

---

# Thwarting Adversarial Examples: An $L_0$ -Robust Sparse Fourier Transform

---

**Mitali Bafna** \*

School of Engineering & Applied Sciences  
Harvard University  
Cambridge, MA USA  
mitalibafna@g.harvard.edu

**Jack Murtagh** \*

School of Engineering & Applied Sciences  
Harvard University  
Cambridge, MA USA  
jmurtagh@g.harvard.edu

**Nikhil Vyas**\*

Department of Electrical Engineering and Computer Science  
MIT  
Cambridge, MA USA  
nikhilv@mit.edu

## Abstract

We give a new algorithm for approximating the Discrete Fourier transform of an approximately sparse signal that is robust to worst-case  $L_0$  corruptions, namely that a bounded number of coordinates of the signal can be corrupted arbitrarily. Our techniques generalize to a wide range of linear transformations that are used in data analysis such as the Discrete Cosine and Sine transforms, the Hadamard transform, and their high-dimensional analogs. We use our algorithm to successfully defend against worst-case  $L_0$  adversaries in the setting of image classification. We give experimental results on the Jacobian-based Saliency Map Attack (JSMA) and the Carlini Wagner (CW)  $L_0$  attack on the MNIST and Fashion-MNIST datasets as well as the Adversarial Patch on the ImageNet dataset.

## 1 Introduction

In the last several years, neural networks have made unprecedented achievements on computational learning tasks like image classification. Despite their remarkable success, neural networks have been shown to be brittle in the presence of adversarial noise [SZS<sup>+</sup>13]. Many effective attacks have been proposed in the context of computer vision that reliably generate small perturbations to input images (sometimes imperceptible to humans) that drastically change the network’s classification of the image [MFF15, GSS15, CW16]. As deep learning becomes more integrated into our everyday technology, the need for systems that are robust to adversarial noise grows, especially in applications to security.

A lot of work has been done to improve robustness and defend against adversarial attacks [PMW<sup>+</sup>16, TKP<sup>+</sup>17, MMS<sup>+</sup>17]. However many approaches rely on knowing the attack strategy in advance and too few proposed methods for robustness offer theoretical guarantees and may be broken by a new attack shortly after they’re published. As such, recent deep learning literature has seen an arms race of back-and-forth attacks and defenses reminiscent of cryptography before it was grounded in firm theoretical foundations.

In this work, we give a framework for improving the robustness of classifiers to adversaries with  $L_0$  noise budgets. That is, adversaries are restricted in the number of features they can corrupt, but may corrupt each arbitrarily. Our framework is based on a new Sparse Discrete Fourier transform (DFT) that is robust to worst-case  $L_0$  noise added in the time domain. We call such transformations

---

\* Authors ordered alphabetically.

$L_0$ -robust. In particular, we show how to recover top coefficients of an approximately sparse signal that has been corrupted by worst-case  $L_0$  noise.

Our theoretical results use techniques from compressed sensing [CRT06, BCDH10]. In fact we provide a much more general framework for building  $L_0$ -robust sparse transformations that applies to many transformations used in practice such as all discrete variants of the Fourier transform, the Sine and Cosine Transforms, the Hadamard transform, and their higher-dimensional generalizations. Our approach can be used to develop algorithms with the following benefits:

- **Provable performance guarantees.** Our approach leverages rigorous results in compressed sensing that allow us to prove theorems about  $L_0$ -robust sparse transformations under mild assumptions that typically hold in practice.
- **Worst-case adversaries.** The guarantees of our algorithms hold for all adversaries that stay within the noise budget, given that the input signals are sparse in the Fourier or related domains. In particular, our defenses do not require prior knowledge of the adversary’s attack strategy. In Section 4.2, we relax this and show that one can design improved algorithms within our framework when given more knowledge of the adversary.
- **Generality.** Our framework is general purpose and is compatible with a variety of basis transformations commonly used in scientific computing.

A notable feature of our framework is the focus on  $L_0$  noise. This threat model has been considered in previous works and  $L_0$  attacks and defenses have been developed [CW16, PMJ<sup>+</sup>15, PMW<sup>+</sup>16]. While  $L_2$  attacks are more commonly studied, many of the most high-profile recent real-world attacks actually fit in the  $L_0$  model, such as the graffiti-like road sign perturbations of [EEF<sup>+</sup>17], the eyeglasses that fool facial recognition software [SBBR16], and the patch that can make almost any image get labeled as a ‘toaster’ by state-of-the-art classifiers [BMR<sup>+</sup>17]. In general, physical obstructions in images or malicious splicing of audio or video files are realistic threats that can be modeled as  $L_0$  noise, whereas  $L_2$  attacks may be more difficult to carry out in the physical world.

The connection between our  $L_0$ -robust transformations and adversarial attacks on images is as follows. Many natural images are sparse in Fourier bases such as the basis used in the Discrete Cosine transform (DCT). Indeed, this is a necessary feature for compression algorithms like JPEG to work. Through this lens, corrupted images can be viewed as noisy signals that are sparse in some domain and our techniques allow us to reconstruct these sparse signals under worst-case/adversarial  $L_0$  noise. This reconstruction allows us to correct the corruptions made by the adversary to produce something close to the original image, which in turn improves the neural network accuracy.

Our results have wide applicability in signal processing since it is well known that audio/video signals are sparse in the Fourier or wavelet domains. Signal processing is important in many areas of science and medicine including MRI, radio astronomy, and facial recognition. Errors are ubiquitous in the above applications, whether due to natural artifacts, sensor failures, or malicious tampering.

In Section 4, we give experimental results that demonstrate the effectiveness of our approach against leading  $L_0$  attacks. For example, in one experiment the network accuracy drops from 88.5% on uncorrupted images to 24.8% on adversarial images with 30 pixels corrupted, but after our correction, network accuracy returns to 83.1%. On another attack, the adversary is free to choose its own budget and network accuracy drops from 87.8% all the way to 0% (the adversary succeeds on every image) but after running our correction algorithm, network accuracy returns to 85.7%.

In Section 2, we set up the problem and discuss related work. We give new theoretical results in Section 3. In Section 4, we evaluate our framework on three leading  $L_0$  attacks in the literature: the JSMA attack of Papernot et al [PMJ<sup>+</sup>15], the  $L_0$  attack from Carlini and Wagner (CW) [CW16], and the adversarial patch from Brown et al [BMR<sup>+</sup>17].

**Notation** For a vector  $v$ , we let  $v_{h(k)}$  and  $v_{t(k)}$  denote the head( $k$ ) and tail( $k$ ) of  $v$ . That is,  $v_{h(k)}$  denotes the vector containing just the  $k$  largest coordinates of  $v$  in absolute value with all other coordinates set to 0 and  $v_{t(k)} = v - v_{h(k)}$ . For example if  $v = [-3, 2, 1]$  then  $v_{h(2)} = [-3, 2, 0]$  and  $v_{t(2)} = [0, 0, 1]$ . We refer to  $v_{h(k)}$  as the “top  $k$  coefficients of  $v$ ”. For a vector  $v$ , we let  $\hat{v} = Fv$ , where  $F$  is the contextual linear transformation. We use the phrase, “projection of  $v$  to its top- $k$   $F$ -coefficients”, to mean the result of  $F^{-1}(Fv)_{h(k)}$ . That is, calculate the top- $k$  coefficients of  $Fv$ , set the remaining coefficients to 0 and then invert the result to back to the original domain by applying

$F^{-1}$ . Unless specified,  $\|\cdot\|$  denotes the  $L_2$  norm of a vector. For a scalar  $c$ ,  $|c|$  denotes the absolute value when  $c \in \mathbb{R}$  and denotes the modulus of  $c$  when  $c \in \mathbb{C}$ . We say that a vector  $v$  is  $k$ -sparse if all but  $k$  of its entries are 0. We say that  $v$  is approximately  $(k, \epsilon)$ -sparse if  $\|x_{t(k)}\| \leq \epsilon \cdot \|x\|$ . We define  $\mathcal{M}_k$  to be the set of all  $k$ -sparse vectors.  $\vec{0}$  denotes the all-zeroes vector.

## 2 Overview

### 2.1 $L_0$ -Robust Sparse Fourier Transform

**Problem Setup:** The key property that we use is that natural images are approximately sparse in frequency bases like the 2D Discrete Fourier basis or the 2D Discrete Cosine basis. This sparsity is exploited in image and video compression algorithms like JPEG and MPEG. The DFT and DCT are just linear transformations (in fact change of bases) from the space of images to a frequency domain. So given a  $d \times d$  image, we model it as approximately sparse in one of these bases, which from now on we will just refer to as the ‘Fourier basis’. Note that once the basis is fixed we can think of the image  $x \in \mathbb{R}^n$  ( $n = d^2$ ) as an approximately sparse vector in the corresponding Fourier basis.

Our goal is to approximate the top- $k$  Fourier coefficients of a vector  $x$  even after it has been corrupted with adversarial  $L_0$  noise. We do not know the locations or magnitudes of the corruptions but we do assume that we know an upper bound on the number of corrupted coordinates. In other words, if  $F$  is the Fourier matrix (the matrix corresponding to the Discrete Fourier linear transformation), we want to approximate  $\hat{x}_{h(k)}$  where  $\hat{x} = Fx$ . This can be modeled as the following problem:

**Problem 2.1 (Main Problem).** Given a corrupted vector  $y = x + e$  where  $x \in \mathbb{R}^n$  is approximately  $k$ -sparse in the Fourier basis and  $e$  is exactly  $t$ -sparse in the time domain (i.e. has  $L_0$  norm bounded by  $t$ ), approximate  $\hat{x}_{h(k)}$ .

We will solve the above problem by splitting  $y$  into  $x' + e' + \beta$  where  $x'$  is **exactly**  $k$ -sparse in the Fourier domain,  $e'$  is  $t$ -sparse in the time domain and  $\beta$  is an error term bounded in  $L_2$  norm by the tail of  $x$ . Our techniques are not limited to Fourier matrices and in fact extend naturally to other transformations like wavelets, but for simplicity we will use the term Fourier throughout.

**Related Work:** Our setting is reminiscent of extensively studied dimensionality reduction techniques like Robust PCA [CLMW11] for recovery of low rank matrices from  $L_0$  corrupted data. These have wide applicability in machine learning although, in that setting, they are not able to handle truly adversarial noise and make some assumptions on the error distribution. Our results on the other hand, can protect against worst-case adversaries bounded in their  $L_0$  noise budget.

Variants of the Sparse Fourier Transform have been studied [HIKP12a, HIKP12b, IKP14] but that work is concerned with recovering  $\hat{x}_{h(k)}$  given an approximately sparse vector  $x$ , using sublinear measurements. Our focus is on recovering  $\hat{x}_{h(k)}$  when some of the measurements might be corrupted and we show a tight tradeoff between the number of measurements corrupted versus the quality of recovery we can ensure.

We model images as being approximately sparse in the Fourier domain and prove that in such approximately sparse signals, it is possible to recover from  $L_0$  budgeted adversaries. In a similar manner, GANs have been used [SKC18, IJA<sup>+</sup>17] to model the distribution of unperturbed images, and to detect adversarial perturbations and recover from them, but these approaches are not able to provide theoretical guarantees of recovery.

**Our Techniques:** Our main result uses techniques from the field of compressed sensing (CS) [CRT06, BCDH10] and properties of Fourier (and related) matrices. Using these we prove that Algorithm 1 converges to a good solution to Problem 2.1, where by a good solution we mean that it is close to the true solution in the  $L_\infty$  norm.

In iteration  $i$  of Algorithm 1,  $\hat{x}^{[i]}$  is an estimate of  $\hat{x}_{h(k)}$  and  $e^{[i]}$  is an estimate of  $e$ . In iteration  $i + 1$ , the algorithm uses the previous estimates,  $e^{[i]}$  and  $\hat{x}^{[i]}$  to update its estimates by solving the linear equation  $y = F^{-1}\hat{x} + e$  and projecting onto the top  $k$  Fourier coefficients of  $y - e^{[i]}$ . Note that while this algorithm is intuitive it does not necessarily converge to the true solution for similar settings. For example, if instead of the  $L_0$  norm,  $e$  was bounded in the  $L_\infty$  norm, then information theoretically, there is no algorithm which can give a good solution and hence this algorithm would

---

**Algorithm 1** Iterative Hard Thresholding (IHT) [BCDH10].

---

**Input:** Positive integers  $k, t$ , and  $T$ .  $y = x + e$ , where  $x \in \mathbb{R}^n$  is approximately  $k$ -sparse in the Fourier basis and  $e \in \mathbb{R}^n$  is exactly  $t$ -sparse in the time domain. Fourier matrix  $F$ .

**Output:**  $\hat{x}_{h(k)}$ , approximation of the top  $k$  Fourier coefficients of  $x$ .

```
1: function IHT( $y = x + e, F, k, t, T$ )
2:    $\hat{x}^{[1]} \leftarrow \vec{0}$ 
3:    $e^{[1]} \leftarrow \vec{0}$ 
4:   for  $i = 1 \cdots T$  do
5:      $\hat{x}^{[i+1]} \leftarrow (F(y - e^{[i]}))_{h(k)}$ 
6:      $e^{[i+1]} \leftarrow (y - F^{-1}\hat{x}^{[i]})_{h(t)}$ 
7:   end for
8:   return  $\hat{x}^{[T+1]}$ 
9: end function
```

---

not be able to either. In our setting though, we can show that Algorithm 1 has an exponentially fast convergence towards a good solution to Problem 2.1 and moreover the guarantees we get are tight in the information theoretic sense. We state our result below for a general class of transformations which includes the DFT, DCT and their higher-dimensional versions.

**Theorem 2.2** (Main Theorem). *Let  $F \in \mathbb{C}^{n \times n}$  be an orthonormal matrix, such that each of its entries  $F_{ij}$ ,  $|F_{ij}|$  is  $O(1/\sqrt{n})$ . Let  $\hat{x} = Fx \in \mathbb{C}^n$  be  $(k, \epsilon)$ -sparse,  $e \in \mathbb{R}^n$  be  $t$ -sparse and  $y = F^{-1}\hat{x} + e$ . Let  $\hat{x}^{[T]} = \text{IHT}(y, F, k, t, T)$ , for  $T = O(\log(\|x\| + \|e\|))$ , then*

1.  $\|\hat{x}^{[T]} - \hat{x}_{h(k)}\|_\infty = O(\sqrt{t/n} \cdot \|\hat{x}_{t(k)}\|) = O(\sqrt{t/n} \cdot \|\epsilon\hat{x}\|)$
2.  $\|\hat{x}^{[T]} - \hat{x}_{h(k)}\| = O(\sqrt{kt/n} \cdot \|\hat{x}_{t(k)}\|) = O(\sqrt{kt/n} \cdot \|\epsilon\hat{x}\|)$  (In fact (1) implies (2).)

Note the strong  $L_\infty - L_2$  guarantee that Theorem 2.2 gives us, with a tight dependence between  $t$ , the  $L_0$  budget of the adversary and  $(\epsilon, k)$ , the sparsity parameters of the inputs. Also, our choice of recovering just the top- $k$  coordinates of  $\hat{x}$ , i.e.  $\hat{x}_{h(k)}$ , instead of all of  $\hat{x}$  is important. In the latter case, no matter what  $t$  is, any solution we recover would incur an  $L_2$  error of  $\Omega(\|\hat{x}_{t(k)}\|)$ , even when the adversary corrupts only one coordinate ( $t = 1$ ), while in our case, with  $t = o(n/k)$ , we get an  $L_2$  error that vanishes with  $n$  and is equal to  $o(\|\hat{x}_{t(k)}\|)$ . (by Theorem 2.2)

## 2.2 Defending against $L_0$ budgeted adversaries

We model images as approximately  $k$ -sparse vectors in the 2D-DCT domain. Using the results from Section 2.1, we can recover the top- $k$  coefficients in the face of a worst-case adversary with an  $L_0$  budget. To apply this to image classifiers, we want to build a neural network to recognize images projected to their top- $k$  2D-DCT coefficients. This motivates the following framework for building classifiers that are robust to  $L_0$  adversaries:

1. Train a neural network on images projected to their top- $k$  2D-DCT coefficients. We refer to such projected images as “compressed images”.<sup>2</sup>
2. On adversarial input images we run our  $L_0$ -robust DCT algorithm to recover the top- $k$  coefficients. Then transform the sparse image back to the original domain.
3. Run the recovered/corrected image through the network.

In Section 2.1, we saw that recovering the top- $k$  projection of an image gives better theoretical bounds than recovering the whole image. Hence it is important that the neural network is also trained to recognize compressed images. Training only on compressed images could possibly reduce the accuracy of neural networks, but as has been observed and used in practice (e.g. the JPEG and MPEG compression algorithms), images contain most of their information in relatively few coefficients. This is validated on our datasets, where we incur a  $< 1\%$  loss in accuracy on MNIST and  $< 2.5\%$  for Fashion-MNIST when training on compressed rather than original images. Note that one still

---

<sup>2</sup>Indeed the JPEG lossy-compression algorithm essentially does such a top- $k$  projection!

needs our correction algorithm for  $L_0$ -corrupted images, since a naive compression of an adversarial example (by taking its top- $k$  projection) will not get classified correctly by a neural network in general. For example, if 1 pixel of the image is corrupted to have an extremely high magnitude, this would propagate into the top- $k$  coefficients of the DCT of the image too and the resulting compressed image will be nowhere close to the original uncorrupted image. Our correction algorithm does not depend on the magnitude of the corruptions, only their number ( $t$ ). Hence both the training of the neural network on compressed images and the correction algorithm are essential to our framework.

### 2.3 Reverse Engineering Attacks

In step 2 of our framework, we use Theorem 2.2 to get strong guarantees on the distance  $\delta$ , between the original compressed image  $x$  and the recovered image. Ideally,  $\delta$  will be so small that no adversarial examples exist in the  $\delta$ -ball around  $x$ . This may not always be achieved in practice though and there might exist a small number of adversarial examples that are in the  $\delta$ -ball from the original image. This leaves open the possibility, that an attacker could reverse engineer our algorithm and design an adversarial example that, when corrected, yields a (potentially different) adversarial example inside the  $\delta$ -ball centered at  $x$  (although it is unclear how one would do this). Such an attack can be prevented by initializing the IHT algorithm with random vectors  $\hat{x}^{[1]}, e^{[1]}$  (instead of all-zeros vectors) so that the resulting recovered image is not deterministic. Since there are only a small number of adversarial examples in the  $\delta$ -ball, this randomization would ensure that a reverse engineering attack would fail to hit an adversarial example, with high probability. The guarantees of the IHT algorithm (Theorem 2.2) are independent of the starting vectors and continue to hold with randomized initialization. The IHT algorithm used for the experiments reported in this work is not randomized, because current attacks were not designed to reverse engineer our defense, and the deterministic IHT itself gives good results.

## 3 Proof of Main Result

In this section we prove Theorem 2.2, which says that Algorithm 1 converges to a good solution (one that is close to the true vector in the  $L_\infty$  norm) to Problem 2.1. Our proof uses techniques from compressed sensing. The main problem studied in compressed sensing is reconstructing a signal  $x$  from few linear measurements. For arbitrary signals, this task is impossible, however the main idea of compressed sensing is that signals that are approximately sparse can be recovered using fewer than  $n$  linear measurements. This is modeled as,

**Problem 3.1.** Given observations  $y = Mx$  where  $x \in \mathbb{C}^n$  is an approximately sparse signal, and  $M$  is an  $m \times n$  matrix with  $m < n$ , recover the vector  $x$ .

A main success in compressed sensing (CS) is that there are efficient algorithms [BD08, NT08] for Problem 3.1 when the matrix  $M$  satisfies a property called the RIP.

**Definition 3.2 (Restricted Isometry Property (RIP)).** An  $m \times n$  matrix  $M$  has the  $(k, \delta)$ -restricted isometry property ( $(k, \delta)$ -RIP) if for all  $k$ -sparse vectors  $v$  we have,

$$(1 - \delta) \cdot \|v\| \leq \|Mv\| \leq (1 + \delta) \cdot \|v\|.$$

Recall that in our main problem (Problem 2.1), we want to recover the top- $k$  coefficients of  $\hat{x} = Fx$ , where  $\hat{x}$  is approximately  $k$ -sparse, given a corrupted vector  $y = x + e$ . The key idea is to notice that we can write  $y$  as  $[F^{-1} \ I] \begin{bmatrix} \hat{x} \\ e \end{bmatrix}$ , where  $\hat{x}$  is approximately  $k$ -sparse and  $e$  is  $t$ -sparse. This is almost the same setup as Problem 3.1. In fact, we have more knowledge about the structure of sparsity of the vector  $\begin{bmatrix} \hat{x} \\ e \end{bmatrix} \in \mathbb{C}^{2n}$  that we want to recover.

The problem of recovery with structured sparsity, has been studied under the heading of Model-Based CS [BCDH10, HIS15, HIS14, BIS17] for structured sparsity models. In our setting we want to model vectors of the form  $\begin{bmatrix} \hat{x} \\ e \end{bmatrix}$ , which have sparsity  $k$  in  $x$  and  $t$  in  $e$ . This motivates the following definition.

**Definition 3.3.** Let  $\mathcal{M}_{k,t} \subseteq \mathbb{C}^{2n}$  be the set of all vectors where the first  $n$  coordinates are  $k$ -sparse and the last  $n$  coordinates are  $t$ -sparse.<sup>3</sup> Formally,

$$\mathcal{M}_{k,t} := \left\{ v = \begin{bmatrix} x \\ e \end{bmatrix} \in \mathbb{C}^{2n} \mid x \text{ is } k\text{-sparse}, e \text{ is } t\text{-sparse} \right\}.$$

We say that a matrix  $M$  has the  $((k, t), \delta)$ -RIP if for all vectors  $v \in \mathcal{M}_{k,t}$ ,

$$(1 - \delta) \cdot \|v\| \leq \|Mv\| \leq (1 + \delta) \cdot \|v\|.$$

Model-Based CS was first introduced in [BCDH10], for general sparsity models, and they proved therein that Iterative Hard Thresholding (IHT) [BD08] indeed converges to a good solution to Problem 3.1, given that the measurement matrix  $M$  satisfies RIP for the model. We use this Model-Based IHT approach to argue that Algorithm 1 finds a good solution to Problem 2.1. In [BCDH10], they proved that the IHT algorithm converges to an approximately correct solution, given that the measurement matrix  $M$  satisfies the RIP for the model at hand. For us this translates to the following theorem.

**Theorem 3.4** ([BCDH10]). *Let  $v \in \mathcal{M}_{k,t}$  and let  $y = Mv + \beta$ , where  $M \in \mathbb{R}^n$  is a full-rank matrix and  $\beta$  is a noise vector. Let  $v^{[T]} = \text{IHT}(y, M^{-1}, k, t, T)$ . If  $M$  is  $((3k, 3t), \delta)$ -RIP, with  $\delta \leq 0.1$ , then*

$$\|v^{[T]} - v\| \leq 2^{-T} \cdot \|v\| + 4 \cdot \|\beta\|.$$

We use the above theorem to prove that Algorithm 1 also converges to a good solution. Another key technique we use in our proofs is an uncertainty principle for specific structured matrices.

**Lemma 3.5** (General Uncertainty Principle). *Let  $F$  be a matrix in  $\mathbb{C}^{n \times n}$  such that each entry  $F_{ij}$  has  $|F_{ij}| \leq \alpha$ . Let  $x$  be a  $k$ -sparse vector in  $\mathbb{C}^n$  and  $y = Fx$ . Then  $\|y\|_\infty \leq \alpha \cdot \sqrt{k} \cdot \|x\|$ .*

Note that when  $F$  is the normalized Fourier matrix, this is the same as the folklore Fourier uncertainty principle with  $\alpha = 1/\sqrt{n}$ . The proof of the above is indeed very similar to the Fourier case and is included in the appendix. One can check that for transformation matrices corresponding to Discrete Cosine and Sine Transforms and their 2D analogs we have  $\alpha = O(\sqrt{1/n})$ .

Finally to prove Theorem 2.2, we first prove that the matrix  $M = [F^{-1} \ I]$  has the RIP (Lemma 3.6 below), which then combined with Theorem 3.4 and the uncertainty principle 3.5, finishes the proof of Theorem 2.2. These proofs can be found in the appendix.

**Lemma 3.6.** *Let  $F$  be an orthonormal matrix, such that each entry  $F_{ij}$  has  $|F_{ij}| = O(1/\sqrt{n})$ . Then the matrix  $M = [F^{-1} \ I] \in \mathbb{C}^{n \times 2n}$  satisfies  $((3k, 3t), \delta)$ -RIP with  $\delta \leq 0.1$ , when  $t = O(n/k)$ .*

*Equivalently, for all vectors  $v = \begin{bmatrix} \hat{x} \\ e \end{bmatrix}$ , such that  $\hat{x}$  is  $3k$ -sparse and  $e$  is at most  $3t = O(n/k)$ -sparse,*

$$(1 - \delta) \cdot \|v\| \leq \|Mv\| \leq (1 + \delta) \cdot \|v\|.$$

## 4 Experiments

### 4.1 Worst case adversaries

We evaluated our framework on three leading  $L_0$  attacks in the literature: the JSMA Attack of Papernot et al [PMJ<sup>+</sup>15], the  $L_0$  attack from Carlini and Wagner (CW) [CW16], and the adversarial patch from Brown et al [BMR<sup>+</sup>17]. We evaluated Algorithm 1 on the JSMA and CW attacks and present these results in this section. We discuss experiments on the adversarial patch in Section 4.2.

We tested both JSMA and CW on two datasets: the MNIST handwritten digits [LeC98] and the Fashion-MNIST [XRV17] dataset of clothing images. For each attack, we used randomly selected targets. For both datasets we used a neural network composed of a convolutional layer (32 kernels of 3x3), max pooling layer (2x2), convolutional layer (64 kernels of 3x3), max pooling layer (2x2), fully connected layer (128 neurons) with dropout (rate = .25) and an output softmax layer (10 neurons). We used the Adam optimizer with cross-entropy loss and ran it for 10 epochs over the training datasets.

<sup>3</sup>Recall that  $\mathcal{M}_k$  was the set of all  $k$ -sparse vectors in  $\mathbb{C}$ . Note that  $\mathcal{M}_{k,t}$  is different from  $\mathcal{M}_{k+t}$  which is the set of all vectors  $\in \mathbb{C}^{2n}$  that are  $k+t$ -sparse.

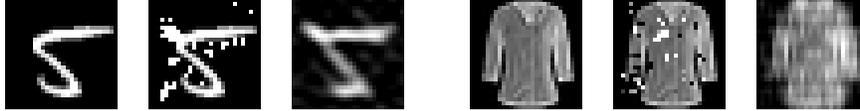


Figure 1: Example experimental conditions. The left 3 images depict an original MNIST image, the image corrupted by JSMA, and our corrected image. The right three images show an original Fashion-MNIST image, the image corrupted by CW, and our corrected image.

For each dataset, we trained our neural network only on images that were projected onto their top- $k$  2D-DCT coefficients. Here  $k$  is a parameter we tuned depending on the dataset (for MNIST  $k = 40$  and for Fashion-MNIST  $k = 35$ ). For each dataset, we fixed its corresponding  $k$  across all experiments reported here.

In all of our evaluations there were three experimental conditions: first we ran uncorrupted images through the network to establish a baseline accuracy. Then we ran the  $L_0$  adversarial examples through the network. Finally, we ran our correction algorithm on the adversarial examples and ran the results through the network. Example images of these conditions can be seen in Figure 1.

For the JSMA, we ran an experiment for several different adversary noise budgets. For each budget, we evaluated the network on the three experimental conditions. The accuracy vs  $L_0$  budget and loss vs  $L_0$  budget graphs can be seen in Figure 2 on the MNIST and Fashion-MNIST datasets. Exact values can be found in Section D of the appendix. The results demonstrate that our correction algorithm successfully defends against the JSMA attack. For example when the adversary corrupts 30 bits, it is able to drop the accuracy of our network on the Fashion-MNIST dataset from 88.5% to 24.8% but after running our recovery algorithm we get back up to 83.1%.

The CW attack works by finding a minimal set of pixels that can be corrupted to fool the network. This means that the adversary’s budget will depend on the particular image being corrupted rather than being fixed in advance. For this reason, we let the CW adversary choose how many pixels to corrupt and allow ourselves to know its budget for each image. Note that the locations and magnitudes of the noise are unknown to us. Since the budget varies across images, a plot like Figure 2 does not make sense and we instead report the overall accuracy and loss of our correction algorithm in Table 3. Again our correction algorithm is effective against CW. For example on the Fashion-MNIST dataset the network’s test accuracy on original images was 87.8%. The CW attack was successful and the network mislabeled every adversarial example. After running our correction algorithm, the accuracy returns to 85.7%.

	MNIST	Adversarial	Corrected	F-MNIST	Adversarial	Corrected
Accuracy	99.0	0.0	72.8	87.8	0.0	85.7
Loss	0.002	0.115	0.095	0.035	0.140	0.040

Figure 3: Experimental results for our algorithm on the CW attack for the MNIST and Fashion-MNIST datasets. Columns 2-4 show results for MNIST data and columns 5 to 7 show Fashion-MNIST.

As images grow larger they become less sparse in Fourier bases but natural images are still block-wise sparse. In such cases our algorithm could be modified to correct images block by block, in which case the network would need to be trained on images compressed block by block (e.g. as in JPEG). This would work with the mild assumption that the corrupted locations are well-distributed across blocks because then our recovery result could be applied to each block separately. Within each block the corrupted locations could still be anywhere and of any magnitude. For images that are too large to be sparse in Fourier bases, the block-wise approach may fail in the case where most of the  $L_0$  noise resides in few blocks because in these blocks there will be too many corrupted coordinates to recover. In the next section we study the extreme case where all of the error is concentrated contiguously. We show that even in this extreme case our framework for  $L_0$ -robust sparse transformations can be used to guard against contiguous noise attacks even in large images.

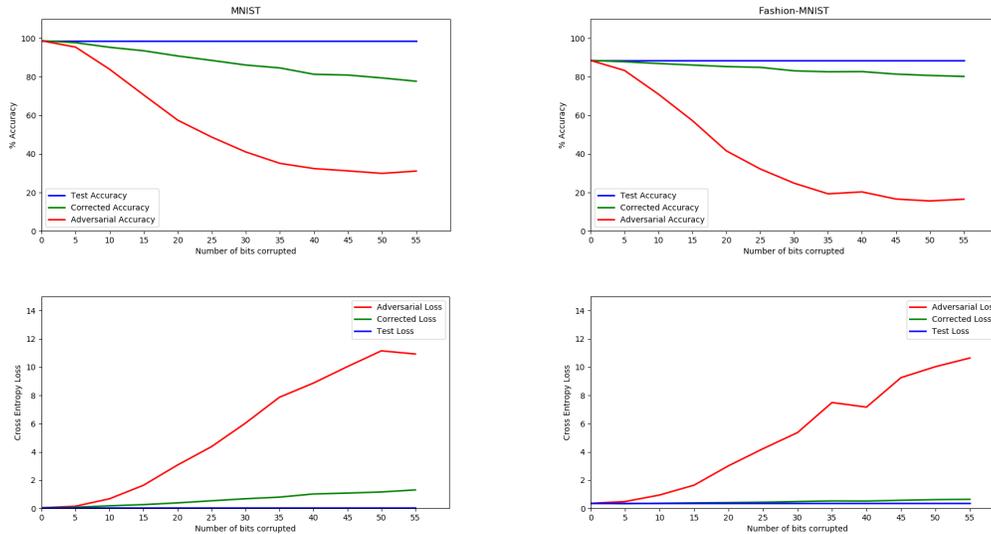


Figure 2: Classification accuracy and loss for JSMA on MNIST (left) and Fashion-MNIST (right). Blue lines show the performance of the network on original images (and hence does not change with the number of coordinates corrupted). Red lines show the performance of the network on uncorrected adversarial examples and green lines show the performance of the network on images that were corrected by Algorithm 1

## 4.2 Adversarial patch

In [BMR<sup>+</sup>17], the authors introduce a method for generating adversarial patches. These are targeted attacks in the form of circular images that get overlaid on input images. They showed that their patch effectively fools leading image classifiers into mislabeling patched images.

Notice that the adversarial patch is an example of  $L_0$  noise and so fits within our framework. The patch attack is only successful when the patch is sufficiently large ( $\sim 80$  pixels in diameter for  $224 \times 224$  images), which is larger than our algorithm can tolerate. Also images of this size are less sparse in the frequency domain and as discussed above, our approach may not be able to correct contiguous noise on such images. Similarly we cannot train the neural networks on compressed images as that would lead to non trivial loss as the images are less sparse. So in this section we use a network that was pretrained on original ImageNet images.

We are able to use the contiguity of the noise with the mild sparsity of large images by using Algorithm 2 to defend against the patch attack. Since image recovery is not possible in this setting, our algorithm instead focuses on *detecting* the location of the contiguous noise. We detect the noise by searching over contiguous blocks in the image and running Algorithm 1 on each block, where we project  $e$  only to the block rather than top- $t$  coordinates. Finally we find the block for which the remaining image ( $y - e$ ) is sparsest in the Fourier domain. We call this Patchwise IHT and a formal description of the algorithm is given in the Appendix. Note that for this particular set of adversarial examples there may be other ways to detect the patch with pre-processing. We do not do any such optimizations that are particular to the adversary and Patchwise IHT is based only on the mild sparsity of the original images.

We took 700 random images from ImageNet and for classification we used pretrained ResNet-50 network [HZRS15]. We ran each image through the network in our three experimental conditions, depicted in Figure 4.

Figure 5 shows the results of our experiment. The patch was a successful attack (Top-5 accuracy dropped from 92.3% to 63.9% and Top-1 from 76.4% to 12.0%). After correcting, Top-5 accuracy jumped to 80.4% (Top-1: 59.7%). Only 1.0% of the original images were labeled as ‘toaster’ (none in the Top-1), but ‘toaster’ was in the Top-5 in 99.0% of the patched images with 85.7% being the



Figure 4: Example of the three image conditions in the patch experiment. Left is the original image classified as ‘banana’ with probability .94. The middle, with the adversarial patch overlayed, is classified as ‘toaster’ with probability .93. The right is the image after our Patchwise IHT algorithm, which gets classified as ‘banana’ with probability .94.

most confident label. Notably, very few corrected images were labeled as ‘toaster’ (Top-5: 7.4%, Top-1: 4.7%).

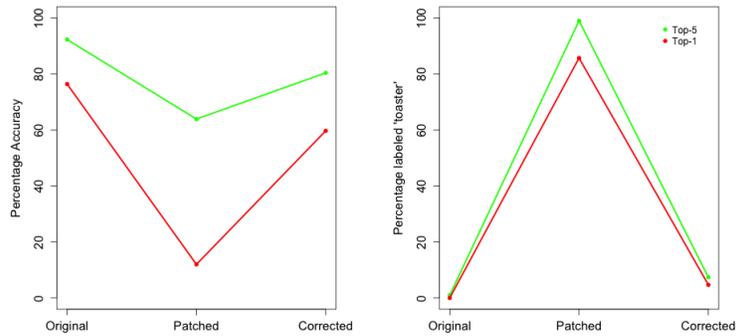


Figure 5: Experimental results for Algorithm 2. The left plot depicts the accuracy of the network in our three experimental conditions. The right plot shows the percentage of images labeled as ‘toaster’ under the same three conditions.

## 5 Acknowledgements

We would like to thank Yaron Singer and Adam Breuer for helpful feedback and encouragement in the early stages of this work. We also want to thank Thibaut Horel for valuable comments on the manuscript. Thanks also to the reviewers for helpful remarks.

## References

- [BCDH10] Richard G. Baraniuk, Volkan Cevher, Marco F. Duarte, and Chinmay Hegde. Model-based compressive sensing. *IEEE Trans. Information Theory*, 56(4):1982–2001, 2010.
- [BD08] Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *CoRR*, abs/0805.0510, 2008.
- [BIS17] Arturs Backurs, Piotr Indyk, and Ludwig Schmidt. Better approximations for tree sparsity in nearly-linear time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 2215–2229, 2017.
- [BMR<sup>+</sup>17] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *CoRR*, abs/1712.09665, 2017.
- [CLMW11] Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, 2011.
- [CRT06] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Information Theory*, 52(2):489–509, 2006.
- [CW16] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. *CoRR*, abs/1608.04644, 2016.
- [EEF<sup>+</sup>17] Ivan Evtimov, Kevin Eykholt, Earlene Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on machine learning models. *CoRR*, abs/1707.08945, 2017.
- [GSS15] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- [HIKP12a] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Nearly optimal sparse fourier transform. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 563–578, 2012.
- [HIKP12b] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Simple and practical algorithm for sparse fourier transform. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1183–1194, 2012.
- [HIS14] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. Nearly linear-time model-based compressive sensing. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 588–599, 2014.
- [HIS15] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. Approximation algorithms for model-based compressive sensing. *IEEE Trans. Information Theory*, 61(9):5129–5147, 2015.
- [HZRS15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- [JA<sup>+</sup>17] Andrew Ilyas, Ajil Jalal, Eirini Asteri, Constantinos Daskalakis, and Alexandros G. Dimakis. The robust manifold defense: Adversarial training using generative models. *CoRR*, abs/1712.09196, 2017.
- [IKP14] Piotr Indyk, Michael Kapralov, and Eric Price. (nearly) sample-optimal sparse fourier transform. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 480–499, 2014.
- [LeC98] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.

- [MFF15] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. *CoRR*, abs/1511.04599, 2015.
- [MMS<sup>+</sup>17] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *CoRR*, abs/1706.06083, 2017.
- [NT08] Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *arXiv preprint arXiv:0803.2392*, 2008.
- [PMJ<sup>+</sup>15] Nicolas Papernot, Patrick D. McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. *CoRR*, abs/1511.07528, 2015.
- [PMW<sup>+</sup>16] Nicolas Papernot, Patrick D. McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 582–597, 2016.
- [SBBR16] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1528–1540. ACM, 2016.
- [SKC18] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-gan: Protecting classifiers against adversarial attacks using generative models. *CoRR*, abs/1805.06605, 2018.
- [SZS<sup>+</sup>13] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013.
- [TKP<sup>+</sup>17] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- [XRV17] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017.

## A Proof of Lemma 3.5: the General Uncertainty Principle

*Proof.* For all  $i \in [n]$  we have

$$\begin{aligned} |(Fx)_i| &= \left| \sum_{j \in [n]} F_{ij} x_j \right| \\ &\leq \sum_{j \in [n]} |F_{ij}| \cdot |x_j| \\ &\leq \alpha \cdot \sum_{j \in [n]} |x_j| \\ &\leq \alpha \cdot \sqrt{k} \cdot \|x\| \end{aligned}$$

where the third line follows from our assumption on the entries in  $F$  and the fourth line follows from the Cauchy-Schwarz inequality. Since  $y = Fx$  we have  $\|y\|_\infty = \|Fx\|_\infty = \max_{i \in [n]} |(Fx)_i| \leq \alpha \cdot \sqrt{k} \cdot \|x\|$  by the above.  $\square$

## B Proof of Lemma 3.6

*Proof.* We will prove that  $[F^{-1} \ I]$  has the RIP for all  $v = \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \in \mathcal{M}_{3k, 3t}$ . That is,

$$0.9 \cdot \left\| \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\| \leq \left\| [F^{-1} \ I] \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\| \leq 1.1 \cdot \left\| \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\|.$$

Note that the right-hand inequality follows immediately: Since  $F^{-1}$  is orthonormal we have that  $\left\| [F^{-1} \ I] \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\| = \|F^{-1}\hat{x} + e\| \leq \|F^{-1}\hat{x}\| + \|e\| = \|\hat{x}\| + \|e\| = \left\| \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\| \leq 1.1 \cdot \left\| \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \right\|.$

Since  $x$  is  $k$ -sparse and  $F$  satisfies the hypotheses of the Uncertainty Principle (Lemma 3.5), with  $\alpha = O(1/\sqrt{n})$ , we have that  $\|F^{-1}\hat{x}\|_\infty \leq \alpha \cdot \sqrt{k} \cdot \|\hat{x}\|$ . Using this, we will prove that,  $\|F^{-1}\hat{x} + e\| \geq 0.9\sqrt{\|\hat{x}\|^2 + \|e\|^2}$ , when the sparsity of  $e$  is at most  $O(n/k)$ .

Since  $e$  is  $t$ -sparse, without loss of generality assume that  $e = [e_1, \dots, e_t, 0, \dots, 0]$  and  $F^{-1}\hat{x} = [x_1, \dots, x_n]$  with  $|x_i| \leq \alpha \cdot \sqrt{k} \cdot \|x\|$ . We have that,

$$\begin{aligned} \|F^{-1}\hat{x} + e\|^2 &\geq \sum_{i=1}^t (|x_i| - |e_i|)^2 + \sum_{i=t+1}^n x_{i+1}^2 \\ &= \|x\|^2 + \|e\|^2 - 2 \sum_{i=1}^t |x_i| \cdot |e_i| \\ &\geq \|x\|^2 + \|e\|^2 - 2 \cdot \alpha \cdot \sqrt{k} \|x\| \sum_{i=1}^t |e_i| && \text{(Uncertainty Principle)} \\ &\geq \|x\|^2 + \|e\|^2 - 2 \cdot \alpha \cdot \sqrt{k} \|x\| \cdot \sqrt{t} \|e\| && \text{(Cauchy-Schwartz inequality)} \\ &= \|x\|^2 + \|e\|^2 - O(\sqrt{kt/n} \|x\| \|e\|) && (1) \end{aligned}$$

We want that  $\|F^{-1}\hat{x} + e\| \geq 0.9\sqrt{\|\hat{x}\|^2 + \|e\|^2}$ . Plugging in equation 1 and moving terms around we get that this happens when  $t = O(n/k)$ . This completes the proof of the lemma.  $\square$

## C Proof of Theorem 2.2

*Proof of the main theorem:* Now we will prove that the RIP property of  $M = [F^{-1} \ I]$  proved above, combined with the Uncertainty Principle (Lemma 3.5) and Theorem 3.4 imply the main theorem.

Consider  $y = F^{-1}\hat{x} + e = F^{-1}\hat{x}_{h(k)} + F^{-1}\hat{x}_{t(k)} + e$ , where  $\|\hat{x}_{t(k)}\| \leq \epsilon \|x\|$  and  $e$  is  $t$ -sparse. Using  $\beta = F^{-1}\hat{x}_{t(k)}$ , we can rewrite this expression as,  $y = [F^{-1} \ I] \begin{bmatrix} \hat{x}_{h(k)} \\ e \end{bmatrix} + \beta = Mv + \beta$ , where  $\|\beta\| = \|F^{-1}\hat{x}_{t(k)}\| = \|\hat{x}_{t(k)}\| \leq \epsilon \|x\|$ , since  $F$  is orthonormal. In Lemma 3.6 we proved that the matrix  $M$  is  $((3k, 3t), \delta)$ -RIP (with  $\delta = 0.1$ ) for the set  $\mathcal{M}_{k,t}$ .

At the  $i^{th}$  iteration of the IHT algorithm let  $\begin{bmatrix} \hat{x}^{[i]} \\ e^{[i]} \end{bmatrix}$  be our estimate of  $\begin{bmatrix} \hat{x}_{h(k)} \\ e \end{bmatrix}$ . At the  $T^{th}$  iteration, by Theorem 3.4 we have the guarantee that,

$$\left\| \begin{bmatrix} \hat{x}^{[T]} \\ e^{[T]} \end{bmatrix} - \begin{bmatrix} \hat{x}_{h(k)} \\ e \end{bmatrix} \right\| \leq 2^{-T} \cdot \sqrt{\|x\|^2 + \|e\|^2} + 4\|\beta\| \approx 4\epsilon \|\hat{x}\| \quad (2)$$

$$\implies \|\hat{x}^{[T]} - \hat{x}_k\|^2 + \|e^{[T]} - e\|^2 \leq 16\epsilon^2 \|\hat{x}\|^2, \quad (3)$$

since we set  $T$  such that  $2^{-T} \cdot \sqrt{\|x\|^2 + \|e\|^2} \approx 0$ . Note that 3 already gives a weak  $L_2$ - $L_2$  guarantee on  $\|\hat{x}^{[T]} - \hat{x}_{h(k)}\|$  but we will derive a stronger  $L_\infty$ - $L_2$  guarantee.

Consider the  $T^{th}$  iteration of IHT, and define a vector  $z := F(y - e^{[i-1]})$ . The IHT algorithm sets  $\hat{x}^{[T]} = z_{h(k)}$ . We have that,

$$y = F^{-1}\hat{x} + e = F^{-1}z + e^{[i-1]} \Leftrightarrow \hat{x} - z = F(e^{[T-1]} - e)$$

Since both  $e, e^{[T-1]}$  are  $t$ -sparse, we have that the vector  $e^{[T-1]} - e$  is  $2t$ -sparse. By the uncertainty principle 3.5, we get that,  $\|\hat{x} - z\|_\infty \leq \sqrt{2t/n} \|e^{[T]} - e\|_2 = O(\epsilon \|\hat{x}\| \sqrt{t/n})$  by Equation 3. This trivially implies that  $\|\hat{x}_{h(k)} - z_{h(k)}\|_\infty = O(\epsilon \|\hat{x}\| \sqrt{t/n})$  which implies  $\|\hat{x}_{h(k)} - z_{h(k)}\| = O(\epsilon \|\hat{x}\| \sqrt{kt/n})$ .  $\square$

## D JSMA experimental results

MNIST Accuracy	0	5	10	15	20	25
Original Images	98.7	98.7	98.7	98.7	98.7	98.7
Adversarial Images	98.7	95.4	83.9	70.6	57.5	48.7
Corrected Images	98.7	97.7	95.3	93.5	90.8	88.5

MNIST Accuracy	30	35	40	45	50	55
Original Images	98.7	98.7	98.7	98.7	98.7	98.7
Adversarial Images	41.0	35.1	32.4	31.2	29.9	31.1
Corrected Images	86.1	84.6	81.3	80.9	79.4	77.7

MNIST Loss	0	5	10	15	20	25
Original Images	0.05	0.05	0.05	0.05	0.05	0.05
Adversarial Images	0.05	0.16	0.68	1.64	3.07	4.37
Corrected Images	0.05	0.08	0.19	0.27	0.39	0.54

MNIST Loss	30	35	40	45	50	55
Original Images	0.05	0.05	0.05	0.05	0.05	0.05
Adversarial Images	6.03	7.86	8.86	10.03	11.15	10.92
Corrected Images	0.69	0.80	1.02	1.09	1.17	1.31

Figure 6: Experimental results for the JSMA attack on MNIST. The columns represent the adversary's budget: the number of pixels corrupted from 0 to 55 in increments of 5, given in the first row. The top two tables show accuracy results on the MNIST for the three experimental conditions. The bottom two tables show the cross-entropy loss across the conditions and budgets.

F-MNIST Accuracy	0	5	10	15	20	25
Original Images	88.5	88.5	88.5	88.5	88.5	88.5
Adversarial Images	88.5	83.3	71.0	57.3	41.6	32.2
Corrected Images	88.5	87.8	86.9	86.1	85.3	84.9

F-MNIST Accuracy	30	35	40	45	50	55
Original Images	88.5	88.5	88.5	88.5	88.5	88.5
Adversarial Images	24.8	19.3	20.3	16.6	15.6	16.5
Corrected Images	83.1	82.6	82.7	81.4	80.7	80.2

F-MNIST Loss	0	5	10	15	20	25
Original Images	0.36	0.36	0.36	0.36	0.36	0.36
Adversarial Images	0.36	0.49	0.95	1.65	3.02	4.23
Corrected Images	0.36	0.34	0.37	0.39	0.41	0.44

F-MNIST Loss	30	35	40	45	50	55
Original Images	0.36	0.36	0.36	0.36	0.36	0.36
Adversarial Images	5.37	7.49	7.17	9.25	10.03	10.64
Corrected Images	0.49	0.53	0.52	0.58	0.62	0.65

Figure 7: Experimental results for the JSMA attack on Fashion-MNIST. The columns represent the adversary’s budget: the number of pixels corrupted from 0 to 55 in increments of 5, given in the first row. The top two tables show accuracy results on the Fashion-MNIST for the three experimental conditions. The bottom two tables show the cross-entropy loss across the the conditions and budgets.

## E Patchwise Iterative Hard Thresholding

---

### Algorithm 2 Patchwise Iterative Hard Thresholding (IHT)

---

**Input:** Positive integers  $k, t, T$ , and  $\ell$ .  $y = x + e$ , where  $x \in \mathbb{R}^n$  is approximately  $k$ -sparse in the Fourier basis and  $e \in \mathbb{R}^n$  is exactly  $t$ -sparse in the time domain. Fourier matrix  $F$ .

**Output:**  $x'$ , approximation of the original signal  $x$ .

```

1: function IHT( $y = F^{-1}\hat{x} + e, F, k, t, T, \ell$ )
2:    $x' \leftarrow \infty$ 
3:   for  $\ell \times \ell$  patch  $p$  in image do
4:      $\hat{x}^{[1]} \leftarrow \vec{0}$ 
5:      $e^{[1]} \leftarrow \vec{0}$ 
6:     for  $i = 1 \dots T$  do
7:        $\hat{x}^{[i+1]} \leftarrow (F(y - e^{[i]}))_{h(k)}$ 
8:        $e^{[i+1]} \leftarrow (y - F^{-1}\hat{x}^{[i]})_{h(t)}$ 
9:     end for
10:    if  $\|\hat{x}^{[T+1]}\| < \|x'\|$  then
11:       $x' = F^{-1}\hat{x}^{[T+1]}$ 
12:    end if
13:  end for
14:  return  $x'$ 
15: end function

```

---