The Phase Problem:
A Mathematical Tour from Norbert Wiener to Random Matrices and Convex Optimization

Thomas Strohmer
Department of Mathematics
University of California, Davis

Norbert Wiener Colloquium
University of Maryland February 22, 2013

Acknowledgements

Research in collaboration with:

Emmanuel Candès (Stanford)
Vlad Voroninski (Berkeley)

This work is sponsored by the Defense Threat Reduction Agency, NSF-DMS, and DARPA.

Hundred years ago ...

In 1912, Max von Laue discovered the diffraction of X-rays by crystals

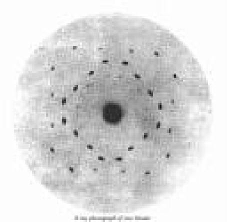

In 1913, W.H. Bragg and his son W.L. Bragg realized one could determine crystal structure from X-ray diffraction patterns

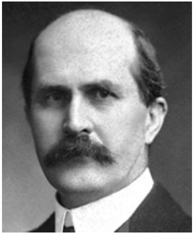

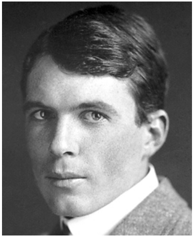

Phase Retrieval Problem

- Signal of interest: $x(t_1, t_2)$
- Fourier transform

$$\hat{x}(\omega_1, \omega_2) = \int x(t_1, t_2) e^{-2\pi i (t_1\omega_1 + t_2\omega_2)}\, dt_1\, dt_2$$

- We measure the intensities of the Fraunhofer diffraction pattern, i.e., the squared modulus of the Fourier transform of the object. The phase information of the Fourier transform is lost.
- Goal: Recover phase of $\hat{x}(\omega_1, \omega_2)$, or equivalently, recover $x(t_1, t_2)$, from $|\hat{x}(\omega_1, \omega_2)|^2$.

Norbert Wiener and Phase Retrieval (1)

- Spectral factorization
- Wiener-Khintchine Theorem (Wiener 1930, Khintchine 1934)
- Wiener-Hopf factorization (1931)

Autocorrelation of a function:

$$\int x(t-s)\overline{x(t)}\,dt \qquad \Longleftrightarrow \qquad |\hat{x}(\omega)|^2$$

($\hat{x}$ denotes the Fourier transform of f)

Norbert Wiener and Phase Retrieval (2)

One particular manifestation:

If x is causal (i.e., $x(t) = 0$, if $t < 0$), and satisfies some regularity conditions, then we can recover x from $|\hat{x}(\omega)|^2$.

Another manifestation:

A singly-infinite positive-definite Toeplitz matrix T has a Cholesky factorization

$$T = C^* C,$$

where C and C^{-1} are upper-triangular matrices.

Patterson function - The workhorse in Phase Retrieval

Patterson: “What do you know about a function, when you know only the amplitudes of its Fourier coefficients?”

Wiener: “You know the Faltung [convolution]”.

Wiener: “The route you are looking for is a corollary of the Wiener-Khintchine Theorem”

The Patterson function is the convolution of the Electron density function with itself

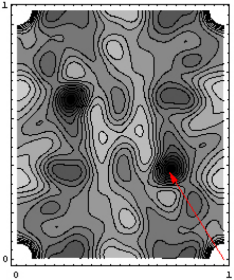

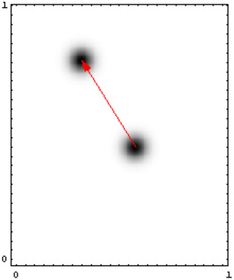

Uncovering the double helix structure of the DNA with X-ray crystallography in 1951.

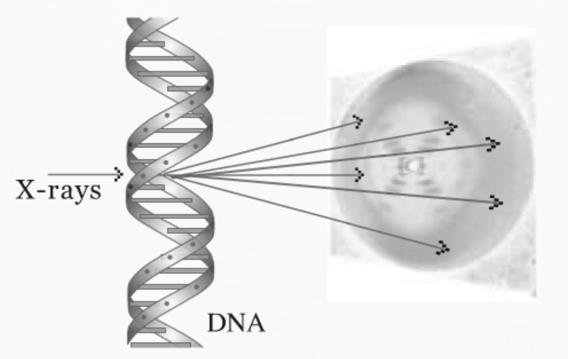

Nobel Prize for Watson, Crick, and Wilkins in 1962 based on work by Rosalind Franklin

Difficult inverse problem:

Determine DNA structure based on diffraction image

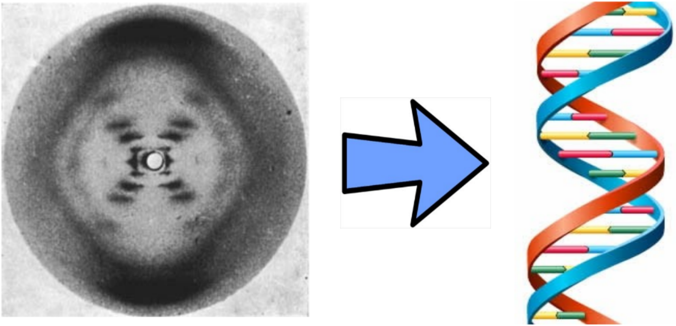

Problem would be easy if we could somehow recover the phase information ("phase retrieval"), because then we could just do an inverse Fourier transform to get DNA structure.

“Shake-and-Bake”

In 1953, Hauptman and Karle developed the Direct method for phase retrieval, based on probabilistic methods and structure invariants and other constraints, expressed as inequalities.

Nobel Prize in 1985.
Method works well for small and sometimes for medium-size molecules (less than a few hundred atoms)

Is phase information really important?

Is phase information really important?

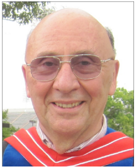

Is phase information really important?

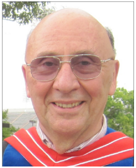

DFT

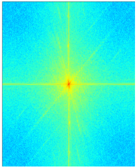

DFT

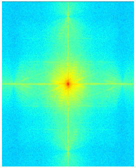

Is phase information really important?

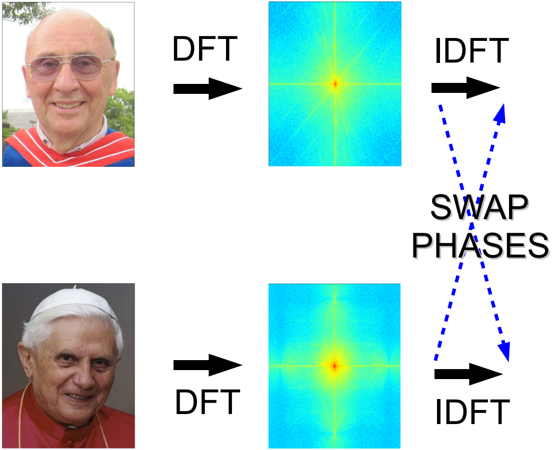

Is phase information really important?

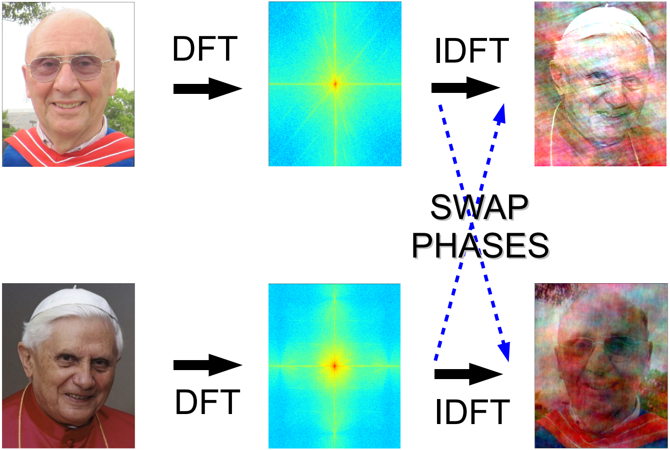

Phase retrieval – why do we care today?

Enormous research activity in recent years due to new imaging capabilities driven by numerous applications.

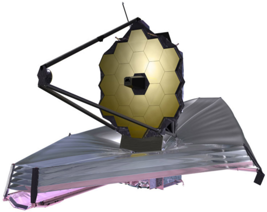

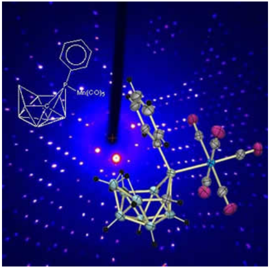

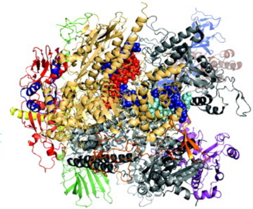

X-ray crystallography

- Method for determining atomic structure within a crystal
- Knowledge of phase crucial to build electron density map
- Initial success of phase retrieval for certain cases by using a combination of mathematics, very specific prior knowledge, and ad hoc “bake-and-shake”-algorithm (1985-Nobel Prize for Hauptman and Karle).
- Very important e.g. in macromolecular crystallography for drug design.

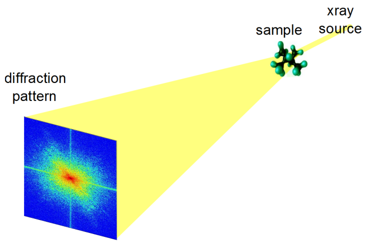

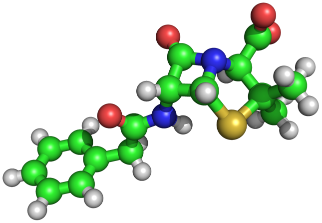

Diffraction microscopy

- X-ray crystallography has been extended to allow imaging of non-crystalline objects by measuring X-ray diffraction patterns followed by phase retrieval.
- Localization of defects and strain field inside nanocrystals
- Quantitative 3D imaging of disordered materials such as nanoparticles and biomaterials
- Potential for imaging single large protein complexes using extremely intense and ultrashort X-ray pulses

Astronomy

Hubble Space Telescope
Wavefront sensing to design and install corrective optics (implemented in 1993); to monitor telescope shrinkage

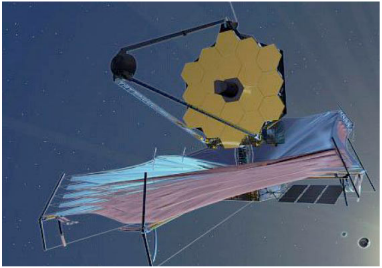

James Webb Space Telescope
Uses deployable segmented optics. Launch in 2018?
Phase retrieval used to align segments of the mirror

An opportunity for mathematics

We spend millions of dollars (and with good reason) on building highly sophisticated instruments and machines that can carry out extremely accurate diffraction experiments

An opportunity for mathematics

We spend millions of dollars (and with good reason) on building highly sophisticated instruments and machines that can carry out extremely accurate diffraction experiments

Yet, we are still stuck with 40-year old fairly simple mathematical algorithms (such as alternating projections by Saxton-Gerchberg) with all their limitations and pitfalls, when attempting to reconstruct images from these high-precision measurements.

Status Quo

Drawbacks of existing phase retrieval methods:

- ad hoc, without any guarantees of recovery of true signal
- need a lot of additional constraints
- unstable in presence of noise
- require user interaction
- do not scale

At the core of phase retrieval lies the problem:

We want to recover a function $\boldsymbol{x}(t)$ from intensity measurements of its Fourier transform, $|\hat{\boldsymbol{x}}(\omega)|^2$.

- Without further information about $\boldsymbol{x}$, the phase retrieval problem is ill-posed. We can either impose additional properties of $\boldsymbol{x}$ or take more measurements (or both)

At the core of phase retrieval lies the problem:

We want to recover a function $\boldsymbol{x}(t)$ from intensity measurements of its Fourier transform, $|\hat{\boldsymbol{x}}(\omega)|^2$.

- Without further information about $\boldsymbol{x}$, the phase retrieval problem is ill-posed. We can either impose additional properties of $\boldsymbol{x}$ or take more measurements (or both)
- We want an efficient phase retrieval algorithm based on a rigorous mathematical framework, for which:
  - (i) we can guarantee exact recovery,
  - (ii) which is stable in the presence of noise.
- Want flexible framework that does not require any prior information about the function (signal, image,...), yet can incorporate additional information if available.

General phase retrieval problem

Suppose we have $\boldsymbol{x}_0 \in \mathbb{C}^n$ or $\mathbb{C}^{n_1 \times n_2}$ about which we have quadratic measurements of the form

$$\mathbb{A}(\boldsymbol{x}_0) = \{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0 \rangle|^2 : k = 1, 2, \ldots, m\}.$$

Phase retrieval:

$$\begin{array}{ll} \text{find} & \boldsymbol{x} \\ \text{obeying} & \mathbb{A}(\boldsymbol{x}) = \mathbb{A}(\boldsymbol{x}_0) := \boldsymbol{b}. \end{array}$$

General phase retrieval problem

Suppose we have $\boldsymbol{x}_0 \in \mathbb{C}^n$ or $\mathbb{C}^{n_1 \times n_2}$ about which we have quadratic measurements of the form

$$\mathbb{A}(\boldsymbol{x}_0) = \{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0 \rangle|^2 : k = 1, 2, \ldots, m\}.$$

Phase retrieval:

$$\begin{array}{ll} \text{find} & \boldsymbol{x} \\ \text{obeying} & \mathbb{A}(\boldsymbol{x}) = \mathbb{A}(\boldsymbol{x}_0) := \boldsymbol{b}. \end{array}$$

Goals:

- Find measurement vectors $\{\boldsymbol{a}_k\}_{k \in \mathcal{I}}$ such that x_0 is uniquely determined by $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0 \rangle|\}_{k \in \mathcal{I}}$.
- Find an algorithm that reconstructs $\boldsymbol{x}_0$ from $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0 \rangle|\}_{k \in \mathcal{I}}$.

When does phase retrieval have a unique solution?

We can only determine $\boldsymbol{x}$ from its intensity measurements $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}\rangle|^2\}$ up to a global phase factor:
If $\boldsymbol{x}(t)$ satisfies $\mathbb{A}(\boldsymbol{x}) = \boldsymbol{b}$, then so does $\boldsymbol{x}(t)e^{2\pi i\varphi t}$ for any $\varphi \in \mathbb{R}$.
Thus uniqueness means uniqueness up to global phase.

When does phase retrieval have a unique solution?

We can only determine $\boldsymbol{x}$ from its intensity measurements $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}\rangle|^2\}$ up to a global phase factor:
If $\boldsymbol{x}(t)$ satisfies $\mathbb{A}(\boldsymbol{x}) = \boldsymbol{b}$, then so does $\boldsymbol{x}(t)e^{2\pi i\varphi t}$ for any $\varphi \in \mathbb{R}$.
Thus uniqueness means uniqueness up to global phase.

Conditions for uniqueness for a general signal $\boldsymbol{x} \in \mathbb{C}^n$:

- $4n - 2$ *generic* measurement vectors are sufficient for uniqueness [Balan-Casazza-Edidin 2007]
- As of Feb. 22, 2013: Bodman gives explicit construction showing $4n - 4$ measurements are sufficient
- About $4n$ measurements are also necessary

Uniqueness does not say anything about existence of feasible algorithm or stability in presence of noise.

Lifting

Following [Balan, Bodman, Casazza, Edidin, 2007], we will interpret quadratic measurements of x as linear measurements of the rank-one matrix $\boldsymbol{X} := \boldsymbol{x}\boldsymbol{x}^*$:

$$|\langle \boldsymbol{a}_k, \boldsymbol{x} \rangle|^2 = \mathrm{Tr}(\boldsymbol{x}^* \boldsymbol{a}_k \boldsymbol{a}_k^* x) = \mathrm{Tr}(\boldsymbol{A}_k \boldsymbol{X})$$

where $\boldsymbol{A}_k$ is the rank-one matrix $\boldsymbol{a}_k \boldsymbol{a}_k^*$. Define linear operator $\mathcal{A}$: $\boldsymbol{X} \to \{\mathrm{Tr}(\boldsymbol{A}_k \boldsymbol{X})\}_{k=1}^{m}$.

Lifting

Following [Balan, Bodman, Casazza, Edidin, 2007], we will interpret quadratic measurements of x as linear measurements of the rank-one matrix $\boldsymbol{X} := \boldsymbol{x}\boldsymbol{x}^*$:

$$|\langle \boldsymbol{a}_k, \boldsymbol{x} \rangle|^2 = \mathrm{Tr}(\boldsymbol{x}^* \boldsymbol{a}_k \boldsymbol{a}_k^* x) = \mathrm{Tr}(\boldsymbol{A}_k \boldsymbol{X})$$

where $\boldsymbol{A}_k$ is the rank-one matrix $\boldsymbol{a}_k \boldsymbol{a}_k^*$. Define linear operator $\mathcal{A}$: $\boldsymbol{X} \to \{\mathrm{Tr}(\boldsymbol{A}_k \boldsymbol{X})\}_{k=1}^m$.
Now, the phase retrieval problem is equivalent to

$$\begin{array}{ll} \text{find} & \boldsymbol{X} \\ \text{subject to} & \mathcal{A}(\boldsymbol{X}) = \boldsymbol{b} \\ & \boldsymbol{X} \succeq 0 \\ & \mathrm{rank}(\boldsymbol{X}) = 1 \end{array} \qquad \text{(RANKMIN)}$$

Having found $\boldsymbol{X}$, we factorize $\boldsymbol{X}$ as $\boldsymbol{x}\boldsymbol{x}^*$ to obtain the phase retrieval solution (up to global phase factor).

Phase retrieval as convex problem?

We need to solve:

$$\begin{array}{ll} \text{minimize} & \operatorname{rank}(\boldsymbol{X}) \\ \text{subject to} & \mathcal{A}(\boldsymbol{X}) = \boldsymbol{b} \\ & \boldsymbol{X} \succeq 0. \end{array} \qquad \text{(RANKMIN)}$$

Note that $\mathcal{A}(\boldsymbol{X}) = \boldsymbol{b}$ is highly underdetermined, thus cannot just invert $\mathcal{A}$ to get $\boldsymbol{X}$.
Rank minimization problems are typically NP-hard.

Phase retrieval as convex problem?

We need to solve:

$$\begin{array}{ll} \text{minimize} & \operatorname{rank}(\boldsymbol{X}) \\ \text{subject to} & \mathcal{A}(\boldsymbol{X}) = \boldsymbol{b} \\ & \boldsymbol{X} \succeq 0. \end{array} \qquad \text{(RANKMIN)}$$

Note that $\mathcal{A}(\boldsymbol{X}) = \boldsymbol{b}$ is highly underdetermined, thus cannot just invert $\mathcal{A}$ to get $\boldsymbol{X}$.
Rank minimization problems are typically NP-hard.
Use trace norm as convex surrogate for the rank functional [Beck '98, Mesbahi '97], giving the semidefinite program:

$$\begin{array}{ll} \text{minimize} & \operatorname{trace}(\boldsymbol{X}) \\ \text{subject to} & \mathcal{A}(\boldsymbol{X}) = \boldsymbol{b} \\ & \boldsymbol{X} \succeq 0. \end{array} \qquad \text{(TRACEMIN)}$$

A new methodology for phase retrieval

Lift up the problem of recovering a vector from quadratic constraints into that of recovering a rank-one matrix from affine constraints, and relax the combinatorial problem into a convenient convex program.

A new methodology for phase retrieval

Lift up the problem of recovering a vector from quadratic constraints into that of recovering a rank-one matrix from affine constraints, and relax the combinatorial problem into a convenient convex program.

PhaseLift

A new methodology for phase retrieval

Lift up the problem of recovering a vector from quadratic constraints into that of recovering a rank-one matrix from affine constraints, and relax the combinatorial problem into a convenient convex program.

PhaseLift

But when (if ever) is the trace minimization problem equivalent to the rank minimization problem?

When is phase retrieval a convex problem?

Theorem: [Candès-Strohmer-Voroninski '11]

Let $\boldsymbol{x}_0$ in $\mathbb{R}^n$ or $\mathbb{C}^n$ and suppose we choose the measurement vectors $\{\boldsymbol{a}_k\}_{k=1}^m$ independently and uniformly at random on the unit sphere of $\mathbb{C}^n$ or $\mathbb{R}^n$. If $m \geq c\, n \log n$, where c is a constant, then PhaseLift recovers $\boldsymbol{x}_0$ exactly from $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0 \rangle|^2\}_{k=1}^m$ with probability at least $1 - 3e^{-\gamma \frac{m}{n}}$, where γ is an absolute constant.

Note that the "oversampling factor" $\log n$ is rather minimal!

First result of its kind: phase retrieval can provably be accomplished via convex optimization with small amount of "oversampling"

When is phase retrieval a convex problem?

Theorem: [Candès-Strohmer-Voroninski '11]

Let $\boldsymbol{x}_0$ in $\mathbb{R}^n$ or $\mathbb{C}^n$ and suppose we choose the measurement vectors $\{\boldsymbol{a}_k\}_{k=1}^m$ independently and uniformly at random on the unit sphere of $\mathbb{C}^n$ or $\mathbb{R}^n$. If $m \geq c\, n \log n$, where c is a constant, then PhaseLift recovers $\boldsymbol{x}_0$ exactly from $\{|\langle \boldsymbol{a}_k, \boldsymbol{x}_0\rangle|^2\}_{k=1}^m$ with probability at least $1 - 3e^{-\gamma\frac{m}{n}}$, where γ is an absolute constant.

Note that the "oversampling factor" $\log n$ is rather minimal!

First result of its kind: phase retrieval can provably be accomplished via convex optimization with small amount of "oversampling"
Recent update: [Candes- Li '12]
Condition $m \geq cn \log n$ can be replaced by $m \geq c_0 n$.

Geometric interpretation

Assume for simplicity that the trace of the solution were known (easy to do in practice), say $\mathrm{Tr}(\boldsymbol{X}) = 1$. In this case our problem reduces to solving the feasibility problem

$$\begin{array}{ll} \text{find} & \boldsymbol{X} \\ \text{such that} & \mathcal{A}(\boldsymbol{X}) = \boldsymbol{b}, \ \boldsymbol{X} \succeq 0 \end{array}$$

(knowledge of $\mathcal{A}$ determines $\mathrm{Tr}(\boldsymbol{X})$)
This is a problem in algebraic geometry since we are trying to find a solution to a set of polynomial equations.

Our main theorem states that $\boldsymbol{x}\boldsymbol{x}^*$ is the unique feasible point. I.e, there is no other positive semidefinite matrix $\boldsymbol{X}$ in the affine space $\mathcal{A}(\boldsymbol{X}) = \boldsymbol{b}$.

Geometric interpretation

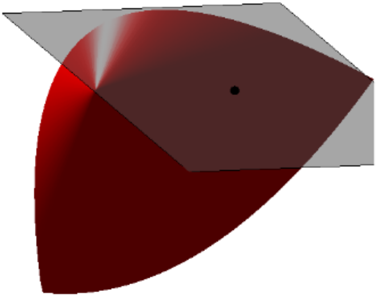

The slice of the (red) positive semidefinite cone $\{X : X \succeq 0\} \cap \{\mathrm{trace}(X) = 1\}$ is quite "pointy" at xx^*. Therefore it is possible for the (gray) affine space $\{\mathcal{A}(X) = b\}$ to be tangent even though it is of dimension about $n^2 - n$.

Sketch of proof

Def: Let $T = T(x)$ be the set of hermitian matrices of the form

$$T = \{\boldsymbol{X} = \boldsymbol{x}_0\boldsymbol{y}^* + \boldsymbol{y}\boldsymbol{x}_0^* : \boldsymbol{y} \in \mathbb{C}^n\}$$

and denote by $T^\perp$ its orthogonal complement.
We use $\boldsymbol{X}_T$ to denote projection of $\boldsymbol{X}$ onto T.
Can further assume that $\boldsymbol{x} = \boldsymbol{e}_1$, since measurement matrix is rotationally invariant.

Sketch of proof

Def: Let $T = T(x)$ be the set of hermitian matrices of the form

$$T = \{\boldsymbol{X} = \boldsymbol{x}_0 \boldsymbol{y}^* + \boldsymbol{y}\boldsymbol{x}_0^* : \boldsymbol{y} \in \mathbb{C}^n\}$$

and denote by $T^\perp$ its orthogonal complement.
We use $\boldsymbol{X}_T$ to denote projection of $\boldsymbol{X}$ onto T.
Can further assume that $\boldsymbol{x} = \boldsymbol{e}_1$, since measurement matrix is rotationally invariant.
Standard duality theory: A sufficient condition for $\boldsymbol{x}\boldsymbol{x}^*$ to be the unique solution to (TRACEMIN) is:

1. the restriction of $\mathcal{A}$ to T is injective:
   $\boldsymbol{X} \in T$ and $\mathcal{A}(X) = 0 \Rightarrow \boldsymbol{X} = 0$.
2. and there exists a dual certificate $\boldsymbol{Y}$ in the range of $\mathcal{A}^*$ obeying $\boldsymbol{Y}_T = \boldsymbol{x}\boldsymbol{x}^*$ and $\boldsymbol{Y}_T^\perp \prec I_T^\perp$.

Showing the existence of such dual certificates is very hard.
Instead we will strengthen the injectivity property, which allows us to relax the conditions on the dual certificate

Key Lemma

Key Lemma: Suppose that the mapping $\mathcal{A}$ obeys the following two properties: for all positive semidefinite matrices X,

$$m^{-1}\|\mathcal{A}(\boldsymbol{X})\|_1 < (1 + 1/9)\|\boldsymbol{X}\|_1; \tag{1}$$

and for all matrices $\boldsymbol{X} \in T$

$$m^{-1}\|\mathcal{A}(\boldsymbol{X})\|_1 > 0.94(1 - 1/9)\|\boldsymbol{X}\|. \tag{2}$$

Suppose further that there exists $\boldsymbol{Y}$ in the range of $\mathcal{A}^*$ obeying

$$\|\boldsymbol{Y}_T - \boldsymbol{x}\boldsymbol{x}^*\|_2 \leq 1/3 \quad \text{and} \quad \|\boldsymbol{Y}_T^{\perp}\| \leq 1/2. \tag{3}$$

Then $\boldsymbol{x}\boldsymbol{x}*$ is the unique minimizer to (TRACEMIN).

Further comments about proof

- Smart choice of approximate dual certificate essential:

$$\boldsymbol{Y} := \frac{1}{m}\mathcal{A}^*\mathcal{A}\mathcal{S}^{-1}(\boldsymbol{x}\boldsymbol{x}^*),$$

  where $\mathcal{S} := \mathbb{E}[\boldsymbol{a}_k\boldsymbol{a}_k^* \otimes \boldsymbol{a}_k\boldsymbol{a}_k^*]$.
  Note that $\boldsymbol{Y} \to \boldsymbol{x}\boldsymbol{x}^*$ as $m \to \infty$.
- Tools to prove conditions rely heavily on non-asymptotic random matrix theory: Concentration of measure for random matrix acting on matrix space, operator-Bernstein inequality, various rather technical moment estimates, ...

Stability in presence of noise

Assume we observe

$$b_i = |\langle \boldsymbol{x}, \boldsymbol{z}_i \rangle|^2 + \nu_i,$$

where ν_i is a noise term with $\|\boldsymbol{\nu}\|_2 \leq \epsilon$.
Consider the solution to

$$\begin{array}{ll} \text{minimize} & \operatorname{trace}(\boldsymbol{X}) \\ \text{subject to} & \|\mathcal{A}(\boldsymbol{X}) - \boldsymbol{b}\|_2 \leq \epsilon \\ & \boldsymbol{X} \succeq 0. \end{array}$$

Theorem: [Candès-S.-Voroninski '11]

Under the same assumptions as in the other theorem, the solution to the noisy, the solution $\hat{\boldsymbol{x}}$ computed via PhaseLift obeys

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2 \leq C_0 \min\left(\|\boldsymbol{x}_0\|_2, \epsilon/\|\boldsymbol{x}_0\|_2\right)$$

Current limitations of our theory

Theorems are not yet completely practical, since most phase retrieval problems involve diffraction, i.e., Fourier transforms, and not unstructured random measurements.

Multiple structured illuminations

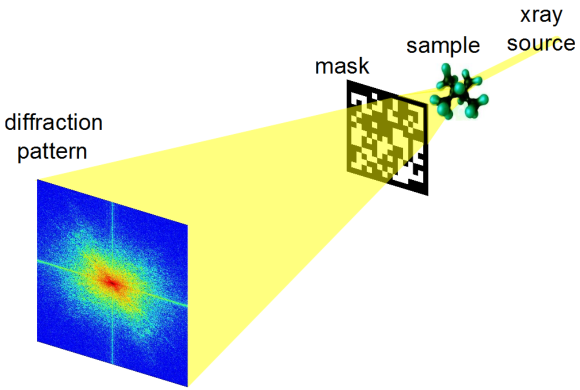

Using different masks (or gratings) generates different illuminations.

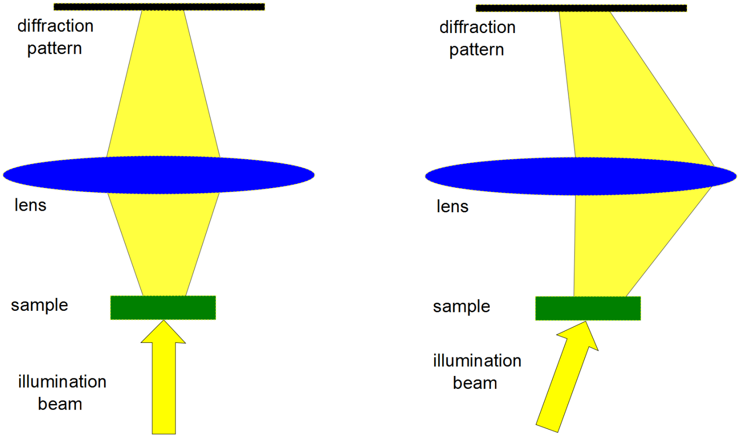

Multiple illuminations using oblique illuminations

Multiple structured illuminations

In all these cases the waveform $\boldsymbol{a}_k$ can be written as

$$\boldsymbol{a}_k(t) \propto w(t) e^{j2\pi\,\langle \omega_k, t\rangle}$$

Say, we use 3 masks, M_1, M_2, M_3, then we measure $|\boldsymbol{F}(\mathrm{diag}(M_1)x_0)|^2, |\boldsymbol{F}(\mathrm{diag}(M_2)x_0)|^2, |\boldsymbol{F}(\mathrm{diag}(M_3)x_0)|^2$, where $\boldsymbol{F}$ is the Fourier transform.

That means we take 3 times as many measurements as with a single diffraction pattern (a single Fourier transform)

Numerical simulations: 1-dim. noisy data

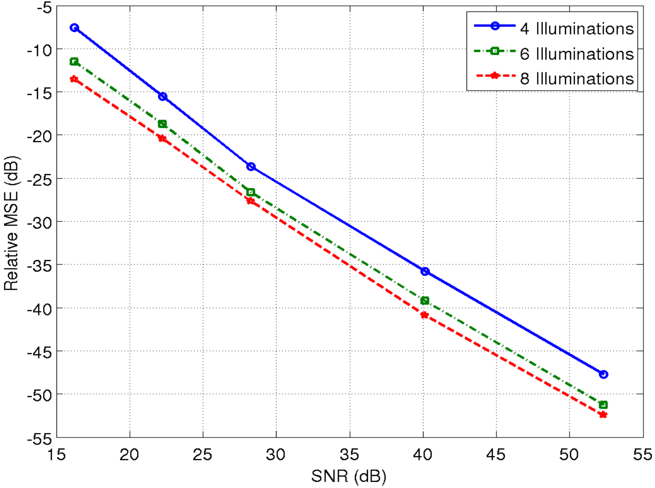

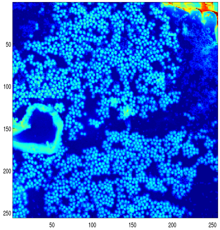

(a) Original image

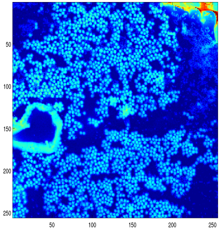

(b) 3 Gaussian masks

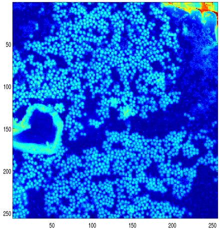

(c) 8 binary masks

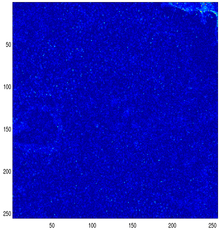

(d) Error btw. (a) and (c)

Thanks to Stefano Marchesini from Berkeley Livermore Labs for data.

TeraHertz Imaging

Potential Applications:

- Chemical mapping of explosives
- Detection of illicit drugs
- Package inspection
- Medical diagnostics

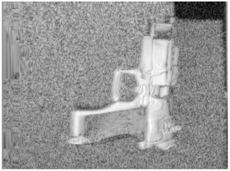

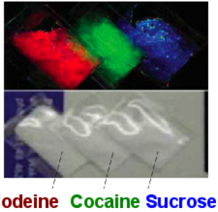

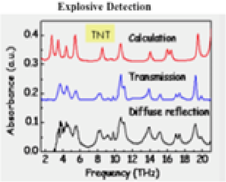

Difficulty: Existing Terahertz imaging systems are too slow for real-time applications due to pixel-by-pixel raster scans

Solution?

- Can take THz measurements in Fourier domain.
- But coherent detectors (i.e., those that can measure phase) are very expensive.
- Non-coherent detectors (i.e., those that only measure intensity) are very cheap.
- PhaseLift TeraHertz camera under construction jointly with Lee Potter (Ohio State).

Conclusion and open problems

PhaseLift: New methodology for phase retrieval

- Can use tools from convex programming
- Works for any signal in any dimension
- Modest amount of "redundant" measurements
- Flexible, noise-aware framework
- Can easily include additional constraints
- Provable results in some regimes
- Other researchers (e.g. S. Mallat) have used PhaseLift successfully, where standard methods failed

- Some deep theoretical questions still open
- Algorithm still slow for large-scale data