^{p \times n}$ and $Y \in \mathbb{R}^{d \times n}$ to denote the stacked feature matrix and label matrix. SLEEC works in two steps. In Step 1, SLEEC reduces the dimension of the labels Y by learning a low-dimensional embedding for the training samples. Let $Z \in \mathbb{R}^{m \times n}$ ($m < d$) be the learned embedding matrix. Note that only the Y matrix is used for learning Z in this step. In Step 2, SLEEC learns a linear mapping $V \in \mathbb{R}^{m \times p}$ such that $Z \approx VX$. To predict the label vector for a new sample $x \in \mathbb{R}^p$, SLEEC uses nearest neighbors method: first computes the embedding Vx , identifies a few nearest neighbors (from the training set) in the embedding space, and uses the sum of their label vectors as prediction.

The method that we propose follows SLEEC’s two-step procedure. The main difference is that in Step 1, we train an autoencoder ℓ_1 -AE to learn embeddings for the labels Y . Note that in XML, the label vectors Y are high-dimensional, sparse, and non-negative. Let $A \in \mathbb{R}^{m \times d}$ be the learned measurement matrix for Y , then the embedding matrix is $Z = AY$. In Step 2, we use the same subroutine as SLEEC to learn a linear mapping from X to Z . To predict the label vector for a new sample, we compared three methods in our experiments: 1) use the nearest neighbor method (same as SLEEC); 2) use the decoder of the trained ℓ_1 -AE (which maps from the embedding space to label space); 3) use an average of the label vectors obtained from 1) and 2). The three methods are denoted as “ ℓ_1 -AE 1/2/3” in Table 3.

In Table 3, we compare the precision score P@1 over two benchmark datasets. The 2nd row is the number of

²⁰AnnexML (Tagami, 2017) is a graph-embedding approach for XML. Some of its techniques (such as better partition of the input data points) can be potentially used with our method.

models combined in the ensemble. According to (Bhatia et al., 2017), SLEEC achieves a precision score 0.7926 for EURLex-4K and 0.8588 for Wiki10-31K, which are consistent with what we obtained by running their code. The embedding dimensions are $m = 100$ for EURLex-4K and $m = 75$ for Wiki10-31K. We set $T = 10$ for ℓ_1 -AE. As shown in Table 3, our method has higher score than SLEEC for EURLex-4K. For Wiki10-31K, a single model of our method has higher score than SLEEC. When 3 or 5 models are ensemble, our method has comparable precision score to SLEEC. More results can be found in Appendix D.3.

Table 3. **Upper:** Description of the datasets. **Lower:** Comparisons of P@1 scores. The 2nd row is the no. of models in the ensemble.

Dataset	Feature dimension	Label dimension	Train/ Test	# labels /point		
EURLex-4K	5000	3993	15539/3809	5.31		
Wiki10-31K	101938	30938	14146/6616	18.64		
Dataset	EURLex-4K			Wiki10-31K		
# models	1	3	5	1	3	5
SLEEC	0.7600	0.7900	0.7944	0.8356	0.8603	0.8600
ℓ_1 -AE 1	0.7655	0.7928	0.7931	0.8529	0.8564	0.8597
ℓ_1 -AE 2	0.7949	0.8033	0.8070	0.8560	0.8579	0.8581
ℓ_1 -AE 3	0.8062	0.8151	0.8136	0.8617	0.8640	0.8630

6. Conclusion

Combining ideas from compressed sensing, convex optimization and deep learning, we proposed a novel unsupervised learning framework for high-dimensional sparse data. The proposed autoencoder ℓ_1 -AE is able to learn an efficient measurement matrix by adapting to the sparsity structure of the given data. The learned measurement matrices can be subsequently used in other machine learning tasks such as extreme multi-label learning. We expect that the learned ℓ_1 -AE can lead to useful representations in various supervised learning pipelines, for datasets that are well represented by large sparse vectors. Investigating the relation between the training data and the learned matrix (see Appendix D.1 for a toy example) is an interesting direction for future research.

7. Acknowledgements

This research has been supported by NSF Grants 1302435, 1564000, and 1618689, DMS 1723052, CCF 1763702, ARO YIP W911NF-14-1-0258 and research gifts by Google, Western Digital and NVIDIA.

References

- Aharon, M., Elad, M., and Bruckstein, A. k-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- Arias-Castro, E., Candès, E. J., and Davenport, M. A. On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory*, 59(1):472–481, 2013.
- Arora, S., Khodak, M., Nikunj, S., and Vodrahalli, K. A compressed sensing view of unsupervised text embeddings, bag-of-n-grams, and lstms. In *International Conference on Learning Representations (ICLR)*, 2018.
- Babbar, R. and Schölkopf, B. Dismec: distributed sparse machines for extreme multi-label classification. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pp. 721–729. ACM, 2017.
- Bah, B., Sadeghian, A., and Cevher, V. Energy-aware adaptive bi-lipschitz embeddings. In *Proceedings of the 10th International Conference on Sampling Theory and Applications*, 2013.
- Baldassarre, L., Li, Y.-H., Scarlett, J., Gözcü, B., Bogunovic, I., and Cevher, V. Learning-based compressive subsampling. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):809–822, 2016.
- Bandeira, A. S., Dobriban, E., Mixon, D. G., and Sawin, W. F. Certifying the restricted isometry property is hard. *IEEE transactions on information theory*, 59(6):3448–3450, 2013.
- 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.
- Bhatia, K., Jain, H., Kar, P., Varma, M., and Jain, P. Sparse local embeddings for extreme multi-label classification. In *Advances in Neural Information Processing Systems*, pp. 730–738, 2015.
- Bhatia, K., Kunal, D., Jain, H., Prabhu, Y., and Varma, M. The extreme classification repository: Multi-label datasets and code, 2017. URL <http://manikvarma.org/downloads/XC/XMLRepository.html>. [Online; last checked 11-May-2019].
- Blumensath, T. and Davies, M. E. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models. In *International Conference on Machine Learning*, pp. 537–546, 2017.
- Boutsidis, C., Garber, D., Karnin, Z., and Liberty, E. Online principal components analysis. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pp. 887–901. Society for Industrial and Applied Mathematics, 2015.
- Boyd, S. Subgradient methods. *Notes for EE364b, Stanford University, Spring 2013–14*, 2014.
- Candès, E. J. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathématique*, 346(9-10):589–592, 2008.
- 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.
- Donoho, D. L. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- Donoho, D. L. and Elad, M. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- Donoho, D. L. and Tanner, J. Sparse nonnegative solution of underdetermined linear equations by linear programming. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27):9446–9451, 2005.
- Donoho, D. L., Maleki, A., and Montanari, A. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., and Bengio, S. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(Feb):625–660, 2010.
- Giryès, R., Eldar, Y. C., Bronstein, A. M., and Sapiro, G. Tradeoffs between convergence speed and reconstruction accuracy in inverse problems. *IEEE Transactions on Signal Processing*, 66(7):1676–1690, 2018.
- Gözcü, B., Mahabadi, R. K., Li, Y.-H., Ilıcak, E., Çukur, T., Scarlett, J., and Cevher, V. Learning-based compressive mri. *arXiv preprint arXiv:1805.01266*, 2018.

- Gregor, K. and LeCun, Y. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pp. 399–406, 2010.
- Hariharan, B., Vishwanathan, S., and Varma, M. Efficient max-margin multi-label classification with applications to zero-shot learning. *Machine learning*, 88(1-2):127–155, 2012.
- Haviv, I. and Regev, O. The restricted isometry property of subsampled fourier matrices. In *Geometric Aspects of Functional Analysis*, pp. 163–179. Springer, 2017.
- He, H., Xin, B., Ikehata, S., and Wipf, D. From bayesian sparsity to gated recurrent nets. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 5560–5570, 2017.
- Hegde, C., Indyk, P., and Schmidt, L. Approximation-tolerant model-based compressive sensing. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pp. 1544–1561. SIAM, 2014.
- Hegde, C., Sankaranarayanan, A. C., Yin, W., and Baraniuk, R. G. Numax: A convex approach for learning near-isometric linear embeddings. *IEEE Transactions on Signal Processing*, 63(22):6109–6121, 2015.
- Hershey, J. R., Roux, J. L., and Wenginger, F. Deep unfolding: Model-based inspiration of novel deep architectures. *arXiv preprint arXiv:1409.2574*, 2014.
- Hinton, G. E. and Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- Hotelling, H. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 448–456, 2015.
- Jain, H., Prabhu, Y., and Varma, M. Extreme multi-label loss functions for recommendation, tagging, ranking & other missing label applications. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 935–944. ACM, 2016.
- Jasinska, K., Dembczynski, K., Busa-Fekete, R., Pfannschmidt, K., Klerx, T., and Hullermeier, E. Extreme f-measure maximization using sparse probability estimates. In *International Conference on Machine Learning*, pp. 1435–1444, 2016.
- Jernite, Y., Choromanska, A., and Sontag, D. Simultaneous learning of trees and representations for extreme classification and density estimation. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1665–1674, 2017.
- Jin, K. H., McCann, M. T., Froustey, E., and Unser, M. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.
- Kellman, M., Bostan, E., Repina, N., and Waller, L. Physics-based learned design: Optimized coded-illumination for quantitative phase imaging. *IEEE Transactions on Computational Imaging*, 2019.
- Khajehnejad, M. A., Dimakis, A. G., Xu, W., and Hassibi, B. Sparse recovery of nonnegative signals with minimal expansion. *IEEE Transactions on Signal Processing*, 59(1):196–208, 2011.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *The International Conference on Learning Representations (ICLR)*, 2014.
- Kulkarni, K., Lohit, S., Turaga, P., Kerviche, R., and Ashok, A. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 449–458, 2016.
- Lewis, D. D., Yang, Y., Rose, T. G., and Li, F. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.
- Li, C. J., Wang, M., and Zhang, T. Diffusion approximations for online principal component estimation and global convergence. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Li, Y.-H. and Cevher, V. Learning data triage: linear decoding works for compressive mri. In *ICASSP*, 2016.
- Liu, J., Chang, W.-C., Wu, Y., and Yang, Y. Deep learning for extreme multi-label text classification. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 115–124. ACM, 2017.
- Lohit, S., Kulkarni, K., Kerviche, R., Turaga, P., and Ashok, A. Convolutional neural networks for non-iterative reconstruction of compressively sensed images. *IEEE Transactions on Computational Imaging*, 2018.
- Mairal, J., Bach, F., Ponce, J., and Sapiro, G. Online dictionary learning for sparse coding. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, pp. 689–696. ACM, 2009.

- Malloy, M. L. and Nowak, R. D. Near-optimal adaptive compressed sensing. *IEEE Transactions on Information Theory*, 60(7):4001–4012, 2014.
- Mardani, M., Monajemi, H., Papyan, V., Vasanawala, S., Donoho, D., and Pauly, J. Recurrent generative adversarial networks for proximal learning and automated compressive image recovery. *arXiv preprint arXiv:1711.10046*, 2017.
- Mardani, M., Sun, Q., Donoho, D., Papyan, V., Monajemi, H., Vasanawala, S., and Pauly, J. Neural proximal gradient descent for compressive imaging. In *Advances in Neural Information Processing Systems*, pp. 9573–9583, 2018.
- Meghwanshi, M., Jawanpuria, P., Kunchukuttan, A., Kasai, H., and Mishra, B. Mctorch, a manifold optimization library for deep learning. *arXiv preprint arXiv:1810.01811*, 2018.
- Mika, S., Schölkopf, B., Smola, A. J., Müller, K.-R., Scholz, M., and Rätsch, G. Kernel pca and de-noising in feature spaces. In *Advances in neural information processing systems*, pp. 536–542, 1999.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.
- Mineiro, P. and Karampatziakis, N. Fast label embeddings via randomized linear algebra. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 37–51. Springer, 2015.
- Mousavi, A. and Baraniuk, R. G. Learning to invert: Signal recovery via deep convolutional networks. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2017.
- Mousavi, A., Patel, A. B., and Baraniuk, R. G. A deep learning approach to structured signal recovery. In *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*, pp. 1336–1343. IEEE, 2015.
- Mousavi, A., Dasarathy, G., and Baraniuk, R. G. Deepcodec: Adaptive sensing and recovery via deep convolutional neural networks. In *55th Annual Allerton Conference on Communication, Control and Computing*, 2017.
- Mousavi, A., Dasarathy, G., and Baraniuk, R. G. A data-driven and distributed approach to sparse signal representation and recovery. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1xVTjCqKQ>.
- Needell, D. and Tropp, J. A. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- Niculescu-Mizil, A. and Abbasnejad, E. Label filters for large scale multilabel classification. In *Artificial Intelligence and Statistics*, pp. 1448–1457, 2017.
- Olshausen, B. A. and Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- Pan, V. Y. How bad are vandermonde matrices? *SIAM Journal on Matrix Analysis and Applications*, 37(2):676–694, 2016.
- Pennington, J., Socher, R., and Manning, C. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Prabhu, Y. and Varma, M. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 263–272. ACM, 2014.
- Prabhu, Y., Kag, A., Gopinath, S., Dahiya, K., Harsola, S., Agrawal, R., and Varma, M. Extreme multi-label learning with label features for warm-start tagging, ranking & recommendation. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 441–449. ACM, 2018a.
- Prabhu, Y., Kag, A., Harsola, S., Agrawal, R., and Varma, M. Parabel: Partitioned label trees for extreme classification with application to dynamic search advertising. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pp. 993–1002, 2018b.
- Rauhut, H. Compressive sensing and structured random matrices. *Theoretical Foundations and Numerical Methods for Sparse Recovery*, 9:1–92, 2010.
- Seeger, M. W. and Nickisch, H. Large scale bayesian inference and experimental design for sparse linear models. *SIAM Journal on Imaging Sciences*, 4(1):166–199, 2011.
- Shi, W., Jiang, F., Zhang, S., and Zhao, D. Deep networks for compressed image sensing. In *Multimedia and Expo (ICME), 2017 IEEE International Conference on*, pp. 877–882. IEEE, 2017.
- Sprechmann, P., Bronstein, A. M., and Sapiro, G. Learning efficient sparse and low rank models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1821–1833, 2015.

- Tagami, Y. Annexml: Approximate nearest neighbor search for extreme multi-label classification. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 455–464. ACM, 2017.
- Tropp, J. A. and Gilbert, A. C. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(Dec):3371–3408, 2010.
- Wang, Z., Ling, Q., and Huang, T. Learning deep l0 encoders. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Wu, S., Bhojanapalli, S., Sanghavi, S., and Dimakis, A. G. Single pass pca of matrix products. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2585–2593, 2016.
- Xin, B., Wang, Y., Gao, W., Wipf, D., and Wang, B. Maximal sparsity with deep networks? In *Advances in Neural Information Processing Systems (NIPS)*, pp. 4340–4348, 2016.
- Yen, I. E., Huang, X., Dai, W., Ravikumar, P., Dhillon, I., and Xing, E. Ppdspare: A parallel primal-dual sparse method for extreme classification. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 545–553. ACM, 2017.
- Yen, I. E.-H., Huang, X., Ravikumar, P., Zhong, K., and Dhillon, I. Pd-sparse: A primal and dual sparse approach to extreme multiclass and multilabel classification. In *International Conference on Machine Learning*, pp. 3069–3077, 2016.
- Yu, H.-F., Jain, P., Kar, P., and Dhillon, I. Large-scale multi-label learning with missing labels. In *International conference on machine learning*, pp. 593–601, 2014.
- Zhang, J. and Ghanem, B. Ista-net: Iterative shrinkage-thresholding algorithm inspired deep network for image compressive sensing. *arXiv preprint arXiv:1706.07929*, 2017.