Orbit recovery problems

Tamir Bendory

December 29, 2019

Tel Aviv University
School of Electrical Engineering
https://www.tau.ac.il/~bendory

Orbit recovery problems

We observe i.i.d. realizations of the model:

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where

- x is a fixed but unknown element of a vector space $\mathcal{X}$
- $g_1, \ldots, g_n$ are unknown elements of a (compact) group G
- $P : \mathcal{X} \mapsto \mathcal{Y}$ is a linear operator
- $\mathcal{Y}$ is a finite dimensional measurement space
- $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$

Orbit recovery problems

We observe i.i.d. realizations of the model:

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where

- x is a fixed but unknown element of a vector space $\mathcal{X}$
- $g_1, \ldots, g_n$ are unknown elements of a (compact) group G
- $P : \mathcal{X} \mapsto \mathcal{Y}$ is a linear operator
- $\mathcal{Y}$ is a finite dimensional measurement space
- $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$

Goal: Estimate/recover/reconstruct/learn the orbit

$$\{g \circ x : g \in G\}$$

from $y_1, \ldots, y_n$.

First example: multi-reference alignment

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.

First example: multi-reference alignment

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.

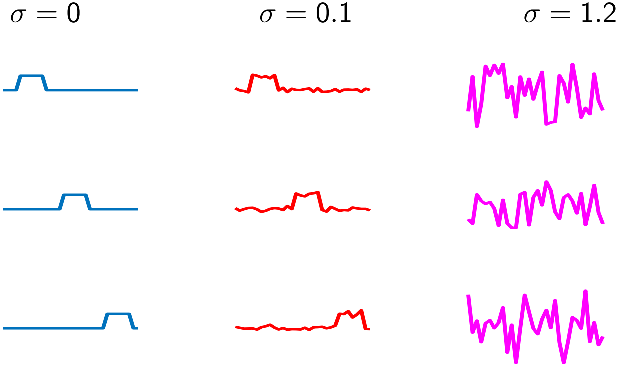

Single particle reconstruction using cryo-electron microscopy

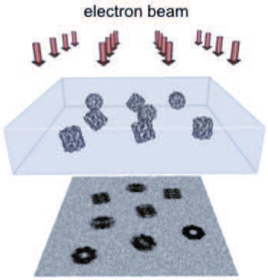

The images of E. coli 50S ribosomal subunit were provided by Dr. Fred Sigworth, Yale Medical School.

Single particle reconstruction using cryo-electron microscopy

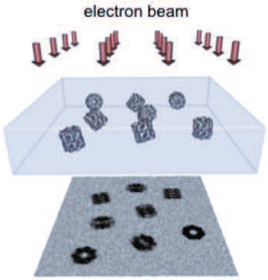

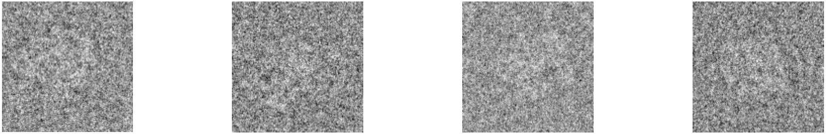

The images of E. coli 50S ribosomal subunit were provided by Dr. Fred Sigworth, Yale Medical School.

The cryo-EM inverse problem

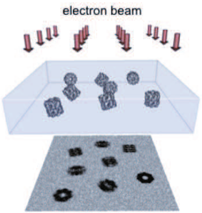

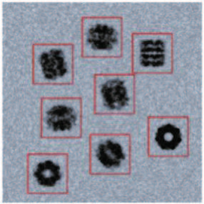

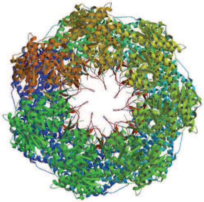

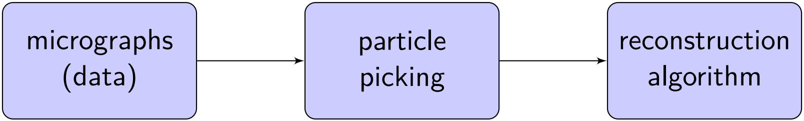

The cryo-EM inverse problem (perfect particle picking)

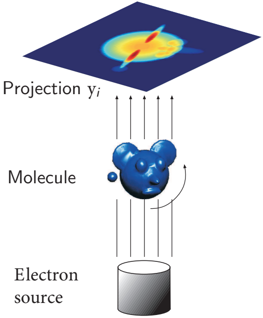

Image formation model

$$y_i = P(g_i \circ x) + \text{noise}, \; g_i \in SO(3)$$

The cryo-EM inverse problem (perfect particle picking)

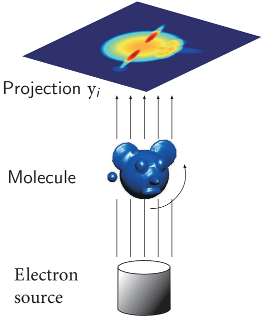

Image formation model

$y_i = P(g_i \circ x) + \text{noise}, \; g_i \in SO(3)$

The cryo–EM problem

Estimate (the orbit of) x from $y_1, \ldots, y_N$

The cryo-EM inverse problem (perfect particle picking)

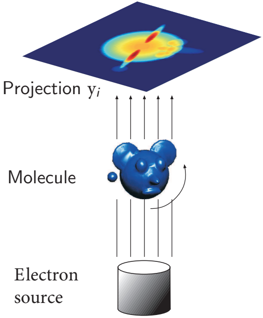

Image formation model

$y_i = P(g_i \circ x) + \text{noise}, \ g_i \in SO(3)$

The cryo–EM problem

Estimate (the orbit of) x from $y_1, \ldots, y_N$

- Can we accurately estimate the rotations?
- Can we accurately estimate the volume x?
- And how?
- What is the optimal estimation rate?

Reconstruction of T20S proteasome

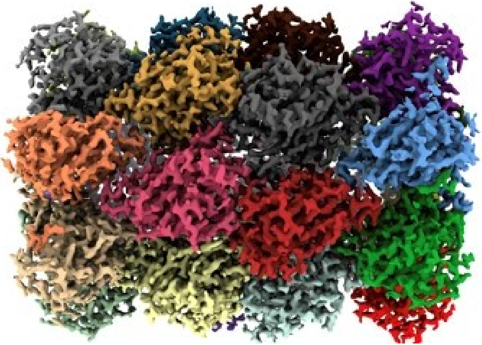

(a)

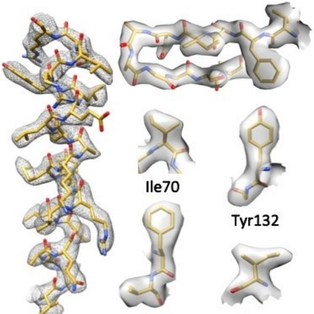

(b)

Taken from: Zhou, Moscovich, Bendory, and Bartesaghi. "Unsupervised particle sorting for high-resolution single-particle cryo-EM." *Inverse Problems*, (2019).

More examples

- Heterogeneous multi-reference alignment [Boumal, Bendory, Lederman, Singer, '18]
- Linear models with permuted data [Pananjady, Wainwright, Courtade, '17]
- Rotations and reflections (the orthogonal group) of a point cloud [Pumir, Singer, Boumal, '19]
- Boolean multi-reference alignment [Abbe, Pereira, Singer, '17]
- Low-rank covariance estimation under unknown translations [Aizenbud, Landa, Shkolnisky, '19; Landa, Shkolnisky, '19]
- Similar algebraic structures [Bendory, Boumal, Leeb, Levin, Singer, '19]

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

High SNR regime: Synchronization! Estimate $g_1, \ldots, g_n$ from $y_1, \ldots, y_n$.

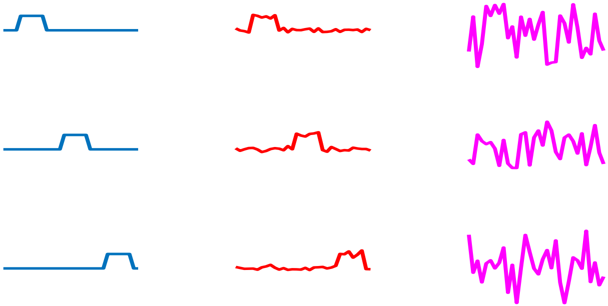

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

High SNR regime: Synchronization! Estimate $g_1, \ldots, g_n$ from $y_1, \ldots, y_n$.

Beautiful theory, near-optimal provable algorithms.

Literature:

- Singer. "Angular synchronization by eigenvectors and semidefinite programming." ACHA, '11.
- Singer, Shkolnisky. "Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming." SIIMS, '11.
- Boumal. "Nonconvex phase synchronization." SIPOT, '16.
- Perry, Wein, Bandeira, Moitra. "Message-Passing Algorithms for Synchronization Problems over Compact Groups." CPAM, '18.
- Zhong, Boumal. "Near-optimal bounds for phase synchronization." SIPOT, '18.
- Chen, Candes. "The projected power method: An efficient algorithm for joint alignment from pairwise differences" CPAM, '18.
- Bandeira, Chen, Singer. "Non-unique games over compact groups and orientation estimation in cryo-EM."

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

Low SNR regime: Accurate estimation of $g_1, \ldots, g_n$ is impossible.

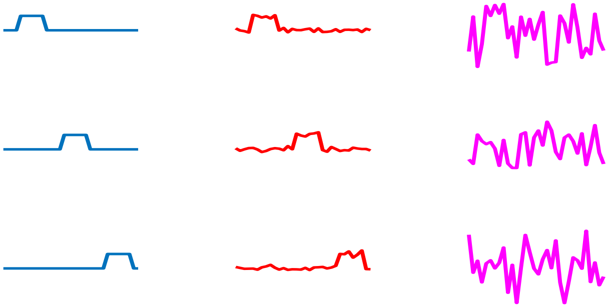

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \dots, n.$$

Low SNR regime: Accurate estimation of $g_1, \dots, g_n$ is impossible.

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

Low SNR regime: Accurate estimation of $g_1, \ldots, g_n$ is impossible.
Solution: Estimate x directly from $y_1, \ldots, y_n$.

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

Low SNR regime: Accurate estimation of $g_1, \ldots, g_n$ is impossible.
Solution: Estimate x directly from $y_1, \ldots, y_n$.
Option 1: Optimize/sample the marginalized likelihood/posterior distribution (EM, MCMC, etc.)

$$L(x; y_1, \ldots, y_n) = p(x) \cdot \prod_i \mathbb{E}_g \mathcal{N}(y_i - P(g \circ x), \sigma^2)$$

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

Low SNR regime: Accurate estimation of $g_1, \ldots, g_n$ is impossible.
Solution: Estimate x directly from $y_1, \ldots, y_n$.
Option 1: Optimize/sample the marginalized likelihood/posterior distribution (EM, MCMC, etc.)

$$L(x; y_1, \ldots, y_n) = p(x) \cdot \prod_i \mathbb{E}_g \mathcal{N}(y_i - P(g \circ x), \sigma^2)$$

Remark: The MLE of

$$L(x, g_1, \ldots, g_n; y_1, \ldots, y_n) = \prod_i \mathcal{N}(y_i - P(g_i \circ x), \sigma^2)$$

might be inconsistent as $n \to \infty$ (Neyman-Scott "paradox").

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \dots, n.$$

Low SNR regime: Accurate estimation of $g_1, \dots, g_n$ is impossible.
Solution: Estimate x directly from $y_1, \dots, y_n$.
Option 2: Method of moments/group invariants

$$\frac{1}{n}\sum_{i=1}^{n} y_i \approx \mathbb{E}_{g,\varepsilon}\{y\} = p_1(x)$$

$$\frac{1}{n}\sum_{i=1}^{n} y_i y_i^T \approx \mathbb{E}_{g,\varepsilon}\{yy^T\} = p_2(x)$$

$$\vdots$$

How to solve the problem?

$$y_i = P(g_i \circ x) + \varepsilon_i, \qquad i = 1, \ldots, n.$$

Low SNR regime: Accurate estimation of $g_1, \ldots, g_n$ is impossible.
Solution: Estimate x directly from $y_1, \ldots, y_n$.
Option 2: Method of moments/group invariants

- One pass over the data
- Consistent
- Amenable to theoretical analysis

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.

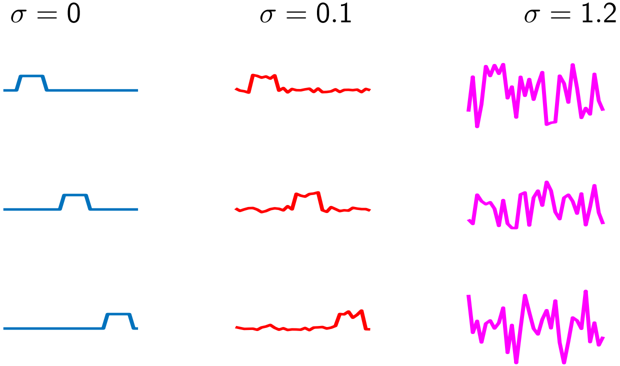

Group invariants

For any $x \in \mathcal{X}$ and $g \in G$, invariants are functions (polynomials) that satisfy

$$h(x) = h(g \circ x)$$

and determine (the orbit of) x uniquely.

Group invariants

For any $x \in \mathcal{X}$ and $g \in G$, invariants are functions (polynomials) that satisfy

$$h(x) = h(g \circ x)$$

and determine (the orbit of) x uniquely.

Example: Let $x \in \mathbb{R}^L$, $G = \mathbb{Z}_2$ with the action $x \mapsto zx$, where $z = \pm 1$.

Group invariants

For any $x \in \mathcal{X}$ and $g \in G$, invariants are functions (polynomials) that satisfy

$$h(x) = h(g \circ x)$$

and determine (the orbit of) x uniquely.

Example: Let $x \in \mathbb{R}^L$, $G = \mathbb{Z}_2$ with the action $x \mapsto zx$, where $z = \pm 1$.

A trivial candidate: $|x|$

Group invariants

For any $x \in \mathcal{X}$ and $g \in G$, invariants are functions (polynomials) that satisfy

$$h(x) = h(g \circ x)$$

and determine (the orbit of) x uniquely.

Example: Let $x \in \mathbb{R}^L$, $G = \mathbb{Z}_2$ with the action $x \mapsto zx$, where $z = \pm 1$.

A trivial candidate: $|x|$

H.W: Is $|x|$ an invariant representation? If not, find an invariant representation h.

Invariants for 1-D MRA

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.
In Fourier space, shifting a signal by α entries is equivalent to

$$\hat{x}[k] \mapsto \hat{x}[k] e^{2\pi i k \alpha / L}.$$

Invariants for 1-D MRA

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.
In Fourier space, shifting a signal by α entries is equivalent to

$$\hat{x}[k] \mapsto \hat{x}[k] e^{2\pi i k \alpha / L}.$$

Thus, it is very easy to construct invariants:

$\hat{x}[0]$	mean
$\hat{x}[k]\hat{x}[-k]$	power spectrum
$\hat{x}[k_1]\hat{x}[k_2]\hat{x}[-k_1 - k_2]$	**bispectrum**
$\hat{x}[k_1]\hat{x}[k_2], \ldots \hat{x}[k_q]\hat{x}[-k_1 - k_2 - \ldots]$	higher-order

Invariants for 1-D MRA

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \dots, n,$$

where $g_i \in G$ is the group of cyclic shifts.
In Fourier space, shifting a signal by α entries is equivalent to

$$\hat{x}[k] \mapsto \hat{x}[k] e^{2\pi i k \alpha / L}.$$

Thus, it is very easy to construct invariants:

$\hat{x}[0]$	mean
$\hat{x}[k]\hat{x}[-k]$	power spectrum
$\hat{x}[k_1]\hat{x}[k_2]\hat{x}[-k_1 - k_2]$	**bispectrum**
$\hat{x}[k_1]\hat{x}[k_2], \dots \hat{x}[k_q]\hat{x}[-k_1 - k_2 - \dots]$	higher-order

Question: What are the invariants of an image under $SO(2)$?

Estimating invariants

Goal: Estimate $x \in \mathbb{R}^L$, up to cyclic shift, from

$$y_i = g_i \circ x + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where $g_i \in G$ is the group of cyclic shifts.
Estimating the invariants:

$$\frac{1}{n}\sum_{i=1}^{n} \hat{y}[0] \to \hat{x}[0] \qquad \text{var}(\sigma^2/n)$$

$$\frac{1}{n}\sum_{i=1}^{n} P_y[k] \to P_x[k] \qquad \text{var}(\sigma^4/n)$$

$$\frac{1}{n}\sum_{i=1}^{n} B_y[k_1, k_2] \to B_x[k_1, k_2] \qquad \text{var}(\sigma^6/n)$$

Many efficient algorithms to recover a signal from its bispectrum [Bendory et al., '17]

Phase retrieval

Phase retrieval is the problem of recovering a signal from its Fourier magnitudes.

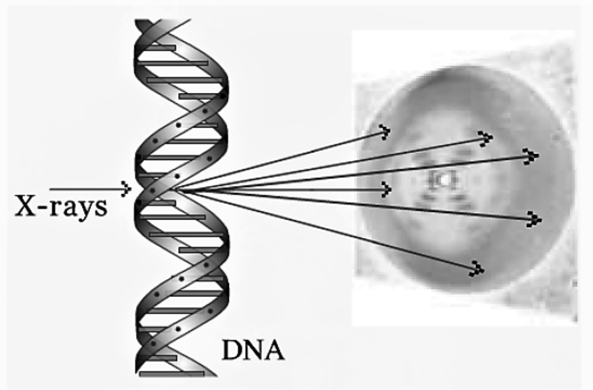

Uncovering the double helix structure of the DNA with X-ray crystallography in 1951. Nobel Prize for Watson, Crick, and Wilkins in 1962 based on work by Rosalind Franklin.

Phase retrieval is also an orbit recovery problem

Let us consider the X-ray crystallography problem: Recover the k-sparse signal $x \in \mathbb{R}^L$ from

$$b = |Fx|,$$

where F is the DFT matrix.

Phase retrieval is also an orbit recovery problem

Let us consider the X-ray crystallography problem: Recover the k-sparse signal $x \in \mathbb{R}^L$ from

$$b = |Fx|,$$

where F is the DFT matrix.

We search for a signal in the intersection of two non-convex sets

$$x \in \mathcal{S} \cap \mathcal{B},$$

where $\mathcal{B} := \{x \in \mathbb{C}^n : b = |Ax|\}$, and $\mathcal{S}$ is the set of all k-sparse signals.

Phase retrieval is also an orbit recovery problem

Let us consider the X-ray crystallography problem: Recover the k-sparse signal $x \in \mathbb{R}^L$ from

$$b = |Fx|,$$

where F is the DFT matrix.

We search for a signal in the intersection of two non-convex sets

$$x \in \mathcal{S} \cap \mathcal{B},$$

where $\mathcal{B} := \{x \in \mathbb{C}^n : b = |Ax|\}$, and $\mathcal{S}$ is the set of all k-sparse signals.

One cannot recover x, but only its orbit. In particular, recovery is possible up to a cyclic shift $\mathbb{Z}_L$, reflection through the origin $\mathbb{Z}_2$ (together, they form the dihedral group D_{2L}), and global sign $\mathbb{Z}_2$ (global phase S^1 in the complex case).

Phase retrieval is also an orbit recovery problem

Let us consider the X-ray crystallography problem: Recover the k-sparse signal $x \in \mathbb{R}^L$ from

$$b = |Fx|,$$

where F is the DFT matrix.

We search for a signal in the intersection of two non-convex sets

$$x \in \mathcal{S} \cap \mathcal{B},$$

where $\mathcal{B} := \{x \in \mathbb{C}^n : b = |Ax|\}$, and $\mathcal{S}$ is the set of all k-sparse signals.

One cannot recover x, but only its orbit. In particular, recovery is possible up to a cyclic shift $\mathbb{Z}_L$, reflection through the origin $\mathbb{Z}_2$ (together, they form the dihedral group D_{2L}), and global sign $\mathbb{Z}_2$ (global phase S^1 in the complex case).

Intricate geometry!

Single particle reconstruction using X-ray free-electron laser (XFEL)

XFEL $\approx$ cryo-EM + phase retrieval

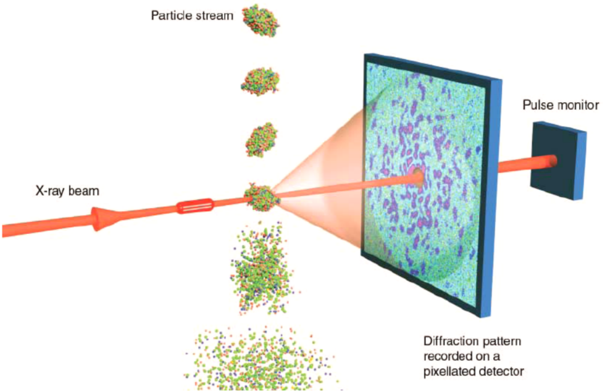

Picture credit: (Gaffney and Chapman, '07)

Thanks for your attention!