PhaseLift: Exact Phase Retrieval via Convex Programming

Thomas Strohmer
Department of Mathematics
University of California, Davis

SIAM Conference on Imaging Science
Albuquerque May 24, 2016

Acknowledgements

Research in collaboration with:

Emmanuel Candès (Stanford)
Yonina Eldar (Technion)
Vladislav 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

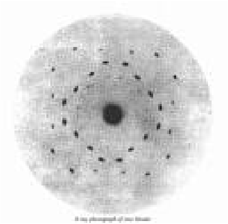

In 1913, W.H. Bragg and his son W.L. Bragg realized one could determine crystal structure from X-ray diffraction patterns

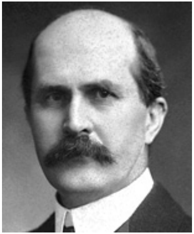

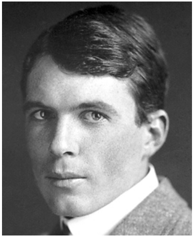

Phase Retrieval Problem

- 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 $|\hat{x}(\omega_1, \omega_2)|^2$.

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

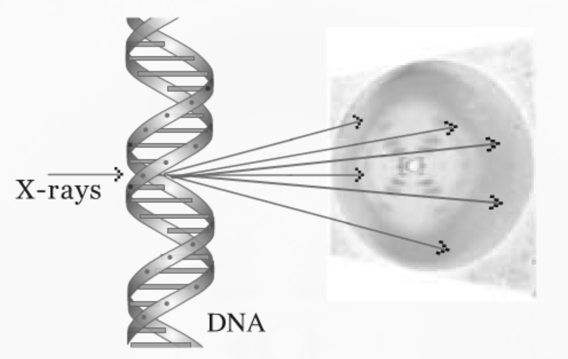

Nobel Prize for Watson, Crick, and Wilkins in 1962 based on work by Rosalind Franklin

Difficult inverse problem:

Determine DNA structure based on diffraction image

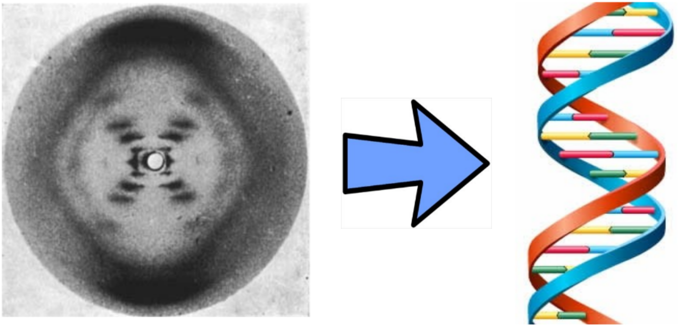

Problem would be easy if we could somehow recover the phase information ("phase retrieval"), because then we could just do an inverse Fourier transform to get DNA structure.

"Shake-and-Bake"

In 1953, Herbert Hauptman, Jerome Karle, and Isabella 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 for H.Hauptman and J.Karle.
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

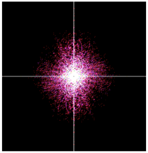

DFT

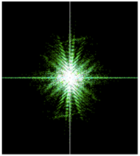

Is phase information really important?

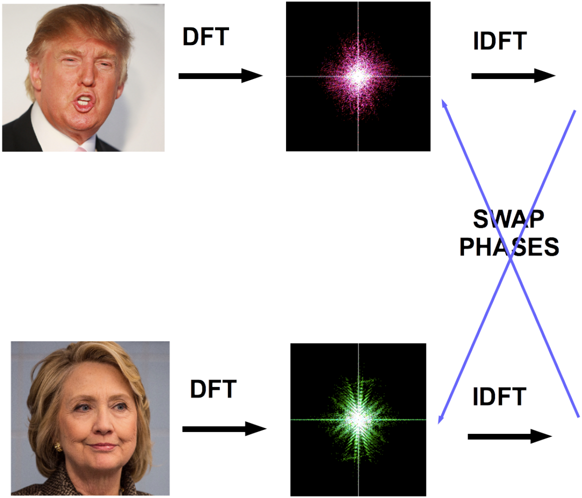

Is phase information really important?

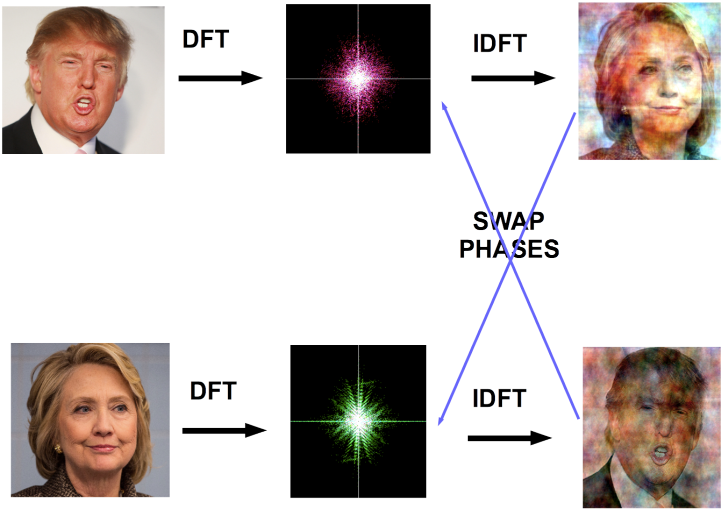

Phase retrieval – why do we care today?

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

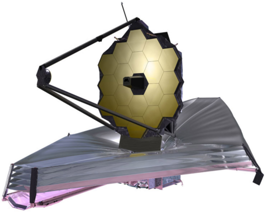

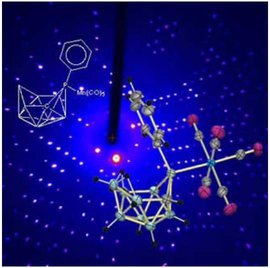

X-ray crystallography

- 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.

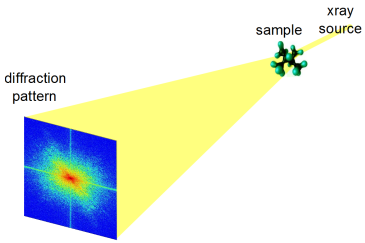

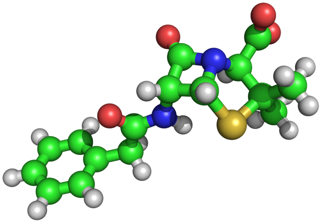

Diffraction microscopy

- 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

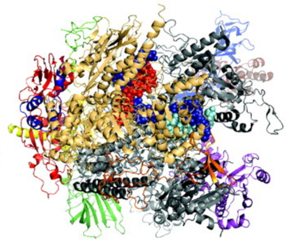

At the core of phase retrieval lies the problem:

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 $\boldsymbol{x}(t)e^{2\pi i\varphi}$ 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 $\boldsymbol{x}(t)e^{2\pi i\varphi}$ 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
- It seems $4n - 4$ measurements are 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 the 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 the linear operator $\mathcal{A}$: $\boldsymbol{X} \rightarrow \{\mathrm{Tr}(\boldsymbol{A}_k \boldsymbol{X})\}_{k=1}^m$.
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"

Proof uses "dual certificate" in semidefinite programming and advanced probability theory (random matrix theory, ...)

Geometric interpretation

Assume for simplicity that the trace of the solution were known (easy to do in practice), say $\text{Tr}(\boldsymbol{X}) = 1$. 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 $\text{Tr}(\boldsymbol{X})$)
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

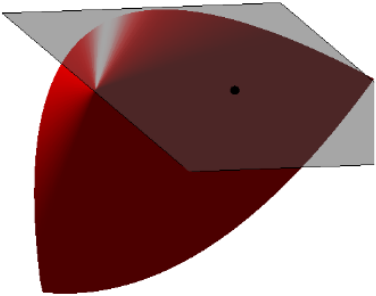

The slice of the (red) positive semidefinite cone $\{X : X \succeq 0\} \cap \{\text{trace}(X) = 1\}$ is quite "pointy" at xx^*. Therefore it is possible for the (gray) affine space $\{\mathcal{A}(X) = b\}$ to be tangent even though it is of dimension about $n^2 - n$.

Multiple structured illuminations

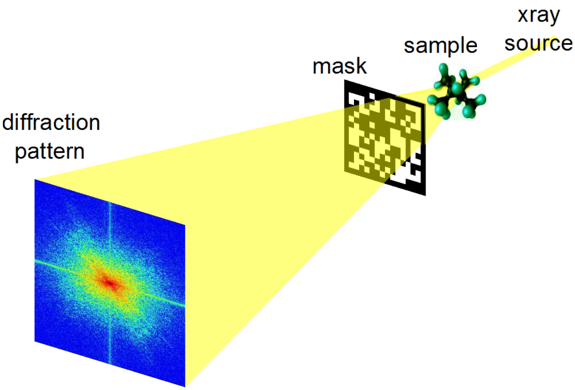

Using random masks to generate multiple illuminations.
[Candes-Eldar-S.-Voroninski, SIAM J. Imaging Sci., 2013].

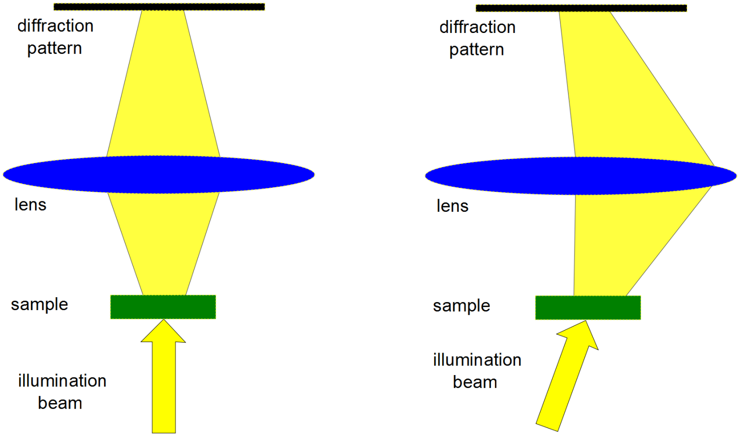

Multiple illuminations using oblique illuminations

Numerical simulations: 1-dim. noisy data

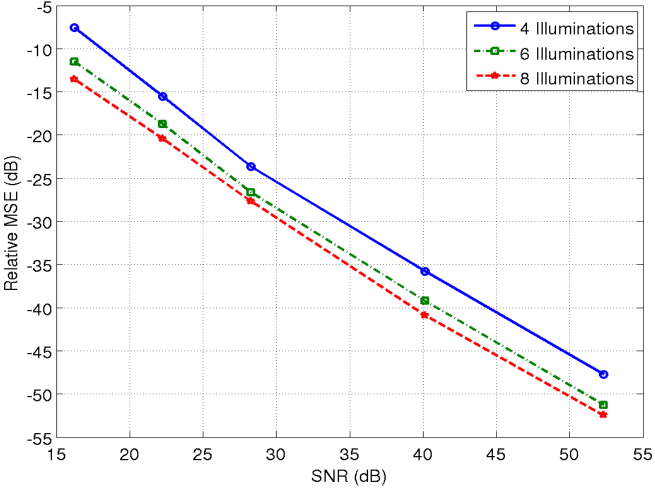

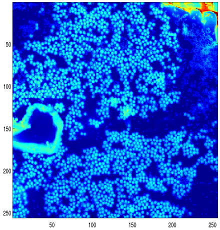

(a) Original image

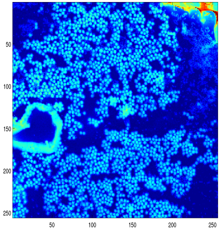

(b) 3 Gaussian masks

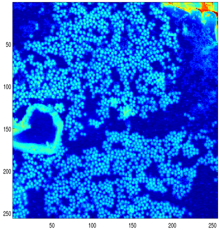

(c) 8 binary masks

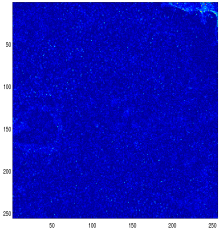

(d) Error btw. (a) and (c)

Thanks to Stefano Marchesini from Berkeley Livermore Labs for data.

Many extensions and improvements of PhaseLift (1)

- Number of measurements $m \geq cn \log n$ can be reduced to $m \geq c_0 n$. (Candes-Li, 2012)
- Theory for random Fourier masks (Candes-Li, 2013)
- Sparse signals (Ohlsson et al. 2012, Li-Voroninsky, 2012, Eldar et al. 2014, Hassibi et al. 2013)
- Expander graphs and PhaseLift (Alexeev-Mixon-Bandeira, 2012)
- PhaseLift and Spherical Designs (Gross-Krahmer-Kueng, 2013)

Many extensions and improvements of PhaseLift (2)

- PhaseLift and STFT measurements, Ptychography (Eldar-Mixon, 2014, Marchesini et al., 2014)
- PhaseLift and Spherical Designs (Gross-Krahmer-Kueng, 2013)
- More general bilinear problems, self-calibration (Schniter 2013, Bresler et al., 2014, Friedlander-Strohmer-Ling, 2014)
- Rigorous nonconvex solvers (Candes, Li, Soltanolkotabi, 2015, Candes-Chen, 2015, Marchesini et al. 2015)
- Blind deconvolution (Ahmed-Recht-Romberg 2014, Bresler et al. 2015, Li-Ling-Strohmer-Wei 2016)