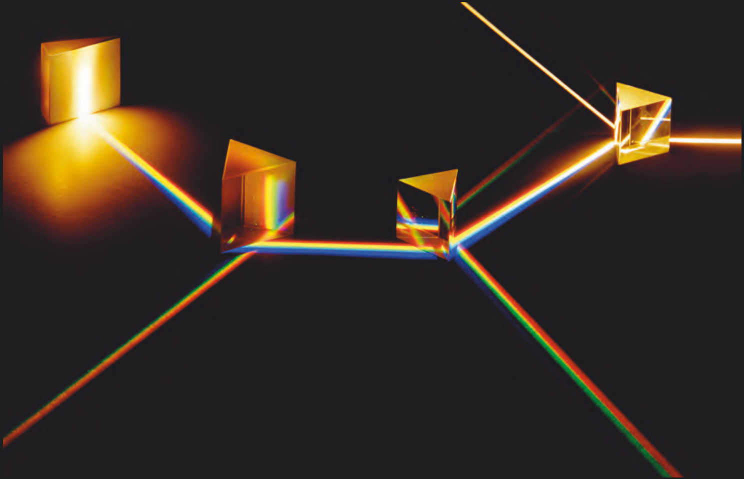

Euclidean Distance Matrices:
A Short Walk Through Theory, Algorithms and Applications

Ivan Dokmanić, Miranda Kreković, Reza Parhizkar, Juri Ranieri and Martin Vetterli

Overview

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

"There are three things that matter in property: location, location, location."

— Lord Harold Samuel, a real estate tycoon in Britain

The 2014 Nobel Prize in Physiology or Medicine

May-Britt Moser, John O'Keefe and Edvard I. Moser
for their discoveries of cells that constitute a positioning system in the brain.

Bat, echoes and navigation

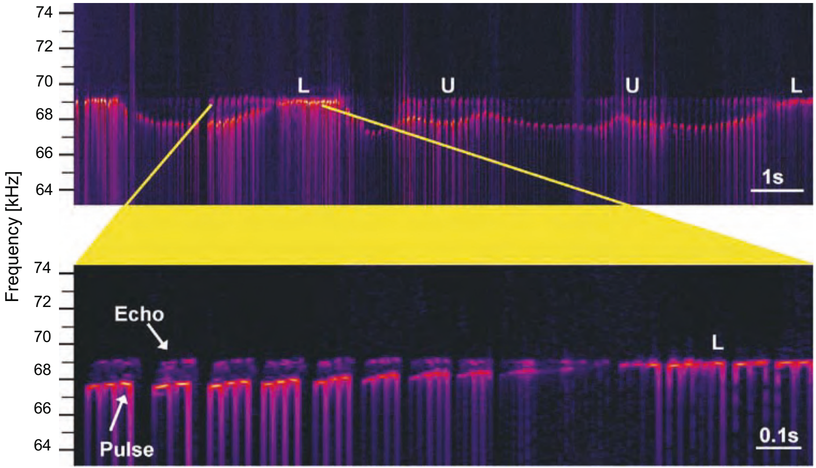

Hiryu, Shizuko, et al. "On-board telemetry of emitted sounds from free-flying bats: compensation for velocity and distance stabilizes echo frequency and amplitude." Journal of Comparative Physiology A 194.9 (2008): 841-851.

Daniel Kish, a human echo-locator

"You took my sonar concept and applied it to every phone in the city. With half the city feeding you sonar, you can image all of Gotham. This is *wrong*."

— Lucius Fox, The Dark Knight

Motivation

- You land at Geneva airport with the Swiss train schedule but no map.

	Lausanne	Geneva	Zürich	Neuchâtel	Bern
Lausanne	0	33	128	40	66
Geneva	33	0	158	64	101
Zürich	128	158	0	88	56
Neuchâtel	40	64	88	0	34
Bern	66	101	56	34	0

Distances in minutes between five Swiss cities

- You land at Geneva airport with the Swiss train schedule but no map.

	Lausanne	Geneva	Zürich	Neuchâtel	Bern
Lausanne	0	33	128	40	66
Geneva	33	0	158	64	101
Zürich	128	158	0	88	56
Neuchâtel	40	64	88	0	34
Bern	66	101	56	34	0

Distances in minutes between five Swiss cities

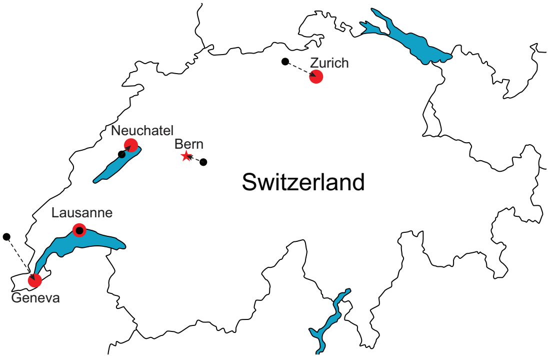

Locations (red) and estimated locations (black) of the cities.

Sensor network localization problem

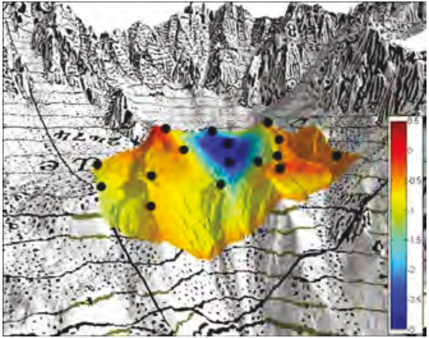

Sensor network deployed on a rock glacier
(SensorScope project)

- We measure the distances between the sensor nodes;
- The distances are noisy and some are missing;
- How do we reconstruct the locations of the sensors?

Molecular conformation problem

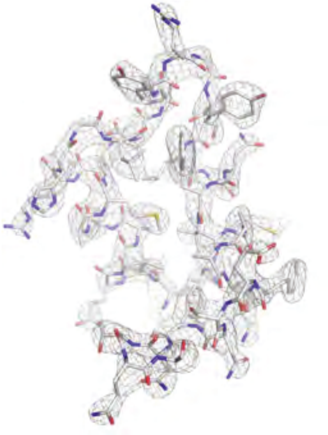

Density map and structure of a molecule [10.7554/elife.01345]

- We measure the distances between the atoms;
- The distances are given as intervals, and some are missing;
- How do we reconstruct the molecule?

Echo Simultaneous Localization and Mapping (Echo SLAM)

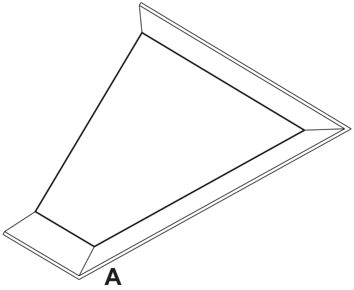

- We consider an unknown room (A)
- A robot moves inside the room autonomously (B)
- At every step, the robot measures the room impulse response (C)
- How do we simultaneously recover the room shape and the robot location?

Echo Simultaneous Localization and Mapping (Echo SLAM)

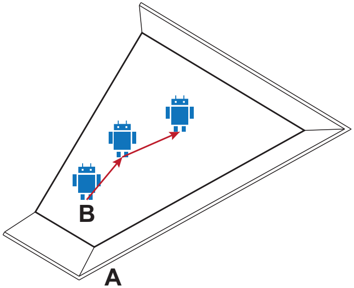

- We consider an unknown room (A)
- A robot moves inside the room autonomously (B)
- At every step, the robot measures the room impulse response (C)
- How do we simultaneously recover the room shape and the robot location?

Echo Simultaneous Localization and Mapping (Echo SLAM)

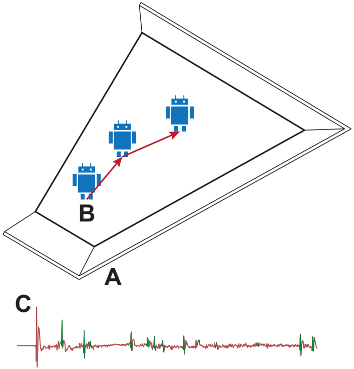

- We consider an unknown room (A)
- A robot moves inside the room autonomously (B)
- At every step, the robot measures the room impulse response (C)
- How do we simultaneously recover the room shape and the robot location?

Echo Simultaneous Localization and Mapping (Echo SLAM)

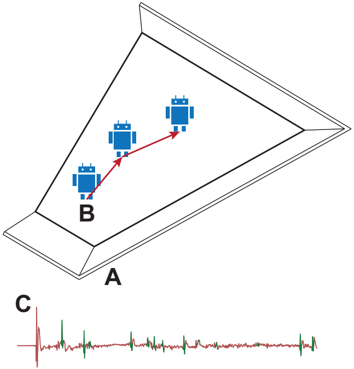

- We consider an unknown room (A)
- A robot moves inside the room autonomously (B)
- At every step, the robot measures the room impulse response (C)
- How do we simultaneously recover the room shape and the robot location?

Overview

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

Euclidean Distance Matrix

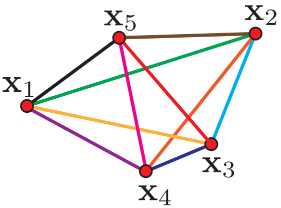

$$\mathbf{X} = \begin{bmatrix} | & & | \\ \mathbf{x}_1 & \cdots & \mathbf{x}_5 \\ | & & | \end{bmatrix} \qquad \mathrm{edm}(\mathbf{X}) =$$

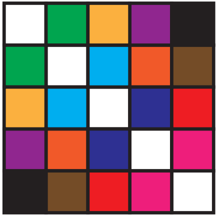

- Consider a set of n points in d dimensions: $\mathbf{X} \in \mathbb{R}^{d \times n}$,
- edm($\mathbf{X}$) contains the squared distances between the points, $\mathrm{edm}(\mathbf{X})_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|^2$,
- Equivalently, edm($\mathbf{X}$) is a linear function of the Gramian $\mathbf{G} = \mathbf{X}^\top \mathbf{X}$.
- Rank of the Gramian $\mathbf{G}$ is d.

Rank testing of Euclidean distance matrix (EDM)

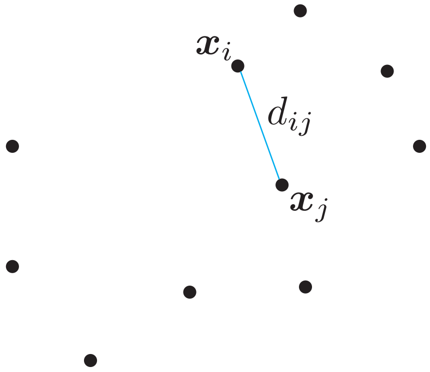

- What can we say about $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ when $\mathbf{x} \in \mathbb{R}^d$?

Rank testing of Euclidean distance matrix (EDM)

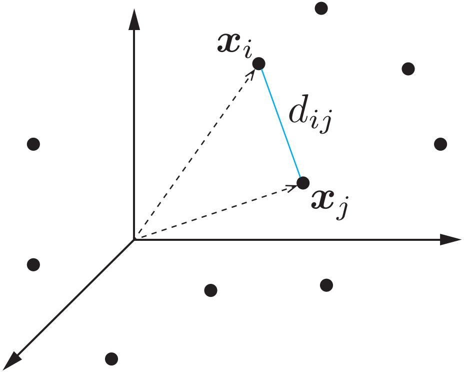

$$\|\mathbf{x}_i - \mathbf{x}_j\|^2 = \langle \mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_j \rangle$$
$$= \langle \mathbf{x}_i, \mathbf{x}_i \rangle + \langle \mathbf{x}_j, \mathbf{x}_j \rangle - 2 \langle \mathbf{x}_i, \mathbf{x}_j \rangle$$

$$g_{ij} \stackrel{\text{def}}{=} \langle \mathbf{x}_i, \mathbf{x}_j \rangle, \quad \mathbf{G} \stackrel{\text{def}}{=} (g_{ij})$$

$$\text{edm}(\mathbf{X}) = -2\mathbf{G} + \mathbf{1}\,\text{diag}(\mathbf{G})^T + \text{diag}(\mathbf{G})\mathbf{1}^T$$

$$\Rightarrow \text{rank}(\text{edm}(\mathbf{X})) \leq \text{rank}(\mathbf{G}) + 1 + 1$$
$$= d + 2$$

- What can we say about $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ when $\mathbf{x} \in \mathbb{R}^d$?

Rank testing of Euclidean distance matrix (EDM)

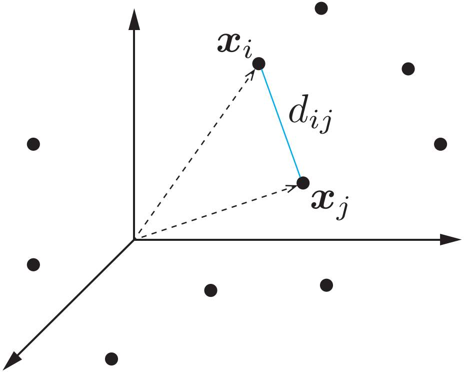

$$\|\mathbf{x}_i - \mathbf{x}_j\|^2 = \langle \mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_j \rangle$$

$$= \underbrace{\langle \mathbf{x}_i, \mathbf{x}_i \rangle}_{g_{ii}} + \underbrace{\langle \mathbf{x}_j, \mathbf{x}_j \rangle}_{g_{jj}} - 2 \underbrace{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}_{g_{ij}}$$

$$g_{ij} \stackrel{\text{def}}{=} \langle \mathbf{x}_i, \mathbf{x}_j \rangle , \quad \mathbf{G} \stackrel{\text{def}}{=} (g_{ij})$$

$$\mathrm{edm}(\mathbf{X}) = -2\mathbf{G} + \mathbf{1}\,\mathrm{diag}(\mathbf{G})^T + \mathrm{diag}(\mathbf{G})\mathbf{1}^T$$

$$\Rightarrow \mathrm{rank}(\mathrm{edm}(\mathbf{X})) \le \mathrm{rank}(\mathbf{G}) + 1 + 1$$

$$= d + 2$$

- What can we say about $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ when $\mathbf{x} \in \mathbb{R}^d$?

Rank testing of Euclidean distance matrix (EDM)

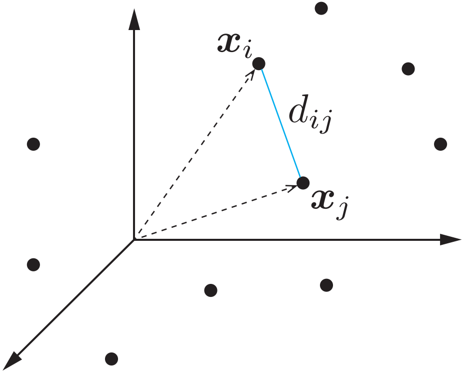

$$\|\mathbf{x}_i - \mathbf{x}_j\|^2 = \langle \mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_j \rangle$$

$$= \underbrace{\langle \mathbf{x}_i, \mathbf{x}_i \rangle}_{g_{ii}} + \underbrace{\langle \mathbf{x}_j, \mathbf{x}_j \rangle}_{g_{jj}} - 2 \underbrace{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}_{g_{ij}}$$

$$g_{ij} \overset{\text{def}}{=} \langle \mathbf{x}_i, \mathbf{x}_j \rangle, \quad \mathbf{G} \overset{\text{def}}{=} (g_{ij})$$

$$\text{edm}(\mathbf{X}) = -2\mathbf{G} + \mathbf{1}\,\text{diag}(\mathbf{G})^T + \text{diag}(\mathbf{G})\mathbf{1}^T$$

$$\Rightarrow \text{rank}(\text{edm}(\mathbf{X})) \leq \text{rank}(\mathbf{G}) + 1 + 1$$

$$= d + 2$$

- What can we say about $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ when $\mathbf{x} \in \mathbb{R}^d$?

Rank testing of Euclidean distance matrix (EDM)

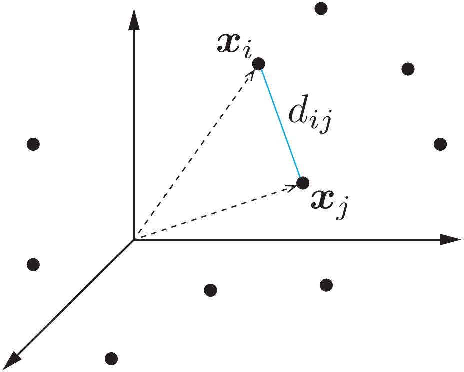

- What can we say about $\|\mathbf{x}_i - \mathbf{x}_j\|^2$ when $\mathbf{x} \in \mathbb{R}^d$?

$$
\begin{aligned}
\|\mathbf{x}_i - \mathbf{x}_j\|^2 &= \langle \mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_j \rangle \\
&= \underbrace{\langle \mathbf{x}_i, \mathbf{x}_i \rangle}_{g_{ii}} + \underbrace{\langle \mathbf{x}_j, \mathbf{x}_j \rangle}_{g_{jj}} - 2 \underbrace{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}_{g_{ij}}
\end{aligned}
$$

$$g_{ij} \stackrel{\text{def}}{=} \langle \mathbf{x}_i, \mathbf{x}_j \rangle , \quad \mathbf{G} \stackrel{\text{def}}{=} (g_{ij})$$

$$\mathrm{edm}(\mathbf{X}) = -2\mathbf{G} + \mathbf{1}\,\mathrm{diag}(\mathbf{G})^T + \mathrm{diag}(\mathbf{G})\mathbf{1}^T$$

$$
\begin{aligned}
\Rightarrow \mathrm{rank}(\mathrm{edm}(\mathbf{X})) &\leq \mathrm{rank}(\mathbf{G}) + 1 + 1 \\
&= d + 2
\end{aligned}
$$

EDM properties: rank

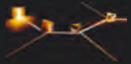

- The rank of an EDM depends only on the embedding dimension of the points:

Theorem (Rank of EDMs)

Rank of an EDM corresponding to points in $\mathbf{R}^d$ is at most $d + 2$.

- Note that the rank depends on the affine dimension of the points.

- EDMs are invariant to rigid transformations (translation, rotations, reflections):

$$\mathrm{edm}(\mathbf{X}) = \mathrm{edm}(\mathbf{R}\mathbf{X} + b\mathbf{1}^\top).$$

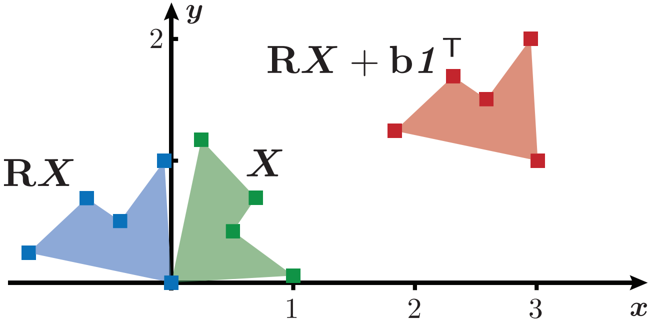

Overview

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

Forward and inverse problem related to EDMs

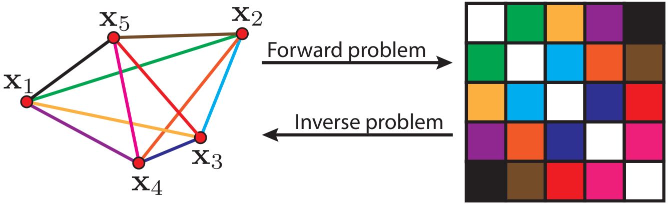

- Forward problem: given the points $\mathbf{X}$, find the EDM $\text{edm}(\mathbf{X})$.
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*;
  - Significantly harder than the forward problem!

Forward and inverse problem related to EDMs

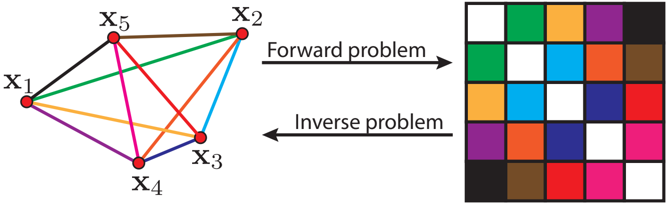

- Forward problem: given the points $\mathbf{X}$, find the EDM $\text{edm}(\mathbf{X})$.
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*;
  - Significantly harder than the forward problem!

Forward and inverse problem related to EDMs

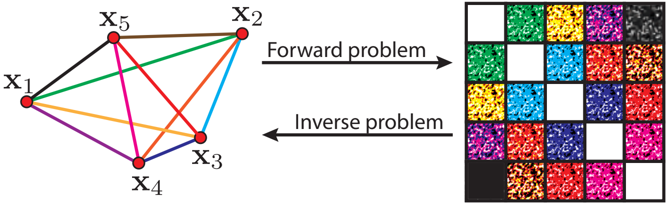

- Forward problem: given the points $\mathbf{X}$, find the EDM edm($\mathbf{X}$).
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*;
  - Significantly harder than the forward problem!

Forward and inverse problem related to EDMs

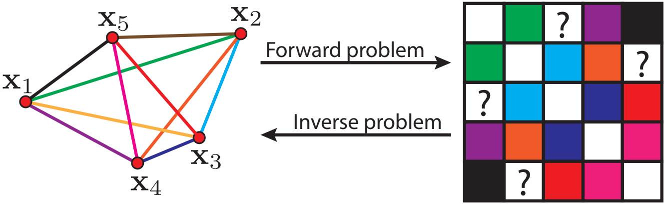

- Forward problem: given the points $\mathbf{X}$, find the EDM $\text{edm}(\mathbf{X})$.
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*;
  - Significantly harder than the forward problem!

Forward and inverse problem related to EDMs

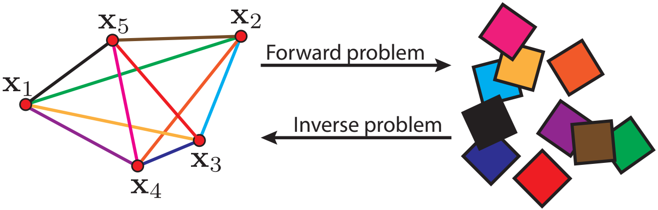

- Forward problem: given the points $\mathbf{X}$, find the EDM $\text{edm}(\mathbf{X})$.
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*:
  - Significantly harder than the forward problem!

Forward and inverse problem related to EDMs

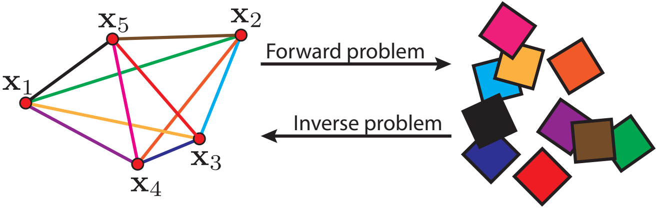

- Forward problem: given the points $\mathbf{X}$, find the EDM edm($\mathbf{X}$).
- Inverse problem: given the distances $\|\mathbf{x}_i - \mathbf{x}_j\|^2$, find the points $\mathbf{X}$.
  - Distances may be *noisy*;
  - Some distances may be *missing*;
  - We may loose the naming of the *distances*:
  - Significantly harder than the forward problem!

Sometimes we measure distances without names

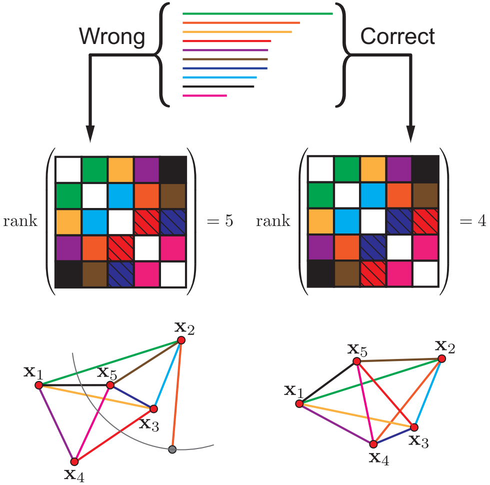

- We measure a set of distances, we lost their labels: combinatorial problem;
- Correct labels generate a low-rank matrix;
- Wrong labels generate a full-rank matrix;
- Hard to solve; harder when the distances are noisy.

Overview

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

Application 1: Can you hear the shape of a room?

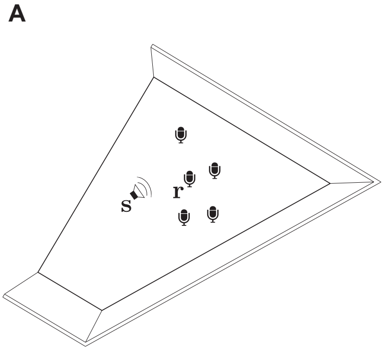

- One cannot hear the shape of a drum.... but can one hear the shape of a room? (Using a source and 5+ microphones)
- Each microphone records the room impulse response (B), the peaks being the reflections on the wall.
- Using the image source model, we transform walls to points.
- Problem 1: unlabeled distance problem
- Problem 2: some peaks do not correspond to any reflection, some reflections may be missing.

Application 1: Can you hear the shape of a room?

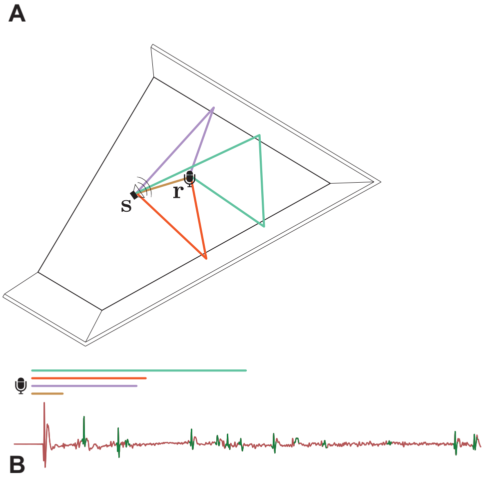

- One cannot hear the shape of a drum.... but can one hear the shape of a room? (Using a source and 5+ microphones)
- Each microphone records the room impulse response (B), the peaks being the reflections on the wall.
- Using the image source model, we transform walls to points.
- Problem 1: unlabeled distance problem
- Problem 2: some peaks do not correspond to any reflection, some reflections may be missing.

Application 1: Can you hear the shape of a room?

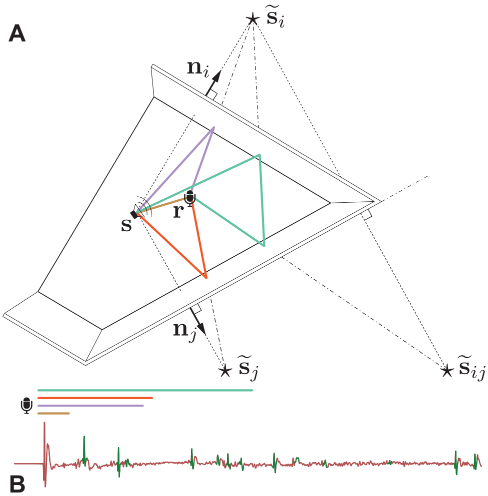

- One cannot hear the shape of a drum.... but can one hear the shape of a room? (Using a source and 5+ microphones)
- Each microphone records the room impulse response (B), the peaks being the reflections on the wall.
- Using the image source model, we transform walls to points.
- Problem 1: unlabeled distance problem
- Problem 2: some peaks do not correspond to any reflection, some reflections may be missing.

Application 1: Can you hear the shape of a room?

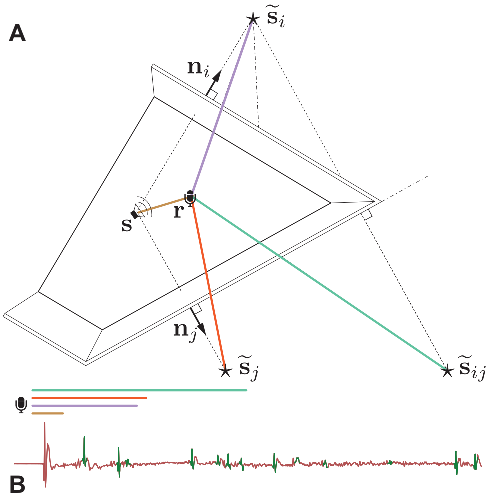

- One cannot hear the shape of a drum.... but can one hear the shape of a room? (Using a source and 5+ microphones)
- Each microphone records the room impulse response (B), the peaks being the reflections on the wall.
- Using the image source model, we transform walls to points.
- Problem 1: unlabeled distance problem
- Problem 2: some peaks do not correspond to any reflection, some reflections may be missing.

Application 1: Echo sorting algorithm

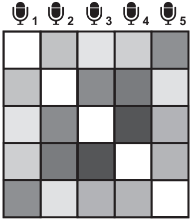

- We know the geometry of the microphone array, that is an EDM.
- For each microphone, we have a set of possible distances to the image sources,
- Location of the image sources: combinatorial search.
- EDM completion regularized by the rank is stable to noise and computationally feasible.

Application 1: Echo sorting algorithm

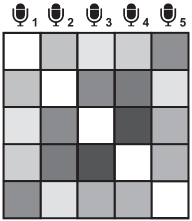

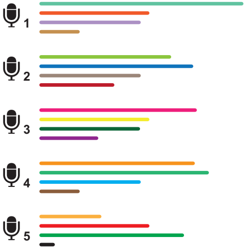

- We know the geometry of the microphone array, that is an EDM.
- For each microphone, we have a set of possible distances to the image sources,
- Location of the image sources: combinatorial search.
- EDM completion regularized by the rank is stable to noise and computationally feasible.

Application 1: Echo sorting algorithm

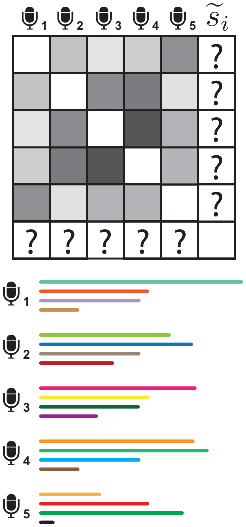

- We know the geometry of the microphone array, that is an EDM.
- For each microphone, we have a set of possible distances to the image sources,
- Location of the image sources: combinatorial search.
- EDM completion regularized by the rank is stable to noise and computationally feasible.

Application 1: Echo sorting algorithm

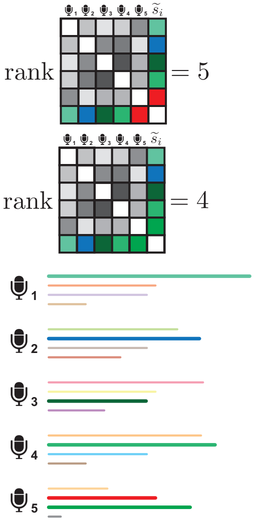

- We know the geometry of the microphone array, that is an EDM.
- For each microphone, we have a set of possible distances to the image sources,
- Location of the image sources: combinatorial search.
- EDM completion regularized by the rank is stable to noise and computationally feasible.

Application 1: It works... at EPFL

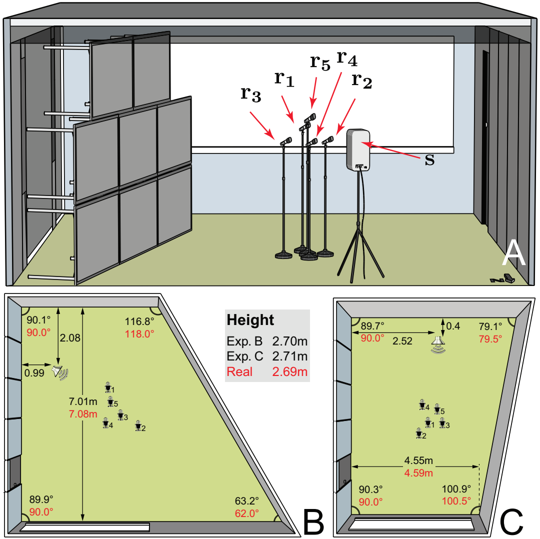

Application 1: It works... at the Lausanne cathedral!

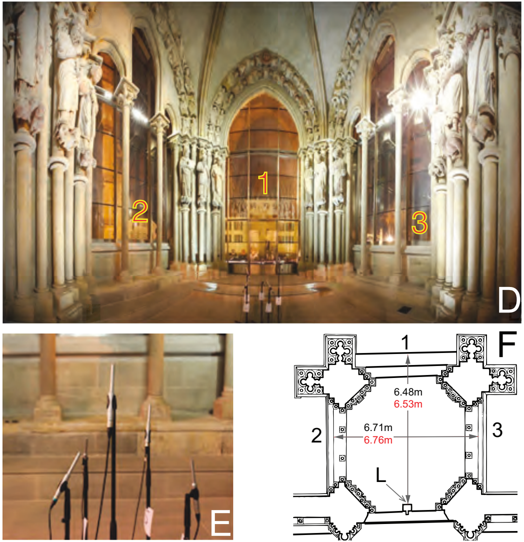

When we know the room, we can rake the cocktail party!

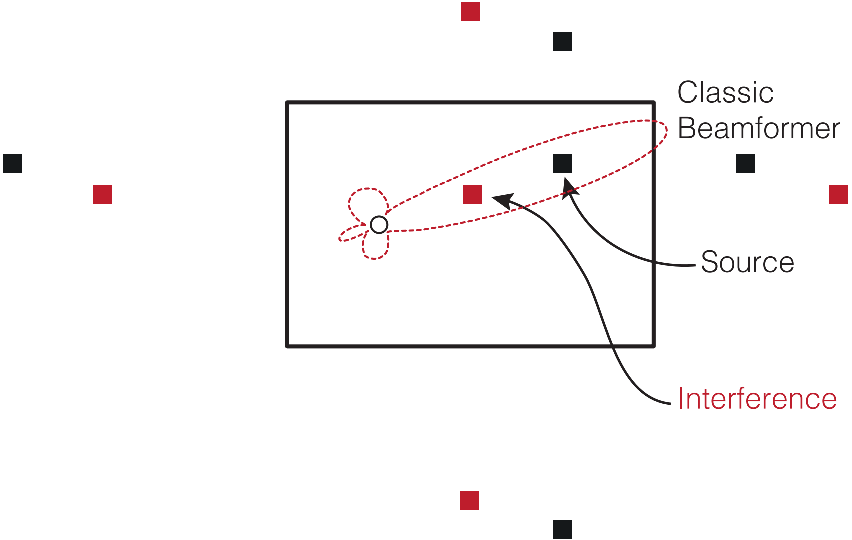

When we know the room, we can rake the cocktail party!

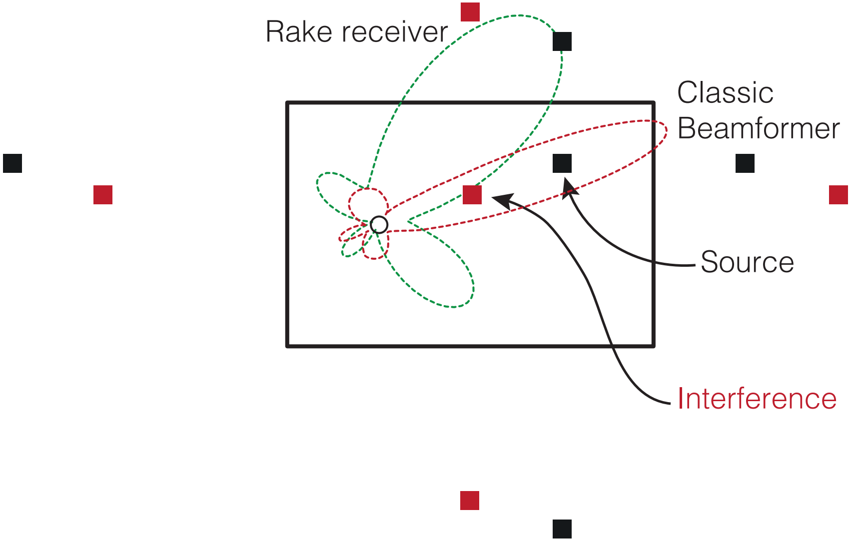

Application 2: Sparse Phase Retrieval

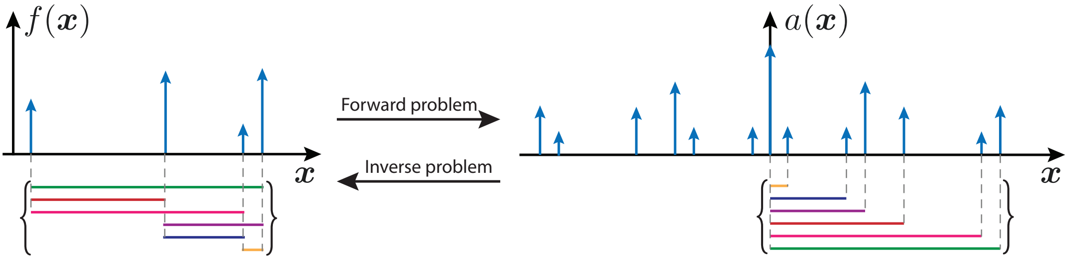

- (A) We measure a sparse auto-correlation function $a(x)$.
- The support of the ACF are the distances between the elements of the original signal.
- The distances are unlabeled and noisy.
- (B) We recover the support of $f(x)$ by labeling the distances.
- Knowing the support, it is easy to recover the amplitudes.

Application 2: Algorithm for Sparse Phase Retrieval

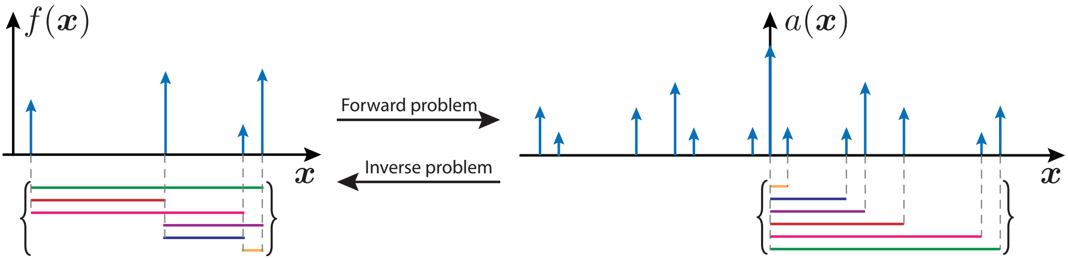

- Distance measurements are shift invariant, we set the first element at the origin.
- The support of the signal is a subset of the support of the autocorrelation function.
- Greedy algorithm selects the best element from the support of the measured ACF.
- Cost function: distance between the support of the ACF of the partial solution and the support of the measured ACF.

Application 2: Algorithm for Sparse Phase Retrieval

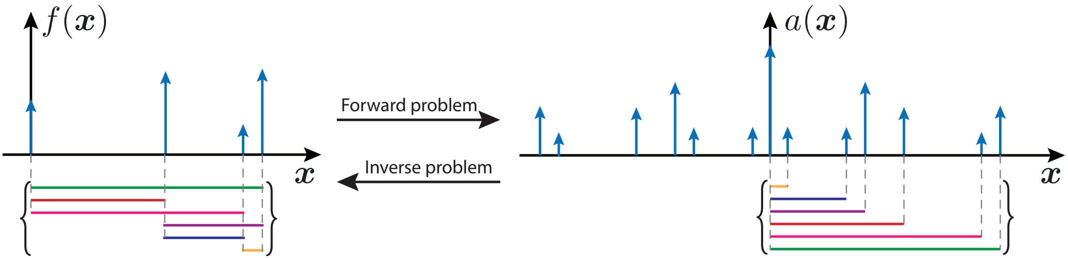

- Distance measurements are shift invariant, we set the first element at the origin.
- The support of the signal is a subset of the support of the autocorrelation function.
- Greedy algorithm selects the best element from the support of the measured ACF.
- Cost function: distance between the support of the ACF of the partial solution and the support of the measured ACF.

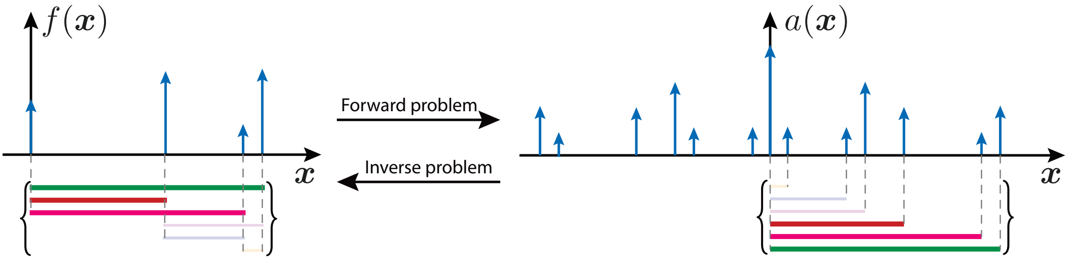

- Distance measurements are shift invariant, we set the first element at the origin.
- The support of the signal is a subset of the support of the autocorrelation function.
- Greedy algorithm selects the best element from the support of the measured ACF.
- Cost function: distance between the support of the ACF of the partial solution and the support of the measured ACF.

Application 2: Algorithm for Sparse Phase Retrieval

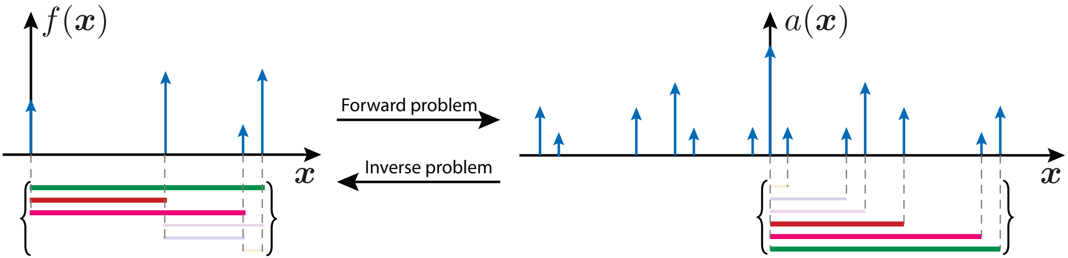

- Distance measurements are shift invariant, we set the first element at the origin.
- The support of the signal is a subset of the support of the autocorrelation function.
- Greedy algorithm selects the best element from the support of the measured ACF.
- Cost function: distance between the support of the ACF of the partial solution and the support of the measured ACF.

Application 2: Algorithm performance

- Percentage of success as a function of the complexity of the signal (# of deltas) and noise affecting the ACF.

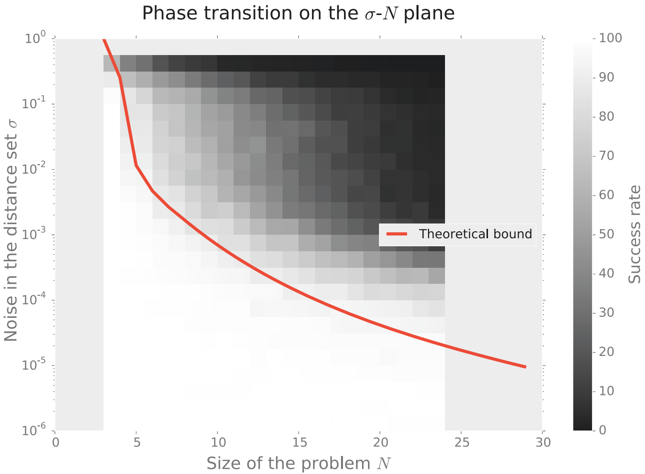

Application 3: Echo SLAM

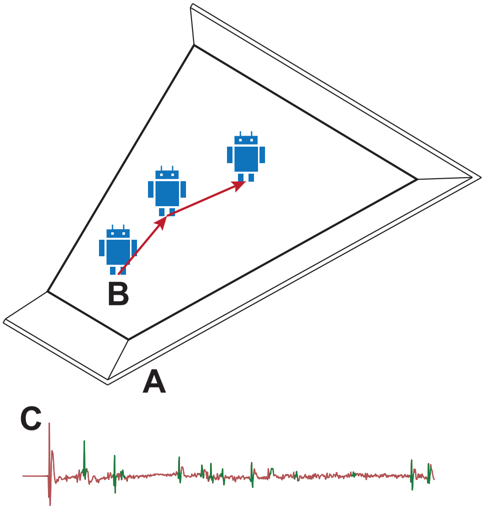

- We would like to localize a robot (B) inside an unknown room (A)
- The robot walks autonomously; we know statistics of the step size and direction.
- After each step, a sound is emitted and measures the room impulse response (C):
  - Source located on the robot or,
  - Source fixed in the room.
- Each peak in the RIR likely corresponds to an image source.
- Echo SLAM: Simultaneous Localization And Mapping based on echoes.

Application 3: Echo SLAM

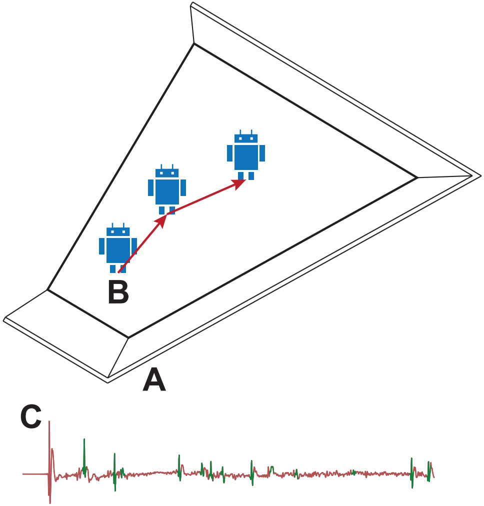

- We would like to localize a robot (B) inside an unknown room (A)
- The robot walks autonomously; we know statistics of the step size and direction.
- After each step, a sound is emitted and measures the room impulse response (C):
  - Source located on the robot or,
  - Source fixed in the room.
- Each peak in the RIR likely corresponds to an image source.
- Echo SLAM: Simultaneous Localization And Mapping based on echoes.

Application 3: Echo SLAM

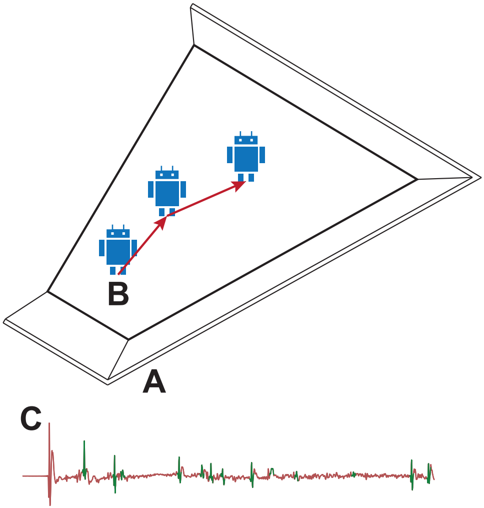

- We would like to localize a robot (B) inside an unknown room (A)
- The robot walks autonomously; we know statistics of the step size and direction.
- After each step, a sound is emitted and measures the room impulse response (C):
  - Source located on the robot or,
  - Source fixed in the room.
- Each peak in the RIR likely corresponds to an image source.
- Echo SLAM: Simultaneous Localization And Mapping based on echoes.

Application 3: Echo SLAM

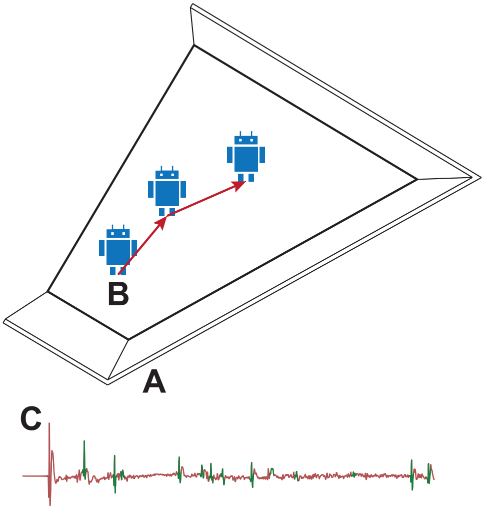

- We would like to localize a robot (B) inside an unknown room (A)
- The robot walks autonomously; we know statistics of the step size and direction.
- After each step, a sound is emitted and measures the room impulse response (C):
  - Source located on the robot or,
  - Source fixed in the room.
- Each peak in the RIR likely corresponds to an image source.
- Echo SLAM: Simultaneous Localization And Mapping based on echoes.

Application 3: Echo SLAM with fixed sound source

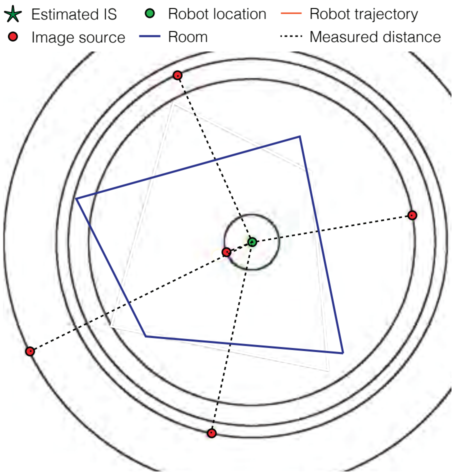

1st step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

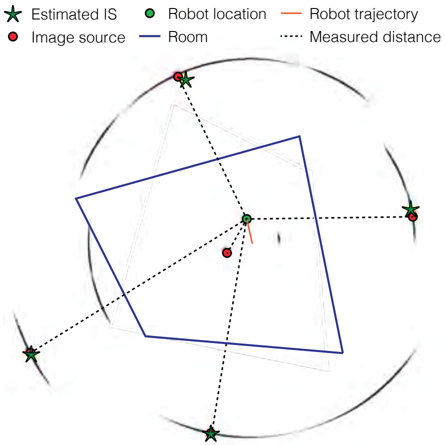

2nd step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

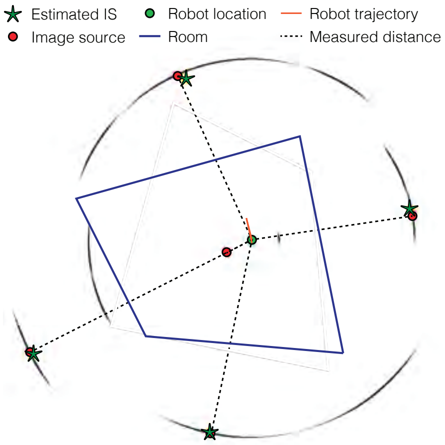

3rd step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

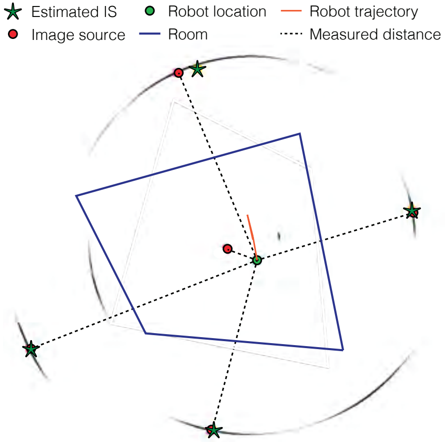

4th step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

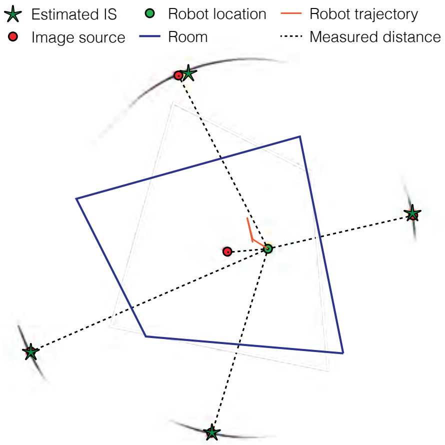

5th step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

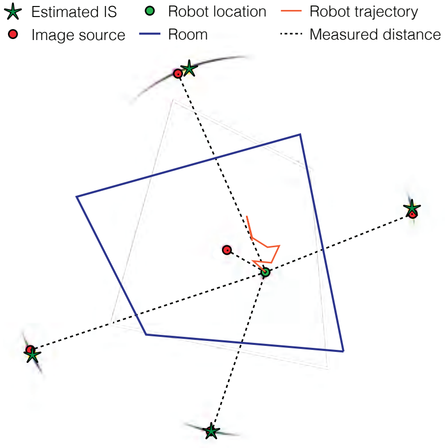

9th step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 3: Echo SLAM with fixed sound source

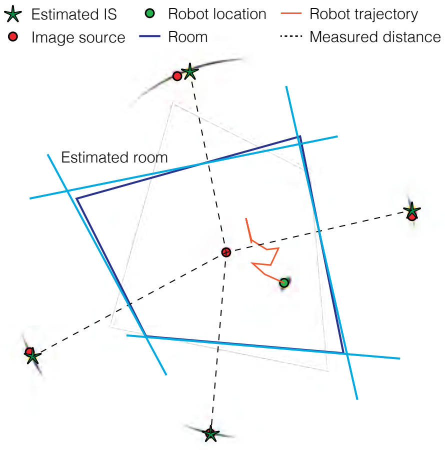

10th step of the robot

- At each step, we measure the RIR and determine the distance from the image sources.
- At the 1st step, initialization of the probability distribution of the image sources.
- At each step, the robot moves autonomously and we update the probability distributions.
- The accuracy of the estimated image sources improves step after step.
- The estimated image sources are then used to estimate the room shape.

Application 4: Multidimensional unfolding

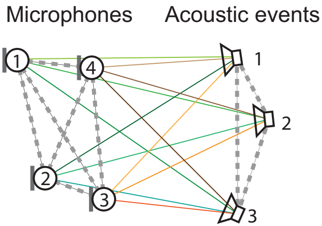

$D =$ 

- Multidimensional unfolding:
  - Divide the points into two subsets,
  - Measure the distances between the points belonging to different subsets,
  - Other distances are unknown.
- Example: calibration of microphones with sources at unknown locations.
- Many distances are missing, the measured ones are noisy.

Application 4b: Calibrate ten microphones with a fingersnap

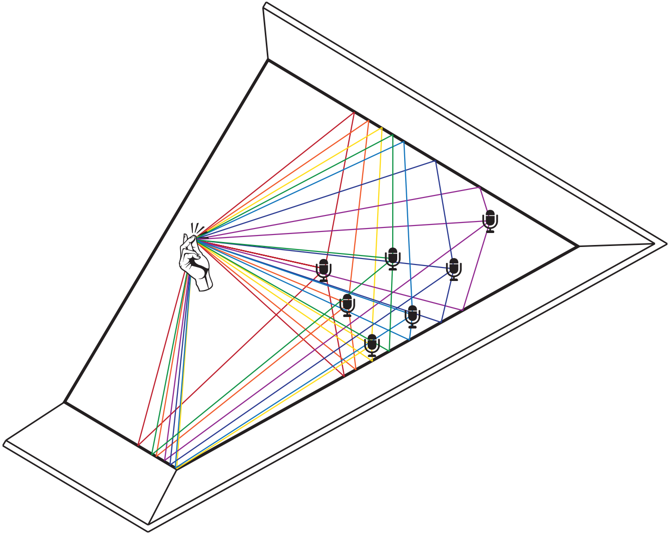

- Unknown room.
- Every microphone measures the echoes coming from the walls (or image sources).
- We measure the unlabeled distances between microphones and sources.
- Sort the echoes to localize the mics!

Application 5: Calibration in ultrasound tomography

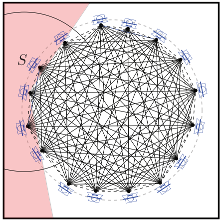

A

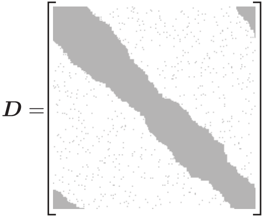

B

- Ultrasound transceivers deployed on a circle (A), their location is unknown.
- Each transceiver can measure the distance w.r.t. the transceivers in front.
- We can fill the EDM, with structured missing entries (B).
- Problem: estimate the locations of the transceiver to improve the imaging performance.

Application 5: Calibration in ultrasound tomography

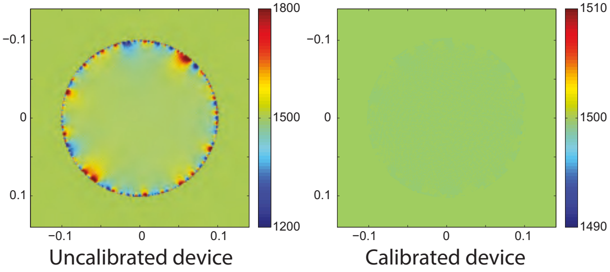

- Solution of the ultrasound tomography when measuring water (1500 m/s).
- Calibration improves significantly the tomographic imaging.

Outline

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

How do we skin the cat?

- Wide array of possible applications,
- We need flexible and efficient algorithms!

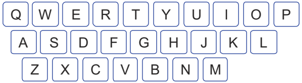

QWERTY Keyboard

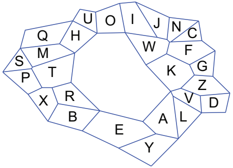

EDM Keyboard

- Large matrix with small rank, having missing and/or noisy entries.
- **Idea:** enforce the rank while minimizing the distance from the measurements.
- **Problem:** Rank constraints are **non-convex**: hard to enforce!

Algorithm 1: Multidimensional scaling (MDS)

Classical Multidimensional scaling (MDS)

- Choose $\mathbf{x}_1 = \mathbf{0}$, then the first column of $\mathbf{D} = \text{edm}(\mathbf{X})$, call it d_1, is $\text{diag}(\mathbf{X}^\top \mathbf{X})$,
- Construct the Gram matrix: $\mathbf{G} = -\frac{1}{2}(\mathbf{D} - \mathbf{1}d_1^\top - d_1\mathbf{1}^\top)$,
- The point set is obtained via SVD $\mathbf{G} = \mathbf{X}^\top \mathbf{X}$.

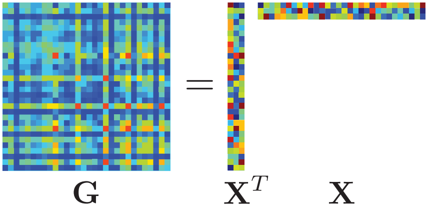

Problem: How do we deal with missing entries?

Classical Multidimensional scaling (MDS)

- Choose $\mathbf{x}_1 = \mathbf{0}$, then the first column of $\mathbf{D} = \text{edm}(\mathbf{X})$, call it d_1, is $\text{diag}(\mathbf{X}^\top \mathbf{X})$,
- Construct the Gram matrix: $\mathbf{G} = -\frac{1}{2}(\mathbf{D} - \mathbf{1}d_1^\top - d_1\mathbf{1}^\top)$,
- The point set is obtained via SVD $\mathbf{G} = \mathbf{X}^\top \mathbf{X}$.

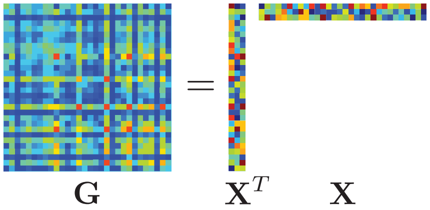

Problem: How do we deal with missing entries?

Algorithm 2: Alternated rank descent

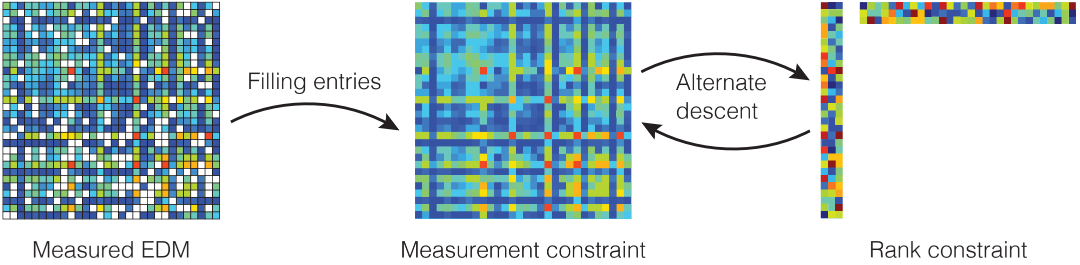

- We measure an EDM with missing and noisy entries.
- We fill the missing entries by enforcing the rank constraint.
- Alternate between enforcing the measured distances and the rank constraint.

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

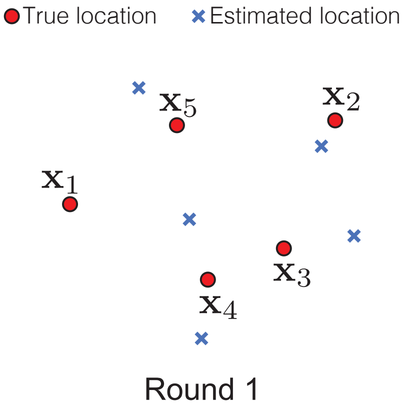

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} \left(\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij}\right)^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

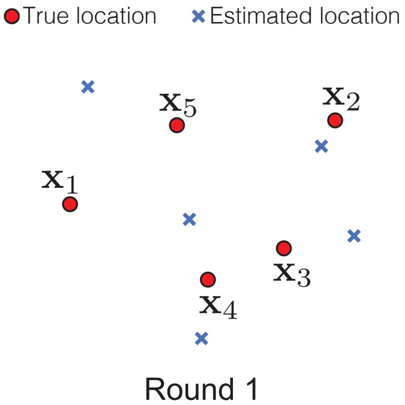

Algorithm 3: Alternating descent of the s-stress function

Alternating descent on the s-stress function

- Minimize the s-stress,
$$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

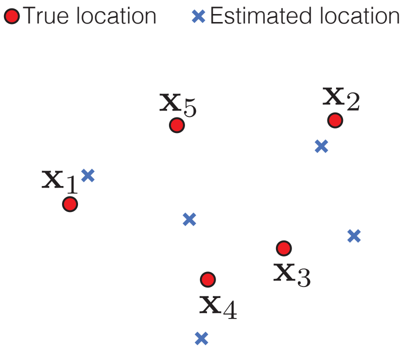

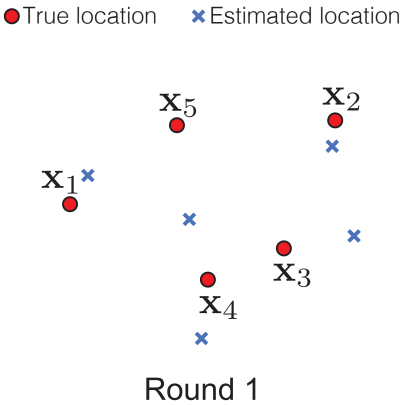

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

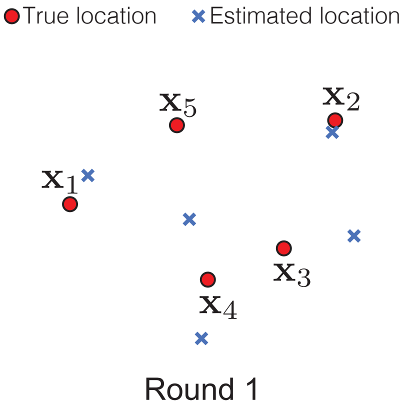

Alternating descent on the s-stress function

- Minimize the s-stress, $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

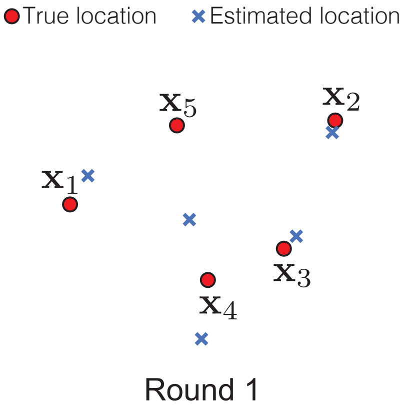

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

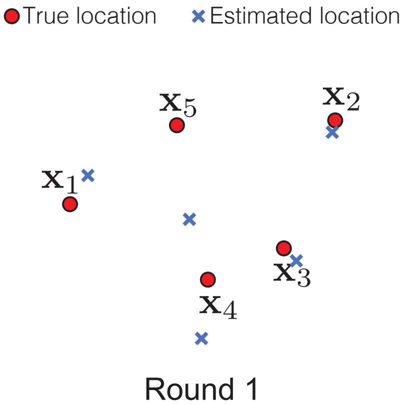

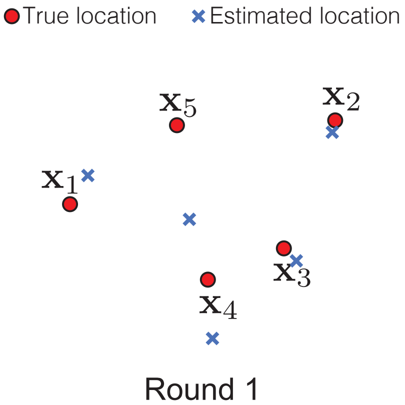

Alternating descent on the s-stress function

- Minimize the s-stress,
$$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} \left(\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij}\right)^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

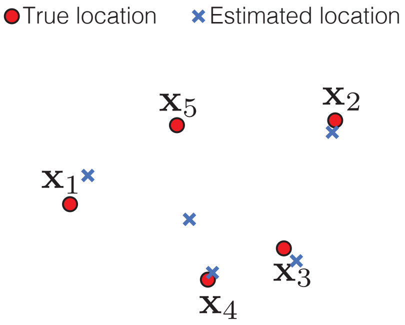

Round 1

Alternating descent on the s-stress function

- Minimize the s-stress, $\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

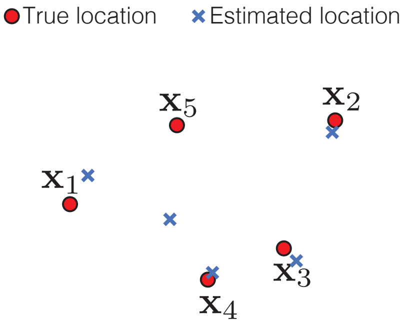

Round 1

Alternating descent on the s-stress function

- Minimize the s-stress, $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

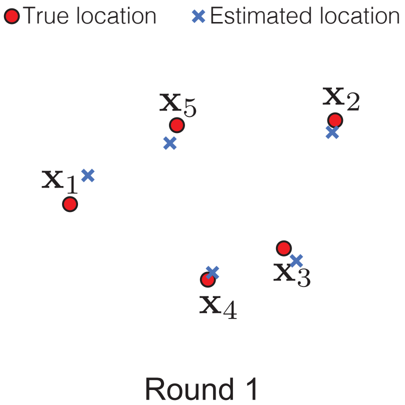

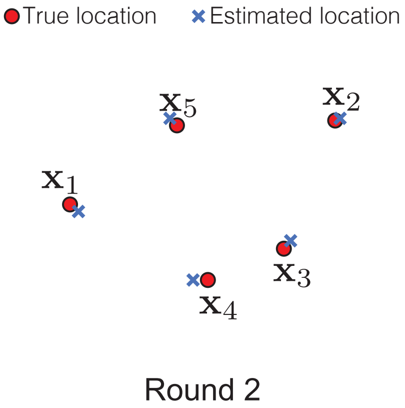

Alternating descent on the s-stress function

- Minimize the s-stress, $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

Algorithm 3: Alternating descent of the s-stress function

Alternating descent on the s-stress function

- Minimize the s-stress,
  $$\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

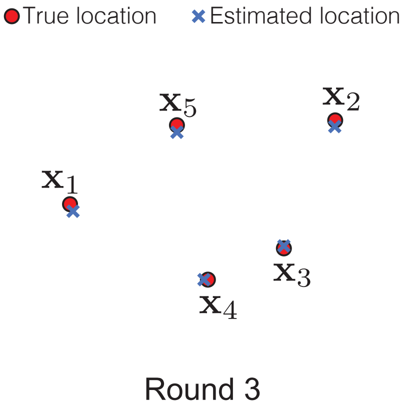

Algorithm 3: Alternating descent of the s-stress function

Alternating descent on the s-stress function

- Minimize the s-stress, $\sum_{(i,j)\in E} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - \widetilde{d}_{ij})^2$
- Alternate between minimizing individual coordinate component of each point $\mathbf{x}_i$

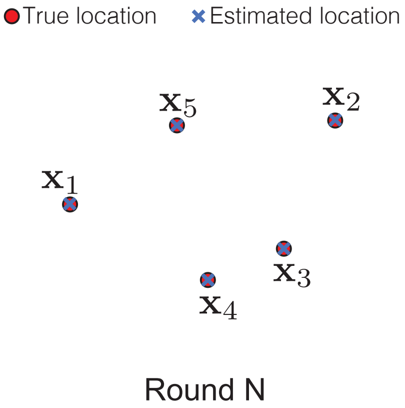

Algorithm 4: Semidefinite relaxation

- The Gram matrix $\mathbf{G}$ is positive semidefinite,
  $\mathbf{W}$ is the mask defining the measured distances,
  $\widetilde{\mathbf{D}}$ contains the measured distances.
- $\mathbf{D}$ is linear in $\mathbf{X}$ (remember,
  $\mathbf{D} = \operatorname{diag}(\mathbf{X})\mathbf{1}^T + \mathbf{1}\operatorname{diag}(\mathbf{X})^T - 2\mathbf{X}$)
- Together with the rank, it's a complete description!

$$\begin{array}{ll} \underset{\mathbf{G}}{\text{minimize}} & \left\| \mathbf{W} \circ \left(\widetilde{\mathbf{D}} - \operatorname{edm}(\mathbf{G}) \right) \right\|_F^2 \\ \text{subject to} & \mathbf{G} \text{ positive semidefinite} \\ & \mathbf{G} \text{ geometrically centered} \\ & \cancel{\operatorname{rank}(\mathbf{G}) \le d} \end{array}$$

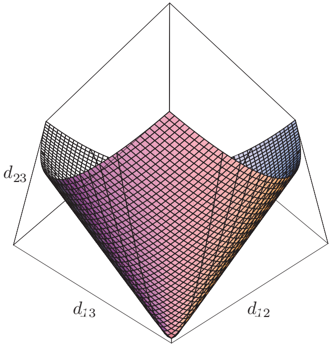

The EDM cone for 3 points

Algorithms comparison

Algorithm	**Pros**	**Cons**
Classical MDS	Very simple Handles noise	No missing distances Unnatural cost function
Alternated rank descent	Missing distances, noise Handles noise	Scales poorly Unnatural cost function
Alternating Descent on s-stress	Simple Good cost function (s-stress) Missing distances, noise	Convergence issues
SDR	Good cost function Missing distances, noise Handles constraints on distances	Slow Can't enforce the embedding dimension

Algorithms: performance w.r.t. missing entries

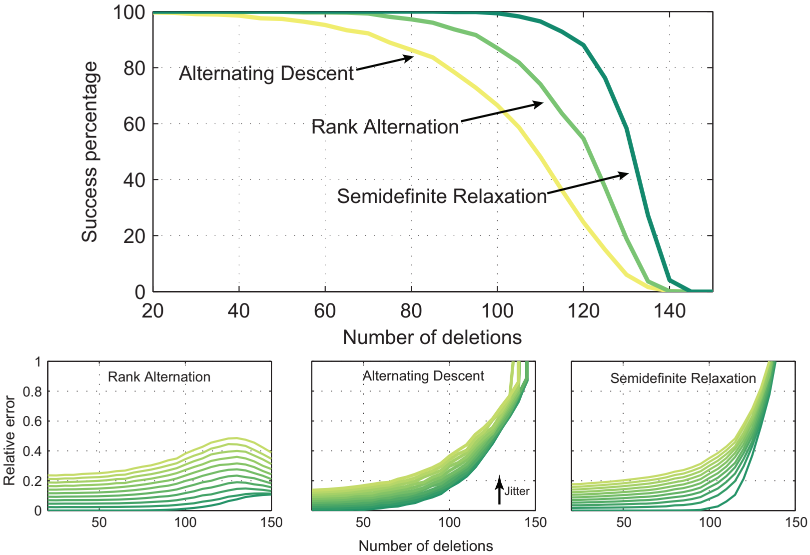

Algorithms: what happens for a specific application?

- Consider the calibration problem with an increasing number of sound sources,
- We have more information, but not every algorithm can exploit it!

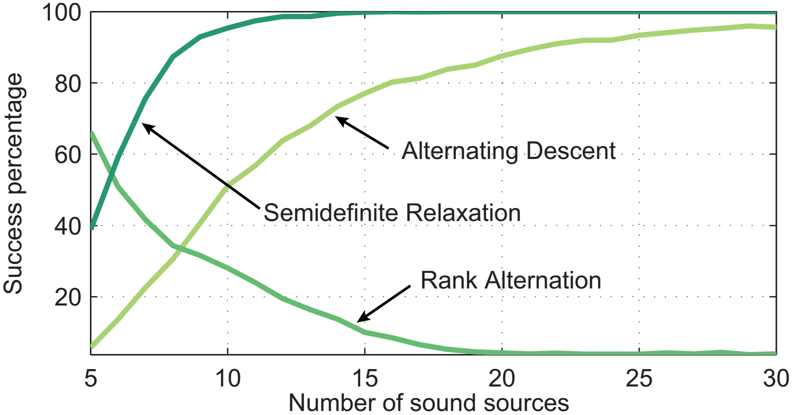

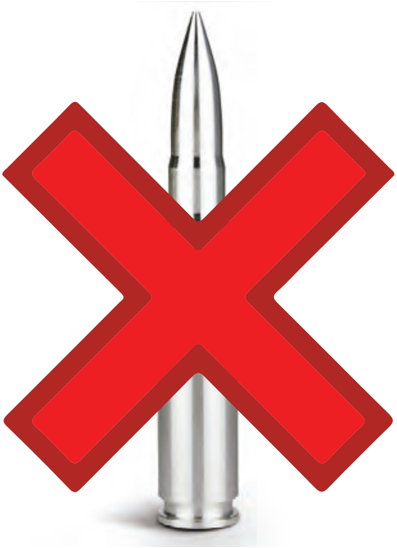

Overview

- Motivation
- Euclidean Distance Matrices (EDM) and their properties
- Forward and inverse problems related to EDMs
- Applications of EDMs
- Algorithms for EDMs
- Conclusions and open problems

A small selection of open problems

- Distance matrices on manifolds

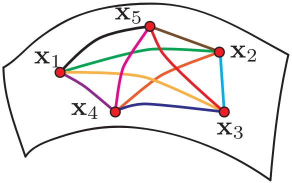

- Projections of EDMs on subspaces

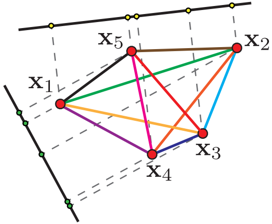

- Efficient algorithms for distance labeling

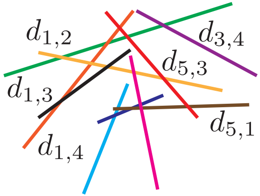

- Analytical local minimum of s-stress

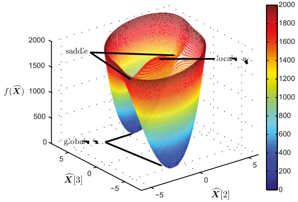

Conclusions

Application	missing distances	noisy distances	unlabeled distances
Wireless sensor networks *	✔	✔	×
Molecular conformation	✔	✔	×
Hearing the shape of a room *	×	✔	✔
Echo SLAM *	×	✔	✔
Indoor localization	×	✔	✔
Calibration *	✔	✔	×
Sparse phase retrieval *	×	✔	✔

* Applications that we have discussed in this talk.

Literature

- I.Dokmanić, R.Parhizkar, J.Ranieri, and M.Vetterli, *Euclidean Distance Matrices: A Short Walk Through Theory, Algorithms and Applications*, to appear (November 2015) in IEEE Signal Processing Magazine.
- R.Parhizkar, A.Karbasi, S.Oh, M.Vetterli, *Calibration Using Matrix Completion with Application to Ultrasound Tomography*, IEEE Transaction on Signal Processing, 2013.
- I.Dokmanić, R.Parhizkar, A.Walther, Y.M.Lu, and M.Vetterli, *Acoustic echoes reveal room shape*, Proceedings of the National Academy of Sciences, 2013.
- J.Ranieri, A.Chebira, Y.M.Lu, and M.Vetterli, *Phase Retrieval for Sparse Signals: Uniqueness Conditions and Algorithms* in preparation.
- R.Parhizkar, *Euclidean Distance Matrices: Properties, Algorithms and Applications*, PhD Thesis EPFL, 2013.
- I.Dokmanić, *Listening to Distances and Hearing Shapes: Inverse Problems in Room Acoustics and Beyond*, PhD Thesis EPFL, 2015.

Acknowledgments

- Farid Movahedi Naini for helping with the numerical experiments.
- Funding resources:
  - Swiss National Science Foundation, Grant 200021 138081.
  - European Research Council SPARSAM, Grant 247006.
  - Google Focused Research Award.
  - Google PhD fellowship.

Thank you for your attention!

Googlecar

Batmobile with echolocation