Phase Retrieval

http://bicmr.pku.edu.cn/~wenzw/bigdata2020.html

Acknowledgement: this slides is based on Prof. Emmanuel Candès 's lecture notes

Outline

X-ray crystallography

Method for determining atomic structure within a crystal

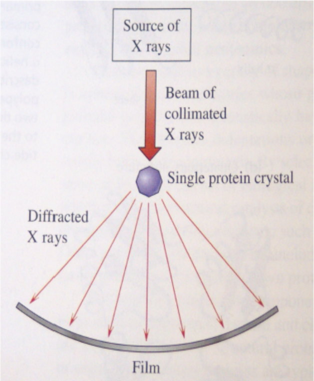

principle

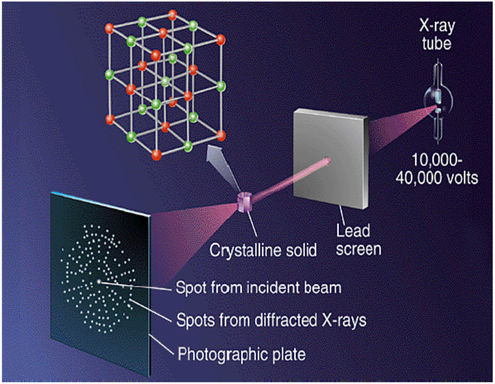

typical setup

10 Nobel Prizes in X-ray crystallography, and counting...

Missing phase problem

Detectors record intensities of diffracted rays $\Longrightarrow$ phaseless data only!

Fraunhofer diffraction $\Longrightarrow$ intensity of electrical $\approx$ Fourier transform

$$|\hat{x}(f_1, f_2)|^2 = \left| \int x(t_1, t_2) e^{-i2\pi(f_1 t_1 + f_2 t_2)} dt_1 dt_2 \right|$$

Electrical field $\hat{x} = |\hat{x}| e^{i\phi}$ with intensity $|\hat{x}|^2$

Phase retrieval problem (inversion)

How can we recover the phase (or signal $x(t_1, t_2)$) from $|\hat{x}(f_1, f_2)|$

Phase and magnitude

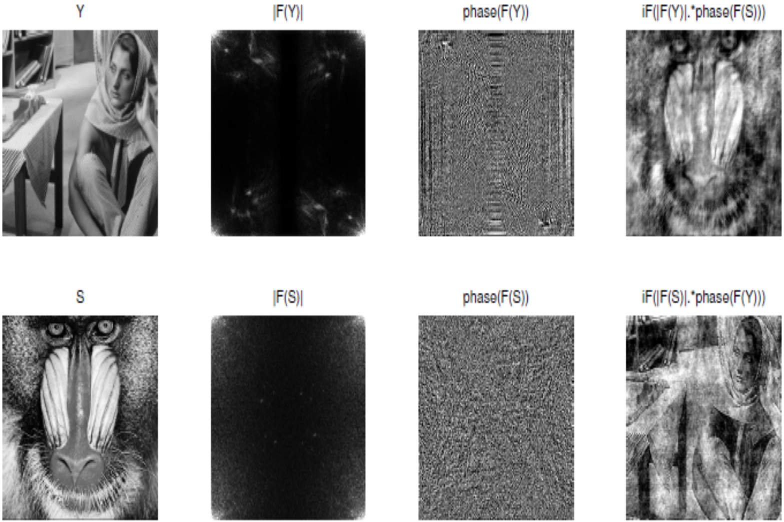

Phase carries more information than magnitude

Outline

Classical Phase Retrieval

Feasibility problem

$$\text{find } x \in S \cap \mathcal{M} \text{ or } \text{ find } x \in S_+ \cap \mathcal{M}$$

- given Fourier magnitudes:

$$\mathcal{M} := \{x(r) \mid |\hat{x}(\omega)| = b(\omega)\}$$

  where $\hat{x}(\omega) = \mathcal{F}(x(r))$, $\mathcal{F}$: Fourier transform
- given support estimate:

$$S := \{x(r) \mid x(r) = 0 \ \text{ for } \ r \notin D\}$$

  or

$$S_+ := \{x(r) \mid x(r) \geq 0 \text{ and } \ x(r) = 0 \text{ if } r \notin D\}$$

Error Reduction

Alternating projection:

$$x^{k+1} = \mathcal{P}_{\mathcal{S}}\mathcal{P}_{\mathcal{M}}(x^k)$$

- projection to S:

$$\mathcal{P}_S(x) = \begin{cases} x(r), & \text{if } r \in D, \\ 0, & \text{otherwise,} \end{cases}$$

- projection to $\mathcal{M}$:

$$\mathcal{P}_{\mathcal{M}}(x) = \mathcal{F}^*(\hat{y}), \text{ where } \hat{y} = \begin{cases} b(\omega)\frac{\hat{x}(\omega)}{|\hat{x}(\omega)|}, & \text{if } \hat{x}(\omega) \neq 0, \\ b(\omega), & \text{otherwise,} \end{cases}$$

Summary of projection algorithms

- Basic input-output (BIO)

$$x^{k+1} = (\mathcal{P}_{\mathcal{S}}\mathcal{P}_{\mathcal{M}} + I - \mathcal{P}_{\mathcal{M}})\,(x^k)$$

- Hybrid input-output (HIO)

$$x^{k+1} = ((1+\beta)\mathcal{P}_{\mathcal{S}}\mathcal{P}_{\mathcal{M}} + I - \mathcal{P}_{\mathcal{S}} - \beta\mathcal{P}_{\mathcal{M}})\,(x^k)$$

- Hybrid projection reflection (HPR)

$$x^{k+1} = \left((1+\beta)\mathcal{P}_{\mathcal{S}_+}\mathcal{P}_{\mathcal{M}} + I - \mathcal{P}_{\mathcal{S}_+} - \beta\mathcal{P}_{\mathcal{M}}\right)(x^k)$$

- Relaxed averaged alternating reflection (RAAR)

$$x^{k+1} = \left(2\beta\mathcal{P}_{\mathcal{S}_+}\mathcal{P}_{\mathcal{M}} + \beta I - \beta\mathcal{P}_{\mathcal{S}_+} + (1-2\beta)\mathcal{P}_{\mathcal{M}}\right)(x^k)$$

- Difference map (DF)

$$x^{k+1} = (I + \beta(\mathcal{P}_{\mathcal{S}}((1-\gamma_2)\mathcal{P}_{\mathcal{M}} - \gamma_2 I) + \mathcal{P}_{\mathcal{M}}((1-\gamma_1)\mathcal{P}_{\mathcal{S}} - \gamma_1 I)))\,(x^k)$$

ADMM

Consider problem

$$\text{find } x \text{ and } y, \text{ such that } x = y, \; x \in \mathcal{X} \text{ and } y \in \mathcal{Y}$$

- $\mathcal{X}$ is either $\mathcal{S}$ or $\mathcal{S}_+$, and $\mathcal{Y}$ is $\mathcal{M}$.
- Augmented Lagrangian function

$$\mathcal{L}(x, y, \lambda) := \lambda^\top (x - y) + \frac{1}{2}\|x - y\|^2$$

- ADMM:

$$\begin{aligned} x^{k+1} &= \arg\min_{x \in \mathcal{X}} \; \mathcal{L}(x, y^k, \lambda^k), \\ y^{k+1} &= \arg\min_{y \in \mathcal{Y}} \mathcal{L}(x^{k+1}, y, \lambda^k), \\ \lambda^{k+1} &= \lambda^k + \beta(x^{k+1} - y^{k+1}), \end{aligned}$$

ADMM

- ADMM

$$
\begin{aligned}
x^{k+1} &= \mathcal{P}_{\mathcal{X}}(y^k - \lambda^k),\\
y^{k+1} &= \mathcal{P}_{\mathcal{Y}}(x^{k+1} + \lambda^k),\\
\lambda^{k+1} &= \lambda^k + \beta(x^{k+1} - y^{k+1}),
\end{aligned}
$$

- ADMM is equivalent to HIO or HPR
  - if $\mathcal{P}_{\mathcal{X}}(x+y) = \mathcal{P}_{\mathcal{X}}(x) + \mathcal{P}_{\mathcal{X}}(y)$

$$x^{k+2} + \lambda^{k+1} = [(1+\beta)\mathcal{P}_{\mathcal{X}}\mathcal{P}_{\mathcal{Y}} + (I - \mathcal{P}_{\mathcal{X}}) - \beta\mathcal{P}_{\mathcal{Y}}](x^{k+1} + \lambda^k)$$

    Hybrid input-output (HIO)

$$x^{k+1} = ((1+\beta)\mathcal{P}_{\mathcal{S}}\mathcal{P}_{\mathcal{M}} + I - \mathcal{P}_{\mathcal{S}} - \beta\mathcal{P}_{\mathcal{M}})\,(x^k)$$

  - if $\beta = 1$

ADMM

- ADMM: updating Lagrange Multiplier twice

$$
\begin{aligned}
x^{k+1} &:= \mathcal{P}_{\mathcal{X}}(y^k - \pi^k), \\
\pi^{k+1} &:= \pi^k + \beta(x^{k+1} - y^k) = -(I - \beta\mathcal{P}_{\mathcal{X}})(y^k - \pi^k), \\
y^{k+1} &:= \mathcal{P}_{\mathcal{Y}}(x^{k+1} + \lambda^k), \\
\lambda^{k+1} &:= \lambda^k + \nu(x^{k+1} - y^{k+1}) = (I - \nu\mathcal{P}_{\mathcal{Y}})(x^{k+1} + \lambda^k),
\end{aligned}
$$

- ADMM is equivalent to ER if $\beta = \nu = 1$

$$x^{k+1} := \mathcal{P}_{\mathcal{X}}(y^k) \text{ and } y^{k+1} := \mathcal{P}_{\mathcal{Y}}(x^{k+1}).$$

- ADMM is equivalent to BIO if $\beta = \nu = 1$

$$x^{k+1} + \lambda^k = (\mathcal{P}_{\mathcal{X}}\mathcal{P}_{\mathcal{Y}} + I - \mathcal{P}_{\mathcal{Y}})\,(x^k + \lambda^{k-1})$$

Numerical comparison

The parameter β in HPR and RAAR was updated dynamically with $\beta_0 = 0.95$. For ADMM, $\beta = 0.5$.

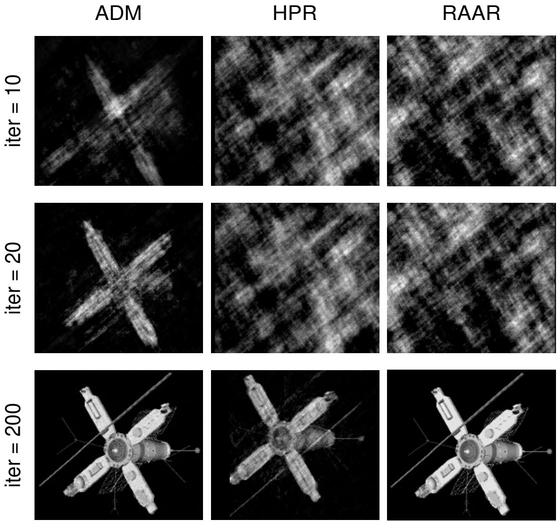

Outline

1. Introduction
2. Classical Phase Retrieval
3. PhaseLift
4. PhaseCut
5. Wirtinger Flows

Discrete mathematical model

- Phaseless measurements about $x_0 \in \mathcal{C}^n$

$$b_k = |\langle a_k, x_0 \rangle|^2, \quad k \in \{1, \ldots, m\}$$

- Phase retrieval is feasibility problem

$$\begin{array}{ll} \text{find} & x \\ \text{s.t.} & |\langle a_k, x_0 \rangle|^2 = b_k, k = 1, \ldots, m \end{array}$$

Solving quadratic equations is NP-complete in general

NP-complete stone problem

Given weights $w_i \in \mathbb{R}$, $i = 1, \ldots, n$, is there an assignment $x_i = \pm 1$ such that

$$\sum_{i=1}^{n} w_i x_i = 0?$$

Formulation as a quadratic system

$$\begin{aligned} |x_i|^2 &= 1, \quad i = 1, \ldots, n \\ \left| \sum_{i=1}^{n} w_i x_i \right|^2 &= 0 \end{aligned}$$

PhaseLift (C., Eldar, Strohmer, Voroninski, 2011)

Lifting: $X = xx^*$

$$b_k = |\langle a_k, x_0\rangle|^2 = a_k^* x x^* a_k = \langle a_k a_k^*, X\rangle$$

Turns quadratic measurements into linear measurements $b = \mathcal{A}(X)$ about xx^*

Phase retrieval problem

$$\begin{array}{ll} \text{find} & X \\ \text{s.t.} & \mathcal{A}(X) = b \\ & X \succeq 0, \operatorname{rank}(X) = 1 \end{array}$$

PhaseLift

$$\begin{array}{ll} \text{find} & X \\ \text{s.t.} & \mathcal{A}(X) = b \\ & X \succeq 0 \end{array}$$

Connections: relaxation of quadratically constrained QP's

- Shor (87) [Lower bounds on nonconvex quadratic optimization problems]
- Goemans and Williamson (95) [MAX-CUT]
- Chai, Moscoso, Papanicolaou (11)

Exact generalized phase retrieval via SDP

Phase retrieval problem

$$\begin{array}{ll} \text{find} & x \\ \text{s.t.} & b_k = |\langle a_k, x_0 \rangle|^2 \end{array}$$

PhaseLift

$$\begin{array}{ll} \text{find} & \operatorname{tr}(X) \\ \text{s.t.} & \mathcal{A}(X) = b, \quad X \succeq 0 \end{array}$$

Theorem (C. and Li ('12); C., Strohmer and Voroninski ('11))

- a_k independently and uniformly sampled on unit sphere
- $m \gtrsim n$

Then with prob. $1 - O(e^{-\gamma m})$, only feasible point is xx^*

$$\{X : \mathcal{A}(X) = b, \text{ and } X \succeq 0\} = \{xx^*\}$$

Extensions to physical setups

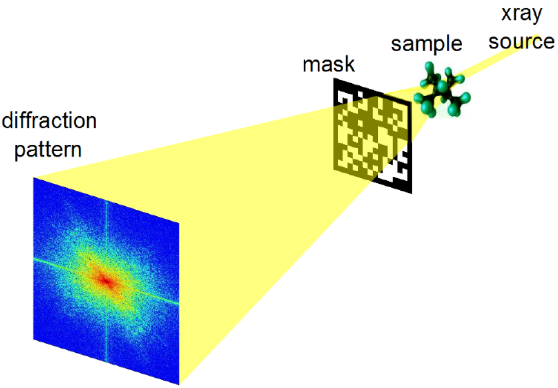

Coded diffraction

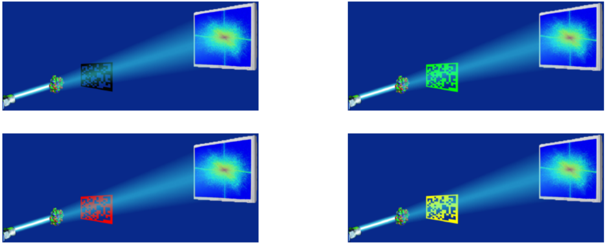

Collect diffraction patterns of modulated samples

$$|\mathcal{F}(w[t]x[t])|^2 \qquad w \in \mathcal{W}$$

Makes problem well-posed (for some choices of $\mathcal{W}$)

Exact recovery

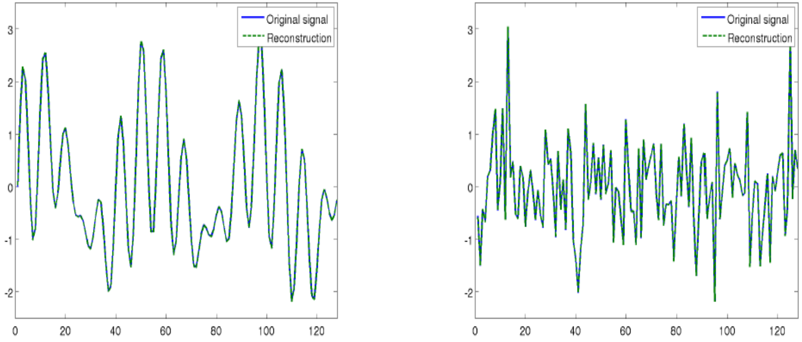

(a) Smooth signal (real part)

(b) Random signal (real part)

Figure: Recovery from 6 random binary masks

Numerical results: noiseless 2D images

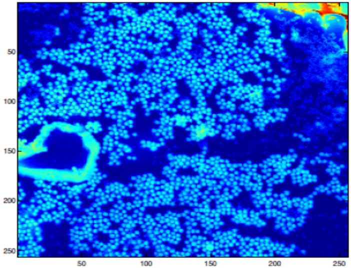

original image

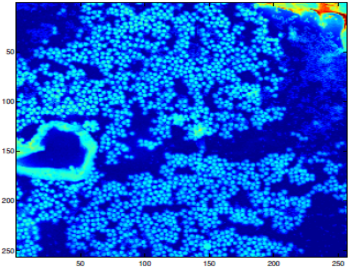

3 Gaussian masks

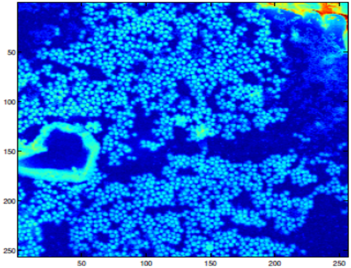

8 binary masks

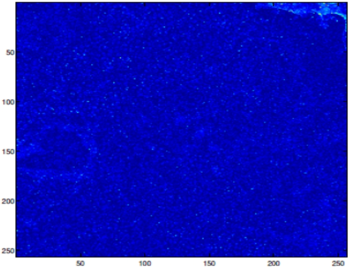

error with 8 binary masks

Outline

PhaseCut

- Given $A \in \mathcal{C}^{m\times n}$ and $b \in \mathbb{R}^m$

$$\text{find } x, \ \text{s.t. } |Ax| = b.$$

(Candes et al. 2011b, Alexandre d'Aspremont 2013)

- An equivalent model

$$\min_{x\in\mathcal{C}^n, y\in\mathbb{R}^m} \quad \frac{1}{2}\|Ax-y\|_2^2$$
$$\text{s.t. } |y| = b.$$

PhaseCut

- Reformulation:

$$\min_{x\in\mathbb{C}^n, u\in\mathbb{C}^m} \frac{1}{2}\|Ax - \operatorname{diag}(b)u\|_2^2$$
$$s.t. \quad |u_i| = 1, , i = 1, \ldots, m.$$

- Given u, the signal variable is $x = A^\dagger \operatorname{diag}(b)u$. Then

$$\min_{u\in\mathbb{C}^m} \quad u^* M u$$
$$s.t. \quad |u_i| = 1, i = 1, \ldots, m,$$

  where $M = \operatorname{diag}(b)(I - AA^\dagger)\operatorname{diag}(b)$ is positive semidefinite.

- The MAXCUT problem

$$\min_{U\in S_m} \quad Tr(UM)$$
$$s.t. \quad U_{ii} = 1,\ i = 1, \cdots, m,\ U \succeq 0.$$

Outline

Phase retrieval by non-convex optimization

Solve the equations: $y_r = |\langle a_r, x\rangle|^2, \quad r = 1, 2, ..., m.$

- **Gaussian model:**

$$a_r \in \mathbb{C}^n \overset{i.i.d.}{\sim} \mathcal{N}(0, I/2) + i\mathcal{N}(0, I/2).$$

- **Coded Diffraction model:**

$$y_r = \left|\sum_{t=0}^{n-1} x[t]\bar{d}_l(t)e^{-i2\pi kt/n}\right|^2, \quad r = (l, k), \quad 0 \le k \le n-1, \quad 1 \le l \le L.$$

Nonlinear least square problem:

$$\min_{z \in \mathbb{C}^n} \quad f(z) = \frac{1}{4m}\sum_{k=1}^{m}(y_k - |\langle a_k, z\rangle|^2)^2$$

- Pro: operates over vectors and not matrices
- Con: f is nonconvex, many local minima

Wirtinger flow: C., Li and Soltanolkotabi ('14)

Strategies:

- Start from a sufficiently accurate initialization
- Make use of Wirtinger derivative

$$
\begin{aligned}
f(z) &= \frac{1}{4m}\sum_{k=1}^{m}(y_k - |\langle a_k, z\rangle|^2)^2 \\
\nabla f(z) &= \frac{1}{m}\sum_{k=1}^{m}(|\langle a_k, z\rangle|^2 - y_k)(a_k a_k^*)z
\end{aligned}
$$

- Careful iterations to avoid local minima

Algorithm: Gaussian model

- **Spectral Initialization:**
  1. Input measurements $\{a_r\}$ and observation $\{y_r\}(r = 1, 2, ..., m)$.
  2. Calculate z_0 to be the leading eigenvector of $Y = \frac{1}{m} \sum_{r=1}^{m} y_r a_r a_r^*$.
  3. Normalize z_0 such that $\|z_0\|^2 = n \frac{\sum_r y_r}{\sum_r \|a_r\|^2}$.
- **Iteration via Wirtinger derivatives:** for $\tau = 0, 1, \ldots$

$$z_{\tau+1} = z_\tau - \frac{\mu_{\tau+1}}{\|z_0\|^2} \nabla f(z_\tau)$$

Convergence property: Gaussian model

distance (up to global phase)

$$\mathbf{dist}(z, x) = \arg \min_{\pi \in [0,2\pi]} \|z - e^{i\phi} x\|$$

Theorem

Convergence for Gaussian model (C. Li and Soltanolkotabi ('14))

- *number of samples* $m \gtrsim n \log n$
- *Step size* $\mu \leq c/n (c > 0)$

Then with probability at least $1 - 10e^{-\gamma n} - 8/n^2 - me^{-1.5n}$*, we have dist*$(z_0, x) \leq \frac{1}{8}\|x\|$ *and after* τ *iteration*

$$\mathbf{dist}(z_\tau, x) \leq \frac{1}{8}(1 - \frac{\mu}{4})^{\tau/2}\|x\|.$$

Here γ *is a positive constant.*

Numerical results: 1D signals

Consider the following two kinds of signals:

- **Random low-pass signals:**

$$x[t] = \sum_{k=-(M/2-1)}^{M/2} (X_k + iY_k) e^{2\pi i(k-1)(t-1)/n},$$

  with M=n/8 and X_k and Y_k are i.i.d. $\mathcal{N}(0,1)$.

- **Random Guassian signals:** where $x \in \mathbb{C}^n$ is a random complex Gaussian vector with i.i.d. entries of the form

$$X[t] = X + iY,$$

  with X and Y distributed as $\mathcal{N}(0,1/2)$.

Success rate

- Set $n = 128$.
- Apply 50 iterations of the power method as initialization.
- Set $\mu_\tau = \min(1 - e^{-\tau/\tau_0}, 0.2)$, where $\tau_0 \approx 330$.
- Stop after 2500 iterations, and declare a trial successful if the relative error of the reconstruction $dist(\hat{x}, x)/\|x\|$ falls below 10^{-5}.
- The empirical probability of success is an average over 100 trials.

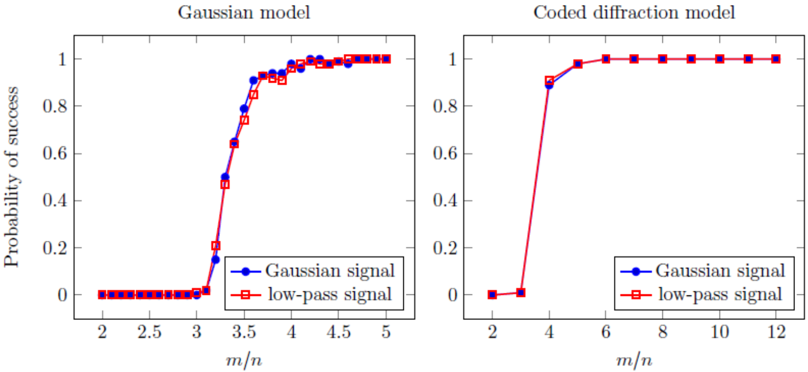

Numerical results: natural images

- View RGB image as $n_1 \times n_2 \times 3$ array, and run the WF algorithm separately on each color band.
- Apply 50 iterations of the power method as initialization.
- Set the step length parameter $\mu_\tau = min(1 - exp(-\tau/\tau_0), 0.4)$, where $\tau_0 \approx 330$. Stop after 300 iterations.
- One FFT unit is the amount of time it takes to perform a single FFT on an image of the same size.

Numerical results: natural images

Figure: Milky way Galaxy. Image size is 1080×1920 pixels; timing is 1318.1 sec or 41900 FFT units. The relative error is 9.3×10^{-16}.

Recall the main theorems

Theorem

Convergence for Gaussian model (C. Li and Soltanolkotabi ('14))

- *number of samples* $m \gtrsim n \log n$
- *Step size* $\mu \leq c/n (c > 0)$

Then with probability at least $1 - 10e^{-\gamma n} - 8/n^2 - me^{-1.5n}$, *we have* *dist*$(z_0, x) \leq \frac{1}{8}\|x\|$ *and after* τ *iteration*

$$\mathbf{dist}(z_\tau, x) \leq \frac{1}{8}(1 - \frac{\mu}{4})^{\tau/2}\|x\|.$$

Here γ *is a positive constant.*

Regularity condition

Definition

Definition We say that the function f satisfies the regularity condition or $RC(\alpha, \beta, \epsilon)$ if for all vectors $z \in E(\epsilon)$ we have

$$Re\left(\langle \nabla f(z), z - xe^{i\phi(z)} \rangle\right) \geq \frac{1}{\alpha} dist^2(z, x) + \frac{1}{\beta} \|\nabla f(z)\|^2.$$

- $\phi(z) := \arg\min_{\phi \in [0, 2\pi]} \|z - e^{i\phi} x\|$.
- $dist(z, x) := \|z - e^{i\phi(z)} x\|$.
- $E(\epsilon) := \{z \in \mathbb{C}^n : dist(z, x) \leq \epsilon\}$.

Proof of convergence

Lemma 1

Assume that f obeys $RC((\alpha, \beta, \epsilon))$ for all $z \in E(\epsilon)$. Furthermore, suppose $z_0 \in E(\epsilon)$, and assume $0 < \mu \leq 2/\beta$. Consider the following update

$$z_{\tau+1} = z_\tau - \mu \nabla f(z_\tau).$$

Then for all τ we have $z_\tau \in E(\epsilon)$ and

$$dist^2(z_\tau, x) \leq \left(1 - \frac{2\mu}{\alpha}\right)^\tau dist^2(z_0, x).$$

Proof of convergence

Proof.

We prove that if $z \in E(\epsilon)$ then for all $0 < \mu \leq 2/\beta$

$$z_+ = z - \mu \nabla f(z)$$

obeys

$$dist^2(z_+, x) \leq \left(1 - \frac{2\mu}{\alpha}\right) dist^2(z, x).$$

Then the lemma holds by inductively applying the equation above.

Proof of convergence

Simple algebraic manipulations together with the regularity condition give

$$
\begin{aligned}
& \left\| z_+ - xe^{i\phi(z)} \right\|^2 \\
= \; & \left\| z - xe^{i\phi(z)} - \mu \nabla f(z) \right\|^2 \\
= \; & \left\| z - xe^{i\phi(z)} \right\|^2 - 2\mu Re \left(\langle \nabla f(z), z - xe^{i\phi(z)} \rangle \right) + \mu^2 \left\| \nabla f(z) \right\|^2 \\
\leq \; & \left\| z - xe^{i\phi(z)} \right\|^2 - 2\mu \left(\frac{1}{\alpha} \left\| z - xe^{i\phi(z)} \right\|^2 + \frac{1}{\beta} \left\| \nabla f(z) \right\|^2 \right) \\
& + \mu^2 \left\| \nabla f(z) \right\|^2 \\
= \; & \left(1 - \frac{2\mu}{\alpha} \right) \left\| z - xe^{i\phi(z)} \right\|^2 + \mu \left(\mu - \frac{2}{\beta} \right) \left\| \nabla f(z) \right\|^2 \\
\leq \; & \left(1 - \frac{2\mu}{\alpha} \right) \left\| z - xe^{i\phi(z)} \right\|^2 ,
\end{aligned}
$$

which concludes the proof.

Proof of regularity condition

We will make use of the following lemma:

Lemma 2

1. x is a solution obeying $\|x\| = 1$, and is independent from the sampling vectors;
2. $m \geq c(\delta) n \log n$ in Gaussian model or $L \geq c(\delta) \log^3 n$ in CD model.

Then,

$$\left\| \nabla^2 f(x) - \mathbb{E} \nabla^2 f(x) \right\| \leq \delta$$

holds with pabability at least $1 - 10e^{-\gamma n} - 8/n^2$ and $1 - (2L+1)/n^3$ for the Gaussian and CD model, respectively.

- The concentration of the Hessian matrix at the optimizers.

Proof of regularity condition

Based on the lemma above with $\delta = 0.01$, we prove the regularity condition by establishing the local curvature condition and the local smoothness condition.

Local curvature condition

We say that the function f satisfies the local curvature condition or $LCC(\alpha, \epsilon, \delta)$ if for all vectors $z \in E(\epsilon)$,

$$Re\left(\langle \nabla f(z), z - xe^{i\phi(z)} \rangle\right) \geq \left(\frac{1}{\alpha} + \frac{1-\delta}{4}\right) dist^2(z, x) + \frac{1}{10m} \sum_{r=1}^{m} \left| a_r^*(z - xe^{i\phi(z)})\right.$$

The LCC condition states that the function curves sufficiently upwards along most directions near the curve of global optimizers.

For the CD model, LCC holds with $\alpha \geq 30$ and $\epsilon = \frac{1}{8\sqrt{n}}$;

For the Gaussian model, LCC holds with $\alpha \geq 8$ and $\epsilon = \frac{1}{8}$.

Proof of regularity condition

Local smoothness condition

We say that the function f satisfies the local smoothness condition or $LSC(\beta, \epsilon, \delta)$ if for all vectors $z \in E(\epsilon)$ we have

$$\|\nabla f(z)\|^2 \leq \beta \left(\frac{(1-\delta)}{4} dist^2(z, x) + \frac{1}{10m} \sum_{r=1}^{m} \left| a_r^*(z - xe^{i\phi(z)}) \right|^4 \right).$$

The LSC condition states that the gradient of the function is well behaved near the curve of global optimizers. Using $\delta = 0.01$, LSC holds with $\beta \geq 550 + 3n$

$$\beta \geq 550 \quad for \quad \epsilon = 1/(8\sqrt{n}),$$
$$\beta \geq 550 + 3n \quad for \quad \epsilon = 1/8.$$

Proof of regularity condition

In conclusion, when $\delta = 0.01$, for the Gaussian model, the regularity condition holds with

$$\alpha \geq 8, \beta \geq 550 + 3n, \ \ and \ \epsilon = 1/8.$$

while for the CD model, the regularity condition holds with

$$\alpha \geq 30, \beta \geq 550, \ \ and \ \epsilon = 1/(8\sqrt{n}),$$

Therefore, for the Gaussian model, linear convergence holds if the initial points satisfies $dist(z_0, x) \leq 1/8$; for the CD model, linear convergence holds if $dist(z_0, x) \leq 1/(8\sqrt{n})$.

Proof of initialization

Recall the initialization algorithm:

1. Input measurements $\{a_r\}$ and observation $\{y_r\}(r = 1, 2, ..., m)$.
2. Calculate z_0 to be the leading eigenvector of $Y = \frac{1}{m}\sum_{r=1}^{m} y_r a_r a_r^*$.
3. Normalize z_0 such that $\|z_0\|^2 = n\frac{\sum_r y_r}{\sum_r \|a_r\|^2}$.

Ideas:

$$\mathbb{E}\left[\frac{1}{m}\sum_{r=1}^{m} y_r a_r a_r^*\right] = I + 2xx^*,$$

and any leading eigenvector of $I + 2xx^*$ is of the form λx. Therefore, by the strong law of large number, the initialization step would recover the direction of x perfectly as long as there are enough samples.

Proof of initialization

In the detailed proof, we will use the following lemma:

Lemma 3

In the setup of Lemma 2,

$$\left\| I - \frac{1}{m} \sum_{r=1}^{m} a_r a_r^* \right\| \leq \delta,$$

holds with probability at least $1 - 2e^{-\gamma m}$ for the Gaussian model and $1 - 1/n^2$ for the CD model. On this event,

$$(1-\delta)\|h\|^2 \leq \frac{1}{m} \sum_{r=1}^{m} |a_r^* h|^2 \leq (1+\delta)\|h\|^2$$

holds for all $h \in \mathbb{C}^n$.

Proof of initialization

Detailed proof:
Lemma 2 gives

$$\|Y - (xx^* + \|x\|^2 I)\| \leq \epsilon := 0.001.$$

Let $\tilde{z}_0$ be the unit eigenvector corresponding to the top eigenvalue λ_0 of Y, then

$$|\lambda_0 - (|\tilde{z}_0 x|^2 + 1)| = |\tilde{z}_0^*(Y - (xx^* + I))\tilde{z}_0| \leq \|Y - (xx^* + I)\| \leq \epsilon.$$

Therefore, $|\tilde{z}_0^* x|^2 \geq \lambda_0 - 1 - \epsilon$. Meanwhile, since λ_0 is the top eigenvalue of Y, and $\|x\| = 1$, we have

$$\lambda_0 \geq x^* Y x = x^*(Y - (I + x^* x))x + 2 \geq 2 - \epsilon.$$

Combining the above two inequalities together, we have

$$|\tilde{z}_0^* x|^2 \geq 1 - 2\epsilon \;\Rightarrow\; dist^2(\tilde{z}_0, x) \leq 2 - 2\sqrt{1 - 2\epsilon} \leq \frac{1}{256} \;\Rightarrow\; dist(\tilde{z}_0, x) \leq \frac{1}{16}.$$

Proof of initialization

Now consider the normalization. Recall that $z_0 = \left(\sqrt{\frac{1}{m}\sum_{r=1}^{m}|a_r^* x|^2}\right)\tilde{z}_0$.

By Lemma 3, with high probability we have

$$\left|\|z_0\| - 1\right| \le \left|\|z_0\|^2 - 1\right| = \left|\frac{1}{m}\sum_{r=1}^{m}|a_r^* x|^2 - 1\right| \le \delta < \frac{1}{16}.$$

Therefore, we have

$$dist(z_0, x) \le \|z_0 - \tilde{z}_0\| + dist(\tilde{z}_0, x) \le \left|\|z_0\| - 1\right| + dist(\tilde{z}_0, x) \le \frac{1}{8}.$$