

RICE UNIVERSITY

Compressive Phase Retrieval

by

Matthew L. Moravec

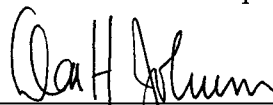
A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:



Richard G. Baraniuk, Chair
Victor E. Cameron Professor
Electrical and Computer Engineering



Don H. Johnson
J.S. Abercrombie Professor
Electrical and Computer Engineering
Statistics



Daniel Mittleman
Associate Professor
Electrical and Computer Engineering

Houston, Texas

April 2008

UMI Number: 1455265

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 1455265

Copyright 2008 by ProQuest LLC.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest LLC
789 E. Eisenhower Parkway
PO Box 1346
Ann Arbor, MI 48106-1346

ABSTRACT

Compressive Phase Retrieval

by

Matthew L. Moravec

In the phase retrieval problem, a signal must be recovered from the magnitude of its Fourier transform. The signal's compressibility can be used as prior knowledge to aid in its recovery. This knowledge also allows for recovery from fewer Fourier modulus measurements than the signal's bandwidth alone would dictate. We introduce a compressibility constraint for phase retrieval, define the number of measurements necessary for phase retrieval of sparse signals, and apply these concepts to terahertz imaging.

Acknowledgments

I would like to thank my advisor, Richard Baraniuk, for all of his helpful ideas, insight, advice, and encouragement. It was also a pleasure to work with Justin Romberg, Daniel Mittleman, and William Chan. I am grateful to Don Johnson for his perspective and feedback. A special thanks is also due to my family, especially my parents and two sisters, to all of my friends, and to my wonderful fiancée Margaret, for their love and support. I thank God for all of you, and I thank Him for the life He has given me through Jesus Christ.

Contents

Abstract	ii
Acknowledgments	iii
List of Illustrations	vi
1 Introduction and Background	1
1.1 Compressive Sensing	1
1.2 Phase Retrieval	4
1.2.1 The Phase Retrieval Problem	5
1.2.2 Phase Retrieval Methods	7
1.2.3 Signal Constraints	8
2 Compressive Phase Retrieval	10
2.1 ℓ_p Compressibility Constraint	10
2.2 Geometric Analysis of Compressibility Constraint	12
2.2.1 Signals Sparse in Space	12
2.2.2 Signals Sparse in Ψ	15
2.2.3 Non-Sparse Compressible Signals	16
2.3 Benefits of the Compressibility Constraint	16
3 Compressive Sensing Phase Retrieval	20
3.1 Sufficient Measurements for a Band-Limited Signal	21
3.2 Sufficient Measurements for a Sparse Signal	22
3.3 CSPR Recovery	25
3.4 CSPR Accuracy	27

4 Compressive Phase Retrieval in Practice	31
4.1 THz Imaging	31
4.2 CPR Recovery from the Entire Diffraction Pattern	32
4.3 CSPR Recovery from a Sub-sampled Diffraction Pattern	36
5 Conclusion	40
Bibliography	41

Illustrations

1.1	Fourier phase-only and magnitude-only reconstructions	5
2.1	Example of phase retrieval with a compressibility constraint	13
2.2	Local and global minima analysis of phase retrieval constraints, varying sparsity	17
2.3	Local and global minima analysis of phase retrieval constraints, varying signal size	18
2.4	Local and global minima analysis of phase retrieval constraints, varying compressibility	18
3.1	Local and global minima analysis of CPR constraints	28
3.2	Necessary number of measurements for CSPR	29
4.1	THz Fourier imaging setup	33
4.2	CPR with THz imaging	34
4.3	Decaying sorted signal magnitude of a THz image	35
4.4	CPR with varying ℓ_1 norm values	37
4.5	CSPR with varying numbers of measurements	38
4.6	CSPR error	39

Chapter 1

Introduction and Background

In science and engineering, signals of interest are often structured in some way. Knowledge of a signal's structure is used for a variety of signal acquisition and processing applications. For example, the bandwidth of a signal is used to determine appropriate sampling rates, and properties of certain classes of signals allow for de-noising or detection.

One way to characterize a signal's structure is how many large values it has when represented in an orthogonal basis Ψ . A *compressible* signal can be represented either exactly, or very well, by a linear combination of a few of the vectors from Ψ . In this thesis, we investigate how the compressibility of a signal can be applied to the phase retrieval problem to aid in signal recovery, and to recover signals from fewer measurements than are traditionally needed. This chapter gives a background of both compressive sensing and phase retrieval. In Chapter 2 we introduce a new compressibility constraint for phase retrieval. In Chapter 3 we show how signal compressibility allows for phase retrieval with fewer measurements. In Chapter 4 we apply both of these concepts to terahertz imaging, and we conclude with Chapter 5.

1.1 Compressive Sensing

Traditionally, digital signal processing involves first uniformly sampling a signal and then processing it in some way that enhances it and/or prepares it for storage or

transmission. The sampling part of this process is governed by the signal's Nyquist rate. The sampling rate must be twice the bandwidth in order to accurately represent the signal.

There are many scenarios in which the signal may have a large bandwidth, but not contain much information. A piecewise smooth signal may have high frequency components necessitating many samples, but can be represented well by a linear combination of only a few wavelets. In digital photography an image field may be sampled at 10 million locations, but this information can be effectively stored with only 100 thousand DCT or wavelet coefficients. *Compressive sensing* (CS) takes the logical step of exploiting a signal's structure to acquire it in less measurements, rather than the observe the whole thing and compress it later. As a model for structured signals, we first consider compressible signals that are k -sparse in some orthogonal basis Ψ . This means that when represented in the basis, the signal has only k non-zero coefficients, where k is much smaller than the signal length N . We are interested in recovering these signals exactly, with as few measurements as possible. The main result of CS is that such signals can be reconstructed perfectly with only $O(k \log(N/k))$ measurements.

The first half of CS is the definition of a special linear, nonadaptive measurement scheme. For a given k -sparse signal $x \in \mathbb{R}^N$, the measurements can be represented mathematically as

$$y = \Phi x. \tag{1.1}$$

The matrix Φ is of size $M \times N$, with $M < N$. We would like this matrix to have the property that no two k -sparse signals can result in having the same measurements y . Mathematically speaking, Φ must be an injective mapping for k -sparse signals. It has been shown that with high probability, random matrices of i.i.d Gaussian or ± 1

entries satisfy this property for sufficiently large M [1, 2, 3]. A third type of Φ matrix that also satisfies the property is a matrix formed from random rows of the discrete Fourier transform (DFT) matrix [2]. The measurements y would therefore merely be a random collection of the Fourier coefficients of x . Unlike with the Gaussian or ± 1 matrices, the measurements y made with the random DFT matrix will be complex-valued. Although these Φ matrices allow for CS measurements of signals that are k -sparse in any orthogonal basis, we will focus on the canonical basis in our discussion of CS recovery.

The second half of CS is recovering a signal \hat{x} from the measurements $y = \Phi x$. This is an ill-posed problem, since an infinite number of potential solutions all will admit the given measurements. However we will choose as \hat{x} the one that has the sparsest representation. To find this we perform the optimization

$$\hat{x} = \arg \min_x \|x\|_0 \text{ s.t. } y = \Phi x. \quad (1.2)$$

This optimization of the ℓ_0 pseudo-norm ($\|x\|_0$ is the number of non-zero elements of x) finds, among all signals that satisfy the linear measurements, the signal that has the fewest number of non-zero elements. Such an optimization is combinatorial and therefore impractical. Instead, we consider the optimization

$$\hat{x} = \arg \min_x \|x\|_1 \text{ s.t. } y = \Phi x, \quad (1.3)$$

where the ℓ_1 norm of a signal is defined as the sum of the absolute values of its components:

$$\|x\|_1 = \sum_i |x_i|. \quad (1.4)$$

This is a convex optimization that can be formulated as a linear program [4]. The attractive property of CS Φ matrices is that this polynomial-time optimization procedure, under the CS measurement scheme, yields perfect recovery of k -sparse signals

when Φ takes $O(k \log(N/k))$ measurements. If x is not sparse, but can be represented well by its k largest components in some basis, then the error of the reconstructed signal \hat{x} is only a constant times the error between x and its k -term approximation.

If the ℓ_1 norm of the signal is given *a priori*, it can be used as a projection constraint in a projection onto convex sets (POCS) procedure [5]. Since the ℓ_1 ball and the hyperplane of potential solutions intersect at the optimal point \hat{x} , this approach has the same accuracy as ℓ_1 optimization, but with less computational complexity. It also has similarities to methods currently used for phase retrieval.

1.2 Phase Retrieval

In many scientific fields, including crystallography, astronomy, and wavefront sensing, measurements of a complex-valued signal must be made with sensors that can only observe its intensity. These magnitude-only measurements are acceptable in instances like photography, since our eyes are only sensitive to the intensity of a light field. However in certain applications, the phase of a signal is very important information.

One such important application is when the signal in question is an object's diffraction pattern. With visual light, objects can be illuminated and then recovered with a lens and sensing media. When other parts of the electromagnetic spectrum are used for illumination, lenses must be prohibitively large or small. A microscopic collection of atoms cannot be imaged with visible light, whose wavelength is too large to resolve atoms, but may be illuminated by X-rays that have a much smaller wavelength [6]. Rather than use a tiny lens and exposure media, the X-rays pass through the specimen and are collected downstream with a CCD. The signal acquired is the specimen's *diffraction pattern*. Other useful applications involving diffraction patterns include a landscape aerially illuminated by a laser [7] and a mask illuminated by terahertz




Figure 1.1 : *The phase of an object's Fourier transform is necessary if an inverse transform is to be performed. An (a) image of the earth is transformed into the frequency domain using the fast Fourier transform. It is then reconstructed using (b) magnitude only and (c) phase only information from the transform.*

(THz) rays [8, 9].

An object's diffraction pattern closely approximates the object or scene's Fourier transform. Sensors such as CCD's, and in some cases, THz receivers, can only observe the magnitude of the pattern. We wish to recover the object from its transform, but we cannot perform an inverse transform without the phase of the Fourier coefficients (see Figure 1.1). *Phase retrieval* is the process of recovering the phase of a signal's Fourier transform when given just the magnitude, thereby recovering the signal itself. In our discussion we consider two-dimensional signals, though all of the concepts extend to higher dimensions.

1.2.1 The Phase Retrieval Problem

For a complex-valued signal x_a sampled to be of size $N \times N$, phase retrieval can be defined mathematically as the search for the intersection of two signal constraint sets:

- The **Fourier Modulus Matching Constraint Set** is the set of signals whose

Fourier transform modulus match the observed Fourier transform modulus of x_a :

$$S_F = \{x \in \mathbb{C}^{2N \times 2N} \text{ s.t. } |\mathcal{F}x| = |X_a|\}, \quad (1.5)$$

where $|X_a|$ is the modulus of the Fourier transform of the signal x_a . The 2-D discrete Fourier transform is defined as:

$$\mathcal{F}x = X[u, v] = \frac{1}{2N} \sum_{m=0}^{2N-1} \sum_{n=0}^{2N-1} x[m, n] e^{-i(\frac{2\pi um}{N} + \frac{2\pi vn}{N})}. \quad (1.6)$$

Because $|X_a|$ is related to the autocorrelation x_a via a Fourier transform, and because this autocorrelation is, in general, of size $(2N - 1) \times (2N - 1)$, $|X_a|$ must also be of size $(2N - 1) \times (2N - 1)$ to avoid aliasing.

- **The Non-Aliasing Constraint Set** are signals of size $2N \times 2N$ that are nonzero only in an $N \times N$ window. Such signals do not alias their autocorrelations, and by extension, their Fourier transform modulus. We denote this constraint set S_{NA} .

The intersection of these two sets contains signals of size $2N \times 2N$ whose Fourier transform modulus is $|X|$, but are nonzero only in an $N \times N$ window, ensuring that their Fourier transform modulus was not aliased. Surprisingly, such a signal in the set intersection is equivalent to x_a^* , excepting a shift in space, flip about one of the axes, or a constant phase shift, since this information is irretrievably lost with the phase. This result is a consequence of the irreducibility of arbitrary complex polynomials in two or more dimensions [10].

*That is, the nonzero $N \times N$ portion of it is equivalent to x_a .

1.2.2 Phase Retrieval Methods

The goal of phase retrieval is to find the intersection of the two sets. If the two sets were convex, alternatively projecting the two would find the nearest intersection. However, the set S_F is non-convex, and the distance between the two sets has many local minima that projection algorithms could get stuck in. Despite this fact, Fienup discovered that projection methods could still reliably find the intersection and recover the original signal [11, 12]. He introduced several algorithms and evaluated their ability to recover real, positive-valued signals from their Fourier transform modulus:

- The **Error Reduction (ER) Algorithm** starts with an initial value $x^{(0)}$, usually chosen randomly, and iterates in the following way:

$$x^{(n+1)} = P_C P_F x^{(n)} \quad (1.7)$$

The operation P_F projects the current iterate $x^{(n)}$ onto the set S_F by taking the Fourier transform of $x^{(n)}$, fixing the magnitudes to the known values while leaving the phase unchanged, and taking the inverse Fourier transform of the result. The operation P_C projects the result of the Fourier projection onto the set S_{NA} by setting to zero the values of the signal outside an $N \times N$ window. At each step, the algorithm finds an iterate on the set S_{NA} which is the same distance or closer to the set S_F than the previous iterate. This is a favorable property until a local minimum between the sets is reached because the algorithm cannot move away from it.

- The **Hybrid Input-Output (HIO) Algorithm** uses the same projections but applies feedback in an effort to keep from getting stuck in local minima.

$$x^{(n+1)} = \left\{ \begin{array}{ll} P_F x^{(n)} & \text{pixels of } P_F x^{(n)} \text{ in } S_{NA} \\ 1 - \beta P_F x^{(n)} & \text{pixels of } P_F x^{(n)} \text{ not in } S_{NA} \end{array} \right\} \quad (1.8)$$

The value of β is usually chosen to be between .9 and 1. While the algorithm is not guaranteed to minimize the distance between the two sets at each step, its feedback mechanism allows it to wander out of local minima in pursuit of the set intersection.

Though it has been more than 20 years since Fienup introduced these algorithms, various projection methods are still the state of the art in phase retrieval [13]. New projection schemes such as the Difference Map [14] and Relaxed Averaged Alternating Reflections (RAAR) [15] continue to be developed. RAAR operates according to the iterative process

$$\hat{x}^{(n+1)} = \left[\frac{1}{2} \beta (R_C R_F + I) + (1 - \beta) P_F \right] \hat{x}^{(n)}. \quad (1.9)$$

where a reflection $R(x)$ is defined as $2Px - x$, and β is usually between .8 and 1.

1.2.3 Signal Constraints

No matter how sophisticated present and future projection algorithms may be, all are subject to the fact that the problem is a non-convex optimization. Fienup and other's methods are only reliable if the set S_{NA} is further constrained, thus reducing the search space. In his seminal work, Fienup considers the recovery of only the real, positive signals [12]. A positivity constraint greatly reduces the size of the set S_{NA} and makes the search considerably easier. Other constraints used to further constrain the set include histogram matching and number of nontrivial (usually nonzero) pixels [16]. In this thesis we will examine, in particular, a **Strict Support Constraint Set**. This constraint set, denoted S_S , is the intersection of S_{NA} with signals that are nonzero on the same support as the original signal. For phase retrieval of complex signals derived from an X-ray diffraction pattern, it is possible to assume and apply

a positivity constraint to the imaginary part of the signal [17]. However, for phase retrieval of general complex signals, it has been observed that for reliable recovery the strict support constraint must be used [18], whether it be known in advance or discovered [19].

Chapter 2

Compressive Phase Retrieval

Needing to know or discover a signal’s support for phase retrieval of complex signals is often too much to ask, so our aim is to develop a more realistic—and hopefully more effective—signal prior for phase retrieval. Our key insight is that many natural signals are compressible in an orthogonal basis Ψ , meaning that they can be represented either exactly or very well by a linear combination of only $k \ll N \times N$ vectors in the basis Ψ . For example, most images are compressible in a wavelet basis. This property is what enables image compression techniques like JPEG and JPEG 2000. Other images, such as a starry sky, are compressible in the canonical basis (Ψ is the identity matrix).

2.1 ℓ_p Compressibility Constraint

Rather than assume that we know the support of $x_{\mathbf{a}}$, the signal to be recovered, we instead assume that the signal is compressible and that we know its ℓ_p norm in the basis Ψ . To leverage this information for phase retrieval we introduce a new *compressibility constraint*:

The ℓ_p **Compressibility Constraint Set**, denoted S_p , limits the search space S_{NA} to be the subset in which the signals have the same ℓ_p norm when represented in the basis Ψ as that of the original signal:

$$S_p = \{x \in S_{NA} \text{ s.t. } \|\Psi^{-1}x\|_p = \|\Psi^{-1}x_{\mathbf{a}}\|_p\},$$

where the ℓ_p norm is defined as

$$\|x\|_p = \left(\sum_i |x_i|^p \right)^{\frac{1}{p}}. \quad (2.1)$$

When $p = 2$, the constraint set considers signals with the same amount of energy as the original signal. For $p = 1$, the constraint set $S_{p=1}$ is the subset of S_{NA} where the signals have the same ℓ_1 norm in the basis Ψ as the original signal x_a . For the limiting case of $p = 0$ (ℓ_0 is technically not a norm), the set $S_{p=0}$ is the subset of S_{NA} where the signals have the same number of nonzero coefficients as x_a in the basis Ψ . In the specific case where Ψ is the canonical basis, $S_{p=0}$ is equivalent to the existing number of nontrivial pixels constraint set [16]. In this chapter we demonstrate which of these values give S_p favorable properties in comparison to S_S , the best known constraint set for phase retrieval of complex-valued signals.

Using S_p as a constraint requires specific projection operations in conjunction with an algorithm like ER. ER works by projecting between the Fourier modulus set and the constraint set, reducing at each step ($n = 1, 2, 3, \dots$) the distance between them:

$$E = \frac{\|P_C P_F \hat{x}^{(n)} - P_F \hat{x}^{(n)}\|_2}{\|P_F \hat{x}^{(n)}\|_2}. \quad (2.2)$$

$P_F \hat{x}^{(n)}$ is the Fourier projection operator. The projection onto a constraint set is denoted $P_C x^{(n)}$. It finds the signal in the constraint set that is closest to $x^{(n)}$. For S_{NA} , it sets to zero all pixel values outside an $N \times N$ window. Additional constraints perform another projection:

S_S it also zeros out any signal values outside of the prescribed support. For $S_{p=1}$ the projection is accomplished by uniformly adding or subtracting a constant value to the magnitudes of the entries of $\Psi^{-1}x$ until $\|\Psi^{-1}x\|_1$ reaches the desired value [5]. For $p = 0$, only the prescribed number of nonzero coefficients are kept while the rest are set to zero, and for $p = 2$ the signal is normalized to have the same energy as x_a .

To visually demonstrate the effectiveness of S_p on a compressible signal, we test it on an image that has 4096 pixels, all nonzero, but has only 300 nonzero coefficients with Ψ being the Daubechies-6 wavelet basis (Figure 2.1). The magnitude of its Fourier transform is recorded, as is the ℓ_1 norm of its wavelet coefficients. For both $S_{p=1}$ and S_S , ER is performed with multiple starting points, followed by several hundred iterations of HIO on the result with the smallest E value. With $S_{p=1}$, the algorithms recover the signal exactly, while with S_S they cannot.

2.2 Geometric Analysis of Compressibility Constraint

To more rigorously compare S_p to S_S , we consider the propensity of each set to introduce local minima in the distance between it and the Fourier modulus constraint set S_F . We reason that the fewer local minima exist between the sets, the more effective the constraint will be for any projection algorithm. In order to analyze these properties, we examine the density of global minima among all minima through a Monte Carlo approach. For a range of signal sizes and number of nonzero elements, we generate multiple complex-valued signals. For each signal, a random starting value is chosen and then the ER algorithm finds the nearest minimum in the distance between S_F and the constraint set to be examined. The proportion of the times this local minimum is a global minimum (E is sufficiently close to zero) is recorded. By comparing this global minimum ratio in tests on S_S and S_p , we can determine which constraint gives the best geometrical setting for recovering compressible signals.

2.2.1 Signals Sparse in Space

We first evaluate the sets' global minimum properties using compressible signals that are sparse in the canonical basis (i.e., there are only a few nonzero values). For




Figure 2.1 : (a) The 64×64 image (zero-padded to be 128×128) has all pixels of nonzero value, but is composed of only 300 Daubechies-6 wavelets. (b) The magnitude of its Fourier transform is recorded. (c) $S_{p=1}$ is used to recover the signal exactly, but (d) the support constraint is ineffective. (For (c) and (d), only the 64×64 nonzero portion of the recovered signal is shown)

a given signal length we vary k , the number of nonzero elements, and determine over 2000 trials what the global minimum ratio is for the distance between the S_F and S_S/S_p (Figure 2.2). S_S has the highest, and therefore most favorable, ratio of global minima. The ratio for $S_{p=1}$ is good considering it only uses the ℓ_1 norm as a constraint, while S_S requires the exact support. It is surprising that $S_{p=1}$ does better than $S_{p=0}$, because it would seem that for a signal with only a few nonzero elements, using that information as prior information (using $S_{p=0}$) would help recover the signal more than just knowing its ℓ_1 norm. However, the ℓ_1 ball being convex and the ℓ_0 ball being highly non-convex may explain why $p = 1$ introduces less local minima and is therefore preferable to $p = 0$. This may also explain why $S_{p=2}$ does not do much better than the basic non-aliasing set S_{NA} . Unlike for $p=0$ or 1, the ℓ_2 ball is very smooth. At the point of its intersection with S_F , there will be many more local minima in the distance between the two sets than compared to the compressibility constraint with the other p values.

Regardless of the constraint, signals that have more nonzero elements introduce more local minima in the search. When using a support constraint S_S , this occurs because the signal space increases. The same applies for $S_{p=0}$. However, for $S_{p=1}$, the number of nonzero elements has no effect on the size of the signal space, which is dictated by the ℓ_1 norm value. Rather than changing the size of the search space, the signal having more nonzero values places it on a less “pointy” location on the ℓ_1 ball. The intersection of this location on the ball with S_F obviously has more local minima in its vicinity. In contrast, a very sparse signal is located on one of the points of the ball, and the intersection of it with S_F is much more favorable.

Keeping k constant and varying N has a minimal effect on phase retrieval recovery (Figure 2.3). This is because if the number of nonzero pixels remains the same, then

changing the overall signal size has no effect on S_S or $S_{p=0}$. For $S_{p=1}$, increasing the signal size has only a slightly adverse effect on the number of local minima if k is held constant.

2.2.2 Signals Sparse in Ψ

When signals are sparse in another basis, such as wavelets, the compressibility constraint excels. As the number of nonzero coefficients increases, the number of nonzero pixels increases at a much quicker rate than before, and the performance of S_S correspondingly degrades. However, $S_{p=1}$ is affected by the number of nonzero coefficients, and not necessarily nonzero pixels. As a result, it is much more effective than S_S when the number of significant coefficients in the basis Ψ is small but the number of nonzero pixels is large. To examine the limiting case where each vector in the basis Ψ is entirely nonzero, we create a basis with vectors drawn from a Gaussian distribution, and then orthogonalize it. S_S is useless for recovering signals generated from this basis because every pixel value is nonzero and the constraint does not limit S_{NA} at all, but $S_{p=1}$ performs as well as if the signals were sparse in space.

In contrast to S_S , the compressibility constraint performs well regardless of the basis Ψ . However, it is not entirely unaffected by different bases. $S_{p=0}$ and $S_{p=1}$ have most favorable phase retrieval geometries for signals that are sparse in space, or are formed by a randomly generated basis. The next best basis for recovery is Haar wavelets, followed by Daubechies-6 wavelets. Unlike with S_S , the performance clearly has nothing to do with the number of nonzero pixel values, for the best two cases are opposite extremes. Instead, the performance appears to be related to the degree of Ψ 's coherence with, or similarity to, the Fourier basis. Spikes, followed by random vectors, are most incoherent with the Fourier basis, followed by Haar and

Daubechies-6 wavelets. Likewise, signals that are sparse in space or in a random basis perform best, followed by wavelets.

2.2.3 Non-Sparse Compressible Signals

In reality, compressible signals are never exactly sparse. Rather, if the magnitude of their coefficients in Ψ were sorted, the values would closely follow an exponential decay. To test the compressibility constraint on this signal class, we generate signals in which the magnitudes of the Daubechies-6 wavelet coefficients α_i have an exponential decay:

$$\alpha_i = \frac{1}{i^n}.$$

We evaluate the performance of $S_{p=1}$ for different values of n . The larger n is, the more quickly the coefficients decay, and the closer the signal is to being truly sparse. As shown in Figure 2.4, the performance is best for larger values of n , but gracefully degrades as n decreases. In all cases S_S and $S_{p=0}$ are totally ineffective because all of the wavelet coefficients, and pixel values, are nonzero.

2.3 Benefits of the Compressibility Constraint

There are many sophisticated phase retrieval algorithms in existence, and we have not provided a new one. Instead, we propose a new kind of constraint. For compressible signals, the constraint set based on the ℓ_p norm of a signal is effective for phase retrieval of complex-valued signals. In particular, the geometry of S_p with $p = 1$ introduces fewer local minima in the distance between it and S_F than does S_S or other values of p when Ψ is a wide-band basis such as wavelets. In such cases, $S_{p=1}$ is even more effective than S_S for phase retrieval. As a result, existing and future phase retrieval methods can perform *compressive phase retrieval* (CPR) by leveraging the




Figure 2.2 : For 8×8 signals that are (a) sparse in space, S_S offers the best hope of recovering the signal, followed by $S_{p=1}$, and then $S_{p=0}$. For signals that are sparse in a wide-band basis such as (b) Haar wavelets and (c) Daubechies-6 wavelets, $S_{p=1}$ is most effective at signal recovery. (d) If each vector of Ψ is generated randomly and is entirely nonzero, $S_{p=1}$ and $S_{p=0}$ are the only effective constraints. In all cases, $S_{p=2}$ is only marginally more effective than using S_{NA} without an additional signal constraint.




Figure 2.3 : For a constant sparsity, changing the signal size $N \times N$ has little effect on the ability to perform phase retrieval.




Figure 2.4 : $S_{p=1}$ performs best when its coefficients in the basis Ψ (in this case, Daubechies-6 wavelets) decay quickly. As the decay moves from i^{-4} to i^{-1} , the performance degrades but is still acceptable. The signal size is 8×8 .

a priori information of a compressible signal's ℓ_p norm to better their performance. In practice, the knowledge (either exactly or approximately) of a signal's ℓ_p norm will not always be available. However, the search for a signal's ℓ_p norm would be a one-dimensional scan, unlike searching for a signal's support.

Chapter 3

Compressive Sensing Phase Retrieval

In many data acquisition and processing problems, the measurement process is time consuming or expensive. For example, high-frequency and ultra wide-band radar is beginning to test the limits of analog to digital conversion. Sensors for THz electromagnetic frequencies cost many times more than those for the visible spectrum. Magnetic resonance imaging (MRI) scans can take over an hour. The irony of these situations is that in each case the wide-band signals involved are undeniably structured, yet aside from general bandwidth considerations, no assumptions about this structure are exploited to ease the sensing process. CS takes the logical step of making measurements that directly take a signal's structure into account. CS takes only as many measurements as the structure dictates. By structured, we mean that the signal has only a few non-zero coefficients when represented in terms of some basis, or can be approximated well by a few coefficients. CS has already seen practical success such as a camera with a single pixel, [20] reduction in MRI scan time, [21] and analog-to-digital conversion [22].

Several types of measurement schemes have been shown to capture the information necessary for CS. They all involve random projections—taking the inner product of the signal with random vectors such as Gaussian vectors, ± 1 Bernoulli/Rademacher vectors, or randomly selected rows from the DFT matrix. Since random DFT coefficients qualify as CS measurements, it would seem that the benefits of CS can apply when we have access to an object's diffraction pattern, since it closely approxi-

mates the image of an object's Fourier Transform. Taking the object's structure into account, we can recover the object from just its randomly sub-sampled diffraction pattern, rather than the entire pattern [8]. We refer to recovering signals from fewer Fourier magnitude measurements as *compressive sensing phase retrieval* (CSPR).

3.1 Sufficient Measurements for a Band-Limited Signal

As discussed in Chapter 2, it is known that a signal's Fourier transform modulus, over-sampled by a factor of two in each dimension, is sufficient information to uniquely recover the signal. The theory behind this uniqueness of recovered results has been developed by Hayes [10]. Observing the magnitude (squared) of a signal's Fourier transform is equivalent to observing its autocorrelation, since they are Fourier transform pairs. Knowing a signal's autocorrelation is equivalent to knowing the z -transform of its autocorrelation. This polynomial is the product of the signal's z -transform with the transform of its time-reversed version. Therefore if a signal's z -transform is irreducible, then it is the only signal that will yield its autocorrelation (excepting a shifted or flipped version, since the absolute position and orientation of the signal is irretrievably lost with the phase). Virtually all polynomials in two or more dimensions are irreducible, since the set of reducible polynomials is a set of measure zero [23]. Therefore if a candidate recovered signal has finite support and has the same autocorrelation as the original signal, then it is indeed the correctly recovered original signal.

Hayes also explains the conditions needed to guarantee that a recovered signal has the same autocorrelation as the original signal [10]. In examining sufficient conditions to guarantee that two signals' autocorrelations are equivalent, he comes to the conclusion that a signal in two or more dimensions is uniquely specified by its twice

over-sampled DFT magnitude.

3.2 Sufficient Measurements for a Sparse Signal

For a band-limited signal, the number of measurements needed to guarantee uniqueness in PR is a function of the bandwidth: the Fourier magnitude must be sampled twice as much in each dimension as the bandwidth of the original. However we know more about the signal than this. We know that natural signals will probably have a relatively small number of significant coefficients in a sparsifying basis, compared to the total size of the signal. CS intuition says that if only a few pieces of information of the signal suffice to represent it well, then only a few measurements should be needed to capture this information. We have access to linear measurements of the signal's autocorrelation, as the intensity (the magnitude-squared) of a signal's Fourier transform is equivalent to the Fourier transform of its autocorrelation. Each Fourier magnitude measurement is a linear projection of the signal's autocorrelation onto a complex sinusoid. The good news is that a random collection of these projections is sufficient to specify a sparse autocorrelation, which is sufficient to uniquely specify the signal.

Lemma 3.1

Suppose $x[n_1, n_2]$ is a two-dimensional sequence of complex-numbers of support $N \times N$, and x has a z -transform that, except for trivial factors, is irreducible and non-symmetric. Then $x[n]$ is uniquely specified by its Fourier transform magnitude, and a $2N \times 2N$ point DFT is sufficient for this unique specification.

Proof 3.1 Hayes proves that sequences of real numbers are uniquely specified by their DFT magnitude, as a consequence of their z -transforms being irreducible [10]. These arguments also follow for complex sequences as long as their z -transforms are also

irreducible. Hayes notes that the irreducible requirement is not strict in two or more dimensions for complex signals, since reducible polynomials correspond to a set of measure zero [23]. \square

This means that if $x \in \mathbb{C}^{N \times N}$ has an irreducible z -transform and there exists an $\hat{x} \in \mathbb{C}^{N \times N}$ such that $|\mathcal{F}\hat{x}| = |\mathcal{F}x|$ on a $2N \times 2N$ lattice, then \hat{x} and x are equivalent. Since the autocorrelation is a Fourier transform pair with a signal's Fourier transform intensity, Lemma 3.1 implies that a signal's autocorrelation is sufficient to uniquely specify it.

Since an arbitrary complex signal is specified by its autocorrelation, we must now consider how many measurements of the signal's autocorrelation are needed to specify it. We consider the case in which the input signal is k -sparse.

Theorem 3.1

Suppose $x[n_1, n_2] \in \mathbb{C}^{N \times N}$ is k -sparse and has an irreducible z -transform. Then with probability of at least $1 - O(N^{-4\rho/\alpha})$ for some fixed $\rho > 0$, $\frac{1}{\alpha}k^2 \log(4N^2/k^2)$ random Fourier magnitude measurements of x are sufficient to uniquely specify it.

Proof 3.2 The autocorrelation of x is of size $2N \times 2N$, which is $4N^2$ total pixels. Since x is k -sparse, the autocorrelation is at most k^2 -sparse. Due to the Uniform Uncertainty Principle of a random Fourier ensemble [3], only $M = \frac{1}{\alpha}k^2 \log(4N^2/k^2)$ random Fourier measurements are needed to specify a k^2 -sparse signal of size $4N^2$, with the probability stated above in the theorem. This means that, with overwhelming probability, if autocorrelation \widehat{R}_x matches R_x at M random Fourier locations, and both \widehat{R}_x and R_x are k^2 -sparse, then $\widehat{R}_x = R_x$. From Lemma 3.1, this implies that \hat{x} and x are equivalent. \square

This approach redefines the number of Fourier magnitude measurements needed to capture the information in a signal and is significant because it scales with the signal’s structure, rather than its bandwidth. The theorem applies to signals that are sparse in space. It does not apply to signals that are sparse in a general basis Ψ because their autocorrelation would not be k^2 -sparse in another basis, but rather in the frame formed by convolving all of the atoms of Ψ with themselves. However, in Section 3.4 we provide numerical evidence that suggests that the results apply in practice to other sparsifying bases, such as wavelets. We also show that in practice the order of measurements needed scales more like $k \log(N^2/k)$ than $k^2 \log(N^2/k^2)$. This may be partially explained with the fact that a k -sparse signal will have *at most* a k^2 -sparse autocorrelation, and would have a much sparser autocorrelation if the non-zero elements were connected. For example, the autocorrelation will be $(2k - 1)$ -sparse if the non-zero elements form a line.

In order to recover the signal from fewer, but still sufficient, number of measurements, we must perform an inverse operation that is ill-posed. We would like to find a k^2 -sparse solution that matches the given random Fourier measurements. This could be accomplished via the optimization

$$\hat{x} = \arg \min_x \|x\|_0 \text{ s.t. } |y| = |\Phi x|. \quad (3.1)$$

As with CS, such an optimization would be combinatorial in its complexity and therefore infeasible. We turn to CS for inspiration to instead find our solution via the optimization

$$\hat{x} = \arg \min_x \|x\|_1 \text{ s.t. } |y| = |\Phi x|. \quad (3.2)$$

The problem for us is that while ℓ_1 minimization is a convex optimization for CS, it is a non-convex optimization for CSPR due to the non-convex feasible solution set.

Rather than perform this intractable optimization, we turn to the projection strategy of phase retrieval algorithms. As we show in Section 3.4 we have found that this method is effective in recovering sparse solutions with a number of measurements that supports what we have proven to be sufficient.

3.3 CSPR Recovery

We know how many Fourier modulus measurements suffice to recover a k -sparse signal $x_{\mathbf{a}}$, so our task is to find a k -sparse signal that matches the measurements. According to our theorem, this will recover the original signal exactly. As with phase retrieval of band-limited signals from all of the Fourier measurements, we use a projection strategy. Mathematically, we again search for the intersection of two sets.

The first is the **Sub-sampled Fourier Modulus Set**, the set of signals whose Fourier transform modulus match the observed Fourier transform modulus of $x_{\mathbf{a}}$ at a sub-sampled set of locations:

$$S_{F_A} = \{x \in \mathbb{C}^{2N \times 2N} \text{ s.t. } |X(\omega)| = |X_{\mathbf{a}}(\omega)|, \omega \in A \subset \Omega\},$$

where A is the set of frequency locations, chosen randomly from the entire set Ω .

As with regular phase retrieval, the second set is the non-aliasing set S_{NA} . We consider several additional sets that further limit it: a strict support set S_S , and a compressibility constraint set S_p for p values 0, 1, and 2. In our analysis of the various constraints we consider both how they perform in finding a signal in their intersection with S_{F_A} , and if the signal they recover is k -sparse or not.

Using the same method of local minima analysis as in Chapter 2, we test the constraints for a signal size of 64 pixels, 6 of them nonzero, while varying M , the number of measurements used for S_{F_A} (see Figure 3.1). We first consider signals

that are sparse in space. S_S has the fewest local minima in the distance between it and S_{F_A} , followed by $S_{p=1}$ and $S_{p=0}$. The rates of convergence drop as the number of measurements drop and the set S_{F_A} increases in size. There is a critical point around $M = 100$ where the intersection is impossible to find, and below this point the compressibility constraints can recover signals again. For values of M less than 100, $S_{p=2}$ and the anti-aliasing constraint S_{NA} both intersect S_{F_A} with few local minima between them.

While all of the signal constraints intersect favorably with S_{F_A} for certain values of M , not all of the constraints recover k -sparse signals. For values of M less than 100, it is easy to find the intersection of S_{F_A} with S_{NA} , but none of the signals in the intersection are k -sparse, and therefore none are a correct recovery. This is to be expected because the anti-aliasing set does not constrain the search space at all, and while there may only be one 6-sparse signal matching the measurements, there are infinitely more non-sparse signals in S_{NA} that also match the measurements. $S_{p=0}$, on the other hand, does restrict the search over k -sparse signals. S_S restricts the search to k -sparse signals with the same nonzero locations as the original signal. As a result, both retrieve k -sparse signals every time they intersect with S_{F_A} for M greater than 50. $S_{p=1}$ also retrieves only k -sparse signals every time it correspondingly intersects S_{F_A} . The result is surprising because it is only placing a restriction on the ℓ_1 norm of the signals and not on the number of nonzero elements. Not only does it recover sparse signals just like $S_{p=0}$, it does so at a significantly higher rate. As with phase retrieval with all of the measurements, this is likely because it is a convex set, in contrast to $S_{p=0}$.

When Ψ is the Haar wavelet basis, the various properties of the sets are the same, with one exception. The exact support constraint set S_S performs much worse, and

the number of measurements needed for the correct recovery is higher. The fact that the signal is 6-sparse in Haar wavelets means that many more than 6 pixels are nonzero. Because there are more nonzero pixels, the support constraint requires more measurements. $S_{p=0}$ and $S_{p=1}$ both correctly recover signals that are k -sparse in wavelets, needing the same number of measurements to do so as when the signals were sparse in space. This suggests that the CSPR theorem, which has only been proven for signals sparse in the canonical basis, also applies to signals sparse in other bases. In the case of Daubechies-6 wavelets, all signals generated have all nonzero pixel values, so only the compressibility constraints can recover them.

3.4 CSPR Accuracy

As seen in the previous section, we find that with CSPR there exists a critical point in the number of measurements, above which converged signals are accurately recovered. For different signal sizes and sparsity rates, we record the number of measurements M needed for consistent (95%) exact recovery of 100 converged solutions. We hold N constant and vary k , and also hold k constant and vary N , in order to empirically understand the dependence of the number of measurements on these values. We compare these results with those found via regular CS if the phases were known, using the SPGL1 solver [24].

For convergent solutions we find that the number of measurements needed does not appear to follow a $k^2 \log(N^2/k^2)$ trend but appears to be closer to $k \log(N^2/k)$, as Figure 3.2 shows. When the signal size is held constant, the number of measurements needed increases linearly. The slope is the same as for CS with the known phases, though more measurements are needed for CSPR. When k is held constant, the number of measurements follows a sub-linear trend, just as CS does, though more

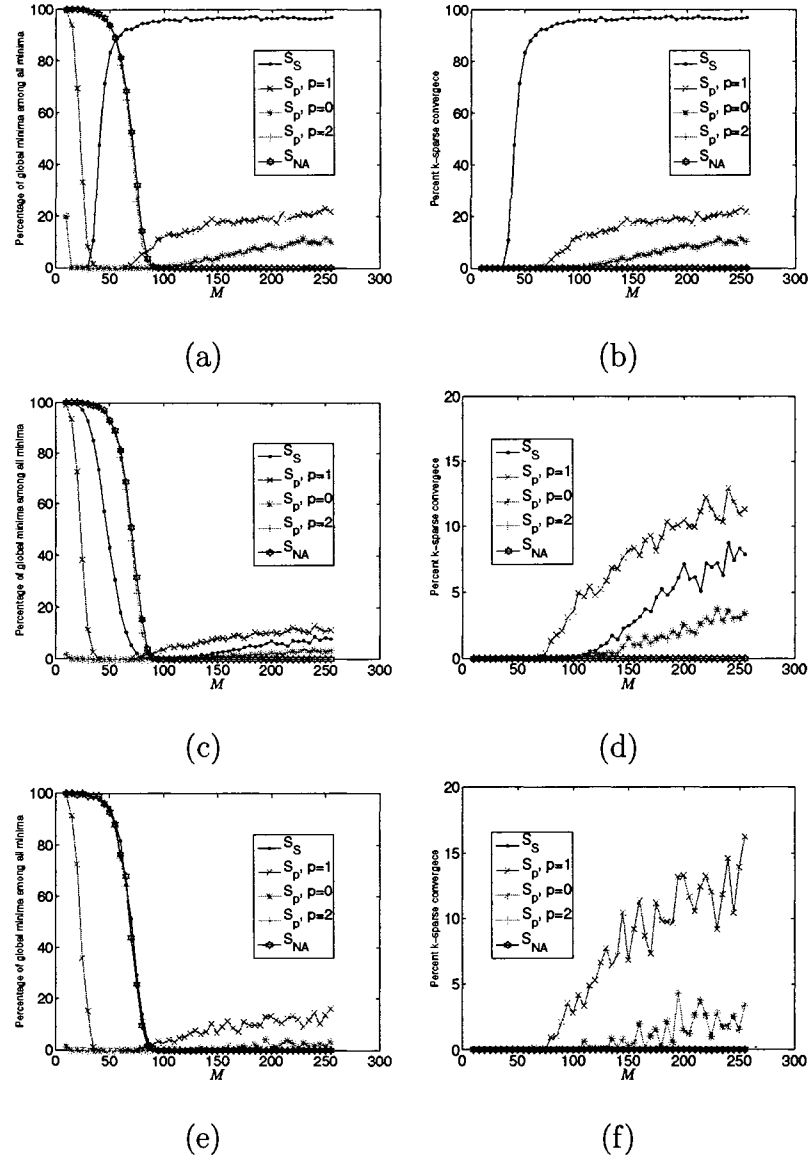


Figure 3.1 : The percentage of global minima among all minima in the distance between S_F and various constraint sets varies with the number of measurements used. (a) For signals that are sparse in space, S_S has the highest density of global minima and is very effective for phase retrieval, followed by $S_{p=1}$ and $S_{p=0}$. For more wide-band bases such as (c) Haar or (e) Daubechies-6 wavelets, the compressibility constraints outperform the support constraint. The compressibility constraints intersect favorably with S_F , and the signals they recover are, without fail, sparse (and correct). As long as the number of measurements is above a critical point, signals are accurately recovered whether they are sparse in (b) space, (d) Haar wavelets, or (f) Daubechies-6 wavelets.




Figure 3.2 : For a given signal size $N \times N$ and sparsity level k , enough trials of CSPR are performed (with the compressibility constraint $S_{p=1}$) on randomly generated signals until 100 convergent solutions have been found. The number of measurements M recorded is the smallest value needed so that at least 95 are perfectly reconstructed. In (a) we hold $N \times N$ constant at 64 and vary k . The number of measurements needed increases linearly, whether the signal is sparse in the canonical basis (I) or the 2D Haar wavelet basis (W). In (b) we hold k constant at 5 non-zero elements (in the canonical basis) and vary the signal size. The increase in the number of measurements needed is approximately linear when plotted on a log scale, and thus it is approximately logarithmic. In each case the number of measurements needed has a trend similar to that of CS with the phases known, though clearly more measurements are required.

measurements are needed than CS would need if the phases were known.

Both of these findings support the CSPR theorem result regarding the number of Fourier modulus measurements sufficient to specify signals that are k -sparse in space. In addition, we also demonstrate in Figure 3.2 that the result appears to hold for signals that are sparse in another basis. When the signal size is held constant, roughly the same number of measurements are needed to recover signals of varying sparsity, whether they are sparse in space or in wavelets. These empirical findings are significant because they suggest that random Fourier modulus measurements may have the CS quality of capturing the information in a signal regardless of the basis in which it may be sparse.

Chapter 4

Compressive Phase Retrieval in Practice

4.1 THz Imaging

The ability to use compressibility to aid in phase retrieval, and perform it with less measurements, improves the recovery of any signal from its diffraction pattern. A particular area of application we investigate is THz imaging. With uses in aerospace, homeland security, medical imaging, and quality control of packaged goods, time-domain THz imaging systems have proven valuable in numerous fields. These systems are generally limited by slow image acquisition rates [25]. In the fastest example of raster-scan THz imaging reported to date, a 400×400 pixel image takes as long as 6 minutes to acquire [26]. Recent developments using more sophisticated image processing approaches, such as the Radon transform [27, 28] and interferometric imaging [29], have shown preliminary successes but also face similar limitations in speed, resolution and/or hardware requirements. Applying compressive phase retrieval allows for THz imaging with less measurements, and hence shorter acquisition time.

Our imaging system consists of a pulsed THz transmitter and receiver, both based on photoconductive antennas, and two lenses, one of which approximately collimates the THz beam while the other focuses the beam (see Figure 4.1). The object mask, placed in between the two lenses, scatters the THz waves. The focusing lens forms the Fourier transform of the object mask at its focal plane. The receiver, mounted on a translation stage, performs a raster scan in the focal plane, over an area of 64×64


mm, at 1 mm intervals. We place a circular aperture (1 mm in diameter) in front of the receiver antenna so that it only samples a small area of the Fourier pattern, rather than relying on the 6 mm receiver aperture [30]. The object mask is made of opaque copper tape on a transparent plastic plate. In our experiments, our object mask is a T-shaped hole, 15 mm by height and by width.

At each detector position, an entire time-domain THz waveform is measured. If the object is exactly 6 cm from the focusing lens, it is possible to observe both the magnitude and phase of the diffraction pattern in the Fourier plane. We compute the power spectrum of each waveform, and select the spectral amplitude and phase at a particular wavelength ($\lambda = 2$ mm) to obtain a (complex) pixel value. In this way, we can assemble a 64×64 Fourier diffraction pattern. Direct 2D Fourier inversion of this signal can reconstruct the object.


4.2 CPR Recovery from the Entire Diffraction Pattern

In the more realistic scenario where the object to be measured is not exactly one focal length distance away from the focusing lens, the acquired Fourier data will have the correct Fourier magnitude but a distorted phase. We can no longer use the phase data to simply perform an inverse Fourier transform and reconstruct the image, so we attempt to recover it via PR. As Figure 4.2(a) shows, the signal cannot be recovered with the basic non-aliasing constraint S_{NA} . Knowing what the true object is, we can make a strict support constraint. Using this constraint S_S with the measured Fourier modulus, we can recover the object via the RAAR phase retrieval algorithm (Figure 4.2(b),(c)).

In our experiment, the signal to be recovered is highly compressible in space, as there are only a few significant pixels values (see Figure 4.3). As described in Chapter



(a)



(b)

Figure 4.1 : (a) The THz Fourier imaging setup. An approximately collimated beam from the THz transmitter illuminates an object, placed one focal length away from the focusing lens. The THz receiver raster-scans and samples the Fourier transform of the object on the focal plane. In the event that the receiver is not exactly one focal length away from the lens, the phase is distorted. (b) The observed diffraction pattern.




Figure 4.2 : From the observed THz diffraction pattern, (a) PR is unsuccessful at recovering the signal with only the non-aliasing constraint S_{NA} . (b) An exact support constraint S_S is needed (c) to recover the signal. With the ℓ_1 norm calculated from the result, (d) the signal can be recovered using the $S_{p=1}$ compressibility constraint.




Figure 4.3 : When sorted, the magnitudes of the “T” signal used in our experiments have a sharp decay, indicating the signal is highly compressible.

2, such a signal can be recovered using a compressibility constraint, instead of the exact support constraint. Rather than use the object’s support to aid in recovering it, we use its ℓ_1 norm. For our experiment, we calculate the ℓ_1 norm value obtained from the recovery via S_S , and use it in $S_{p=1}$ to reconstruct the signal (Figure 4.2(d)). The fact that the signals recovered via S_S and $S_{p=1}$ are nearly identical is remarkable, given that one was given the exact support of the actual signal as prior information, while the other had only the ℓ_1 norm as a constraint.

We derive the value of the signal’s ℓ_1 norm by recovering the signal first with S_S . In most applications, one will not have the luxury of having the exact value of the norm of the signal to be recovered. However, even if the value of the actual signal’s ℓ_1 norm is not known exactly, having a good estimate suffices if the signal is sparse enough. Figure 4.4 shows the signal recovered with the true ℓ_1 norm as a constraint,

along with signals recovered with other norm values used as constraints. A large range of values can be used to obtain acceptable results.

These reconstruction results demonstrate the applicability of our imaging scheme not only to pulsed THz imaging systems but also to continuous-wave systems, in which phase information is typically not available. The latest research on THz CW imaging requires high-power sources (>10 mW), such as quantum-cascade lasers (QCLs) operating at low temperature (~ 30 K), because the focal-plane microbolometer array used for imaging has low sensitivity at THz frequencies [31]. In contrast, our Fourier imaging technique can use a single-pixel THz detector with much higher sensitivity to enable imaging with a low-power CW source.

4.3 CSPR Recovery from a Sub-sampled Diffraction Pattern

The compressibility of a THz signal not only allows for phase retrieval recovery when the phase is distorted, but also recovery from less measurements. We reconstruct the signal with varying numbers of measurements, randomly sub-sampled from the entire Fourier modulus set. As Figure 4.5 shows, signal quality remains high for as low as 600 measurements, at which point it seriously degrades. Because none of the signals exactly converged, we cannot accurately compare the error between them directly. Rather, we calculate the error between their Fourier transform modulus and the true Fourier transform modulus. The error decreases with a corresponding increase in measurements (Figure 4.6).

Using less measurements with THz imaging system designs can significantly reduce the image acquisition time. Traditional imaging systems scan point by point in the space domain. By application of the CSPR theorem, only $O(k^2 \log(4N^2/k^2))$ locations in the frequency (modulus) domain need to be scanned if the signal is sparse




Figure 4.4 : CPR is performed using the true ℓ_1 norm value of 1.4662, along with a range of other norm values.




Figure 4.5 : The results of CSPR with varying levels of M Fourier modulus measurements, randomly selected from the entire diffraction pattern.




Figure 4.6 : *The percent error of the Fourier modulus of the recovered result, compared to the full diffraction pattern, decreases as the number of measurements increases.*

in space and has k nonzero values (empirical evidence from Chapter 3 suggests that the measurements scales according to k , and the basis need not be canonical). Because a scanner must stop at every location to make a measurement, taking fewer measurements will result in less scan time. In addition, the scanner does not have to travel as far, either. The path length traveled by a scanner to k^2 random points on an $n \times N$ square lattice is, on average, $.93kN$ [32]. For $k \ll N$, this is a significant improvement over the N^2 distance required for a full raster scan.

Chapter 5

Conclusion

The compressibility of a signal is prior knowledge that can be used effectively to aid in the recovery of signals from just the magnitude of their Fourier transform. Compared to existing prior signal constraints in PR, such as an exact support constraint, a compressibility constraint is more general, but is as powerful. In the context of PR of signals compressible in a wavelets basis, the signal's compressibility is a better constraint for PR than even exact support. Compressibility can also be effectively leveraged to recover a signal from a number of measurements that scales with its structure, rather than its bandwidth. Mathematical theory, empirical support, and the use of our concepts on a THz imaging problem show that compressive phase retrieval can positively impact any scientific application involving a signal's Fourier transform modulus.

Bibliography

- [1] D. Donoho, “Compressed sensing,” *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [2] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Info. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [3] E. Candès and T. Tao, “Near optimal signal recovery from random projections and universal encoding strategies,” *IEEE Trans. Info. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [4] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit,” *SIAM J. Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.
- [5] E. Candès and J. Romberg, “Stable signal recovery from incomplete observations,” in *Proc. SPIE 5914*, 59140S, 2005.
- [6] J. Miao, P. Charalambous, J. Kirz, and D. Sayre, “Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens,” *Nature*, vol. 400, pp. 342–344, 1999.
- [7] J. R. Fienup, “Lensless coherent imaging by phase retrieval with an illumination pattern constraint,” *Optics Express*, vol. 14, no. 2, pp. 498–508, 2006.

- [8] M. Moravec, J. Romberg, and R. Baraniuk, "Compressive phase retrieval," in *Proc. SPIE 6701*, p. 670120, 2007.
- [9] W. Chan, M. Moravec, R. Baraniuk, and D. Mittleman, "Terahertz imaging with compressed sensing and phase retrieval," in *CLEO-2007*, May 2007.
- [10] M. H. Hayes, "The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform," *IEEE Trans. ASSP*, vol. ASSP-30, no. 2, pp. 140–154, 1982.
- [11] J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," *Optics Letters*, vol. 3, no. 1, pp. 27–29, 1978.
- [12] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Applied Optics*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [13] S. Marchesini, "A unified evaluation of iterative projection algorithms for phase retrieval," *Rev. Sci. Instrum*, vol. 78, 2007.
- [14] V. Elser, "Phase retrieval by iterated projections," *J. Opt. Soc. Am. A*, vol. 20, no. 1, pp. 40–55, 2003.
- [15] D. R. Luke, "Relaxed averaged alternating reflections for diffraction imaging," *Inverse Problems*, vol. 21, pp. 37–50, 2005.
- [16] H. He, "Simple constraint for phase retrieval with high efficiency," *J. Opt. Soc. Am. A*, vol. 23, pp. 550–556, 2006.
- [17] J. Miao, D. Sayre, and H. Chapman, "Phase retrieval from the magnitude of the Fourier transforms of nonperiodic objects," *J. Opt. Soc. Am. A*, vol. 15, pp. 1662–1669, 1998.

- [18] J. R. Fienup, "Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint," *J. Opt. Soc. Am. A*, vol. 4, pp. 118–123, 1987.
- [19] J. S. Wu, U. Weierstall, and J. C. H. Spence, "Iterative phase retrieval without support," *Optics Letters*, vol. 29, no. 23, pp. 2737–2739, 2004.
- [20] D. Takhar, J. Laska, M. Wakin, M. Duarte, D. Baron, S. Sarvotham, K. Kelly, and R. Baraniuk, "A new compressive imaging camera architecture using optical-domain compression," in *Proc. of Computational Imaging IV at SPIE Electronic Imaging*, (San Jose), January 2006.
- [21] M. Lustig, J. Santos, J. Lee, D. Donoho, and J. Pauly, "Compressed sensing for rapid MR imaging," in *Proceedings of Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, (Rennes, France), 2005.
- [22] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud, "Theory and implementation of an analog-to-information converter using random demodulation," in *Proc. IEEE Int. Symp. on Circuits and Systems (ISCAS)*, (New Orleans), May 2007.
- [23] M. H. Hayes and J. H. McClellan, "Reducible polynomials in more than one variable," *Proc. IEEE*, vol. 70, no. 2, pp. 197–198, 1982.
- [24] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," Tech. Rep. TR-200801, Department of Computer Science, University of British Columbia, January 2008. Preprint available at www.optimization-online.org.

- [25] W. L. Chan, J. Deibel, and D. M. Mittleman, "Imaging with terahertz radiation," *Rep. Prog. Phys.*, vol. 70, pp. 1325–1379, 2007.
- [26] D. Zimdars, "High speed terahertz reflection imaging," *Proc. SPIE*, vol. 5692, pp. 255–259, 2005.
- [27] S. Wang and X. Zhang, "Pulsed terahertz tomography," *J. Phys. D: Appl. Phys.*, vol. 37, no. 964, pp. R1–36, 2004.
- [28] J. Pearce, H. Choi, and D. M. Mittleman, "Terahertz wide aperture reflection tomography," *Optics Letters*, vol. 30, pp. 1653–1655, 2005.
- [29] A. Bandyopadhyay, A. Stepanov, B. Schulkin, M. D. Federici, A. Sengupta, D. Gary, and J. F. Federici, "Terahertz interferometric and synthetic aperture imaging," *J. Opt. Soc. Am. A*, vol. 23, pp. 1168–1178, 2006.
- [30] M. T. Reiten, S. A. Harmon, and R. A. Cheville, "Terahertz beam propagation measured through three-dimensional amplitude profile determination," *J. Opt. Soc. Am. B*, vol. 20, pp. 2215–2225, 2003.
- [31] A. W. M. Lee, Q. Qin, S. Kumar, B. S. Williams, Q. Hu, and J. L. Reno, "Real-time terahertz imaging over a standoff distance (>25 meters)," *Appl. Phys. Lett.*, vol. 89, p. 141125, 2006.
- [32] A. Chakraborti and B. Chakrabarti, "The travelling salesman problem on randomly diluted lattices: Results for small-size systems," *Eur. Phys. J. B*, vol. 16, pp. 677–680, 2000.