Dihedral multi-reference alignment

Dan Edidin[1]

Department of Mathematics, Univeristy of Missouri

Codex Seminar, November 30, 2021

[1]Partially support by Simons collaboration grant 315460

My collaborators

Based on joint work with Tamir Bendory, Will Leeb, and Nir Sharon in the paper arXiv:2107.0526

Physical Motivation: Cryo-electron Microscopy

Cryo-EM is an experimental method in molecular imaging that has come in to wide use in the last 10-15 years.
The 2017 Nobel prize in Chemistry was awarded to Jacques Dubochet, Joachim Frank and Richard Henderson for *for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution.*
In the cryo-EM set up, biological molecules are suspended in a liquid solution and rapidly frozen to extremely low temperatures, approximately 100Kelvin, in a process called vitrification that prevents the formation of ice crystals. A low intensity electron beam passes through the frozen sample and multiple 2D tomographic projections called micrographs are measured.

The method of vitrification in elctron micrcoscopy was first developed about 40 years ago by Jacques Dubochet and his collaborators but until quite recently the resolution of images was relatively low. Since 2013, thanks to advances in hardware and data processing techniques, there has been an explosion in the number of very high resolution molecular structures determined by cryo-EM.

Schematic of cryo-EM image analysis

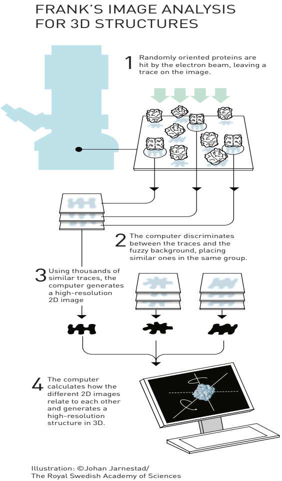

Taken from `https://www.nobelprize.org/prizes/chemistry/2017/press-release/`

The mathematical formulation of the cryo-EM problem

In slightly simplified form, the cryo-EM problem can be described mathematically as follows [7, 5]:

Recover a function $\phi\colon \mathbb{R}^3 \to \mathbb{R}$ (the energy density function of the desired molelcule) from a series of measurements of the form

$$PR_\omega \phi + \text{“noise”}. \tag{1}$$

Here P is a chosen tomographic projection, R_ω is an uknown random rotation in $SO(3)$ and the noise level is high. Because the rotations are unknown, the goal is to recover ϕ up to the action of $SO(3)$; ie recover the $SO(3)$ orbit of the funtion ϕ.

The mutli-reference alignment problem

The *multi-reference alignment problem (MRA)* is a mathematical abstraction of the cryo-EM problem [3]. Precisely it is the problem of recovering the orbit of an unknown signal x in a vector space V from measurements of the form

$$y_i = P(g_i \circ x) + \epsilon_i \tag{2}$$

where P is a linear operator, the g_i are unknown elements of a compact group G and ϵ_i is a Gaussian noise term with standard deviation σ.
In recent years the MRA problem has attracted the attention of many researchers with much work focused on the case where $P = Id$, and $G = \mathbb{Z}_N$ acting on $\mathbb{R}^N$ by cyclic translation.

Synchronization in the low noise regime

If the signal to noise ratio $|x|^2/\sigma^2$ is high then the orbit of x can be recovered using the method of synchronization.

In this method we attempt to synchronize the measurements by determining $g_{ij} = g_i g_j^{-1}$ as the solution to the optimization problem

$$\min_{g \in G} |y_i - g y_j|$$

The (orbit of the) signal can then be approximated by averaging over the synchronized measurements.
However, if the SNR is low the g_{ij} are essentially random since they are more determined by the noise than the translates of the unknown vector x.

Failure of synchronization in the high noise regime from [1]

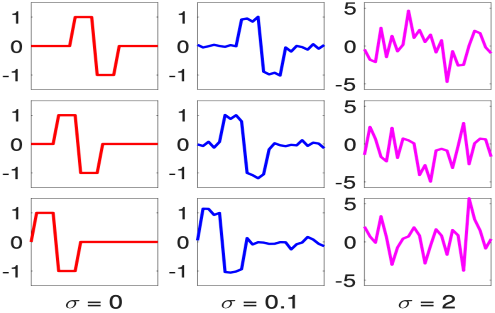

Fig. I.1. The figures illustrates the MRA measurements according to (I.1). The left column presents three measurements with different translations in the absence of noise. In this case, because the solution is defined up to translation, each measurement is a solution. The middle and right columns show measurements with the same translations and low and high noise levels, respectively (note the different scales of the y-axis).

The method of moments

In the high noise regime method of moments can be used instead to recover a signal. The ℓ-th emperical moment of the n measurements $y_i = g_i \circ x + \epsilon_i$ is

$$M^{\ell}_{est} = \frac{1}{n} \sum_{i=1}^{n} y_i^{\otimes \ell} \tag{3}$$

The law of large numbers implies that as $n \to \infty$ the emperical moments converge to the probabilistic moment

$$M^{\ell} = \mathbb{E}(gx^{\otimes \ell}) \tag{4}$$

where the expectation is taken over a probability distribution on the group G.

Sample Complexity

An important theoretical problem is to determine the *sample complexity* of an MRA problem.

The sample complexity of an MRA problem is a bound on the number of measurements as a function of the noise level σ.

For the emperical moments to be an accurate approximation of the probabilistic moments the number of samples needed grows $\sim \sigma^2$. Thus to compute the sample complexity of approximating the ℓ-th moment is $\sim \sigma^{2\ell}$.

Thus if an orbit can be recovered from the first ℓ moments then we can bound the sample complexity of MRA problem by $\sigma^{2\ell}$. Ideally we'd like ℓ to be small, like 2 or 3.

Random vs random

In the MRA problem the elements of the group G are chosen *randomly*. What does *random* mean? One reasonable possibility is that this means that each element of the group occurs equally likely; i.e. the distribution is uniform.

In this case the the entries of the ℓ-th moments are G-invariant polynomial functions on V of degree ℓ. **Hilbert's finiteness theorem** states that there exists ℓ such that the orbit of all $x \in V$ can be recovered from the first ℓ moments.

Bounds for recovering generic vectors

In general computing the degrees of the invariants needed to generate the invariant ring $\mathbb{K}[V]^G$ is hard. However, if G is finite, then the following theorem holds.

Theorem (Baindeira, Blum-Smith, Kileel, Perry, Weed, Wein, [2])

If V contains a copy of the regular representation of G then a generic signal can be recovered from at most 3 moments.

(The regular representation of a group G is the set of functions $G \to \mathbb{K}$ where the group acts by $gf(h) = f(g^{-1}h)$.)

Corollary (Bandeira, Rigolet, Weed [4])

If $\mathbb{Z}_N$ acts on $\mathbb{R}^N$ by cyclic shifts and the probability distribution on $\mathbb{Z}_N$ is uniform, then an orbit can be recovered from the first 3 moments, so the sample complexity of MRA is $\sim \sigma^6$.

Non-uniform distributions

In experimental situations the probability distribution is not necessarily uniform. For example the possible rotations in a cryo-EM are not distributed uniformly in the group $SO(3)$.

If G is finite then a probability distribution on G is simply a function $\rho \colon G \to \mathbb{R}_{\geq 0}$ satisfying $\sum_{g \in G} \rho(g) = 1$.
It follows that the set of probability distributions is a *Zariski dense* subset of the affine linear subspace of the regular representation $\mathrm{reg}(G)$ defined by the equation $\sum_{g \in G} f(g) = 1$.
Thus if we view the distribution as variable then ℓ-th probabalistic moment

$$\sum_{g \in G} \rho(g) g x^{\otimes \ell} \tag{5}$$

is a collection of G-invariant functions of bidegree $(1, \ell)$ on $\mathrm{reg}(G) \times V$.

Generic distributions

We can now ask about how many moments are needed to recover a *generic* vector with a *generic* probability distribution. For the cyclic group we have the following result:

Theorem (Abbe, Bendory, Leeb, Pereira, Sharon, Singer [1])

For a generic probability distribution on the cyclic group $\mathbb{Z}_N$ a generic orbit can recovered from the first two moments. In particular the sample complexity for MRA with generic distribution is $\sim \sigma^4$.

The paper [1] gives a precise condition on the distribution and signal. Precisely the distribution is required to be aperiodic and the DFT has no zero entries.

Moreover, the orbit of a generic signal can be determined using a spectral algorithm. This takes advantage that the action of the cyclic group $\mathbb{Z}_N$ can be diagonalized in the Fourier domain, so the entries of the moment tensors are monomials.

The dihedral group

The dihedral group D_{2N} is a non-abelian group with $2N$ elements which is the symmetry group of the regular N-gon. It is generated by a rotation r of order N and a refelction s of order 2. Rotation and reflection do not commute but satisfy the relation $sr = r^{N-1}s$.
The dihedral group acts on $\mathbb{R}^N$ as follows

$$\begin{aligned} (r \cdot x)[\ell] &= x[(\ell - 1) \bmod N] \\ (s \cdot x)[\ell] &= x[-\ell \bmod N] \end{aligned}$$

The dihedral group naturally occurs as a symmetry group in many physical problems. For example in X-ray crystallography the power spectrum is a function which is invariant under shifts and reflections.
Because the dihedral group is not abelian, the action of D_{2N} cannot be diagonalized, meaning that the terms appearing in the moment tensor cannot be expressed as monomials in any basis.

The main result

Theorem (Bendory, Edidin, Leeb, Sharon [6])

Consider the dihedral MRA problem with a generic probability distribution. The first and second order moments of y are sufficient to recover almost all orbits.

Question. If the distribution is uniform do three moments suffice to recover almost all orbits?
Note that this does not follow from the results of [2] because $\mathbb{R}^N$ is not the regular representation of D_{2N}.

Ideas of proof

As in the proofs of the results in [1] we work in the Fourier domain. Let

$$\hat{M}^1 = \mathbb{E}Fy = FM^1(x, \rho)$$

and

$$\hat{M}^2 = \mathbb{E}Fy(Fy)^* = FM^2(x, \rho)F^*$$

where F is the $N \times N$ DFT matrix. Our main technical result is the following

Theorem

For a generic distribution ρ a generic orbit can be recovered from $\hat{M}^1[0]$ and $\sim 2.5N$ entries of the matrix $\hat{M}^2$

Note that the $\hat{M}^1(\rho, x)[0]$ is just the mean of the entries of x.

Ideas of proof continued

Define

$$M_{i,j} = \hat{p}[i+j]\hat{x}[i]\hat{x}[j] + \hat{q}[N-i-j]\hat{x}[N-i]\hat{x}[N-j], \tag{6}$$

so $M_{i,j}$ is $N\hat{M}^2[i, N-j]$. Our goal is to show that knowledge of $\hat{M}^1[0]$ and $O(N)$ of the entries $M_{i,j}$ determine the orbit of x.
Since ρ is a probability distribution, we note that

$$\hat{p}[0] + \hat{q}[0] = p[0] + \ldots + p[N-1] + q[0] + \ldots + q[N-1] = 1. \tag{7}$$

Thus, $M_{i,-i} = |\hat{x}[i]|^2$. It follows that knowledge of $M^2(x, \rho)$ determines the power spectrum of x.

Idea of proof continued

Replacing x by the vector whose Fourier transform has entries $\hat{x}[i]/|\hat{x}[i]|$, we may assume that each $\hat{x}[i]$ lies on the unit circle. Since $\hat{x}[0]$ is real, we take $\hat{x}[0] = 1$. With this assumption, the formula for $M_{i,j}$ can be written as

$$M_{i,j} = \hat{p}[i+j]\hat{x}[i]\hat{x}[j] + \hat{q}[N-i-j]/(\hat{x}[i]\hat{x}[j]). \tag{8}$$

Showing that the orbit of x is determined by $\hat{M}^1(x,\rho)[0]$ and $\hat{M}^2(x,\rho)$ is equivalent to showing that the system of equations has at most $2N$ solutions:

$$\hat{p}'[i+j]\hat{z}[i]\hat{z}[j] + \hat{q}'[N-i-j]/(\hat{z}[i]\hat{z}[j]) = M_{i,j}, \tag{9}$$

where $\hat{z}[0] = 1$, $|\hat{z}[i]| = 1$ for $i \geq 1$, $\hat{z}$ is the Fourier transform of a vector in $\mathbb{R}^N$, and $\hat{\rho}' = (\hat{p}', \hat{q}')$ is the Fourier transform of a probability distribution on D_{2N}.

Idea of proof continued

Consider the equations

$$\hat{p}'[1]\hat{z}[\ell]\hat{z}[1-\ell] + \hat{q}'[N-1]/(\hat{z}[\ell]\hat{z}[1-\ell]) = M_{\ell,1-\ell}, \quad \ell = 0,\ldots,N-1, \tag{10}$$

where the indices are taken modulo N. For each fixed ℓ, we can view equation (10) as a linear equation in $\hat{p}'[1], \hat{q}'[N-1]$.
For the system to have a solutions it must be consistent. Choosing different values of ℓ, allows to express $\hat{z}[n+1]$ as a rational function of $\hat{z}[1], \hat{z}[2], \hat{z}[n-1], z[n]$. This allows us to show that all $\hat{z}[n]$ for $n \geq 3$ are determined by $\hat{z}[1], \hat{z}[2]$. Finally we show that for each value of $\hat{z}[1]$ there are two possible values of $\hat{z}[2]$ and that $\hat{z}[1]$ satisfies a polynomial equation of degree N. □

References

[1] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *IEEE Transactions on Information Theory*, 65(6):3565–3584, 2018.

[2] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Amelia Perry, Jonathan Weed, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *arXiv preprint arXiv:1712.10163*, 2017.

[3] Afonso S Bandeira, Yutong Chen, Roy R Lederman, and Amit Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *Inverse Problems*, 36(6):064002, 2020.

[4] Afonso S Bandeira, Jonathan Niles-Weed, and Philippe Rigollet. Optimal rates of estimation for multi-reference alignment. *Mathematical Statistics and Learning*, 2(1):25–75, 2020.

[5] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE Signal Processing Magazine*, 37(2):58–76, 2020.

[6] Tamir Bendory, Dan Edidin, William Leeb, and Nir Sharon. Dihedral multi-reference alignment. *arxiv preprint arXiv:2107.05262*.

[7] Amit Singer. Mathematics for cryo-electron microscopy. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures*, pages 3995–4014. World Sci. Publ., Hackensack, NJ, 2018.