Individual Preference Stability for Clustering

Saba Ahmadi*(TTIC) Pranjal Awasthi*(Google) Samir Khuller*(NU)

Matthäus Kleindessner* (Amazon) Jamie Morgenstern* (Google/UW)

Pattara Sukprasert*(NU) Ali Vakilian*(TTIC)

*equal contribution

Clustering

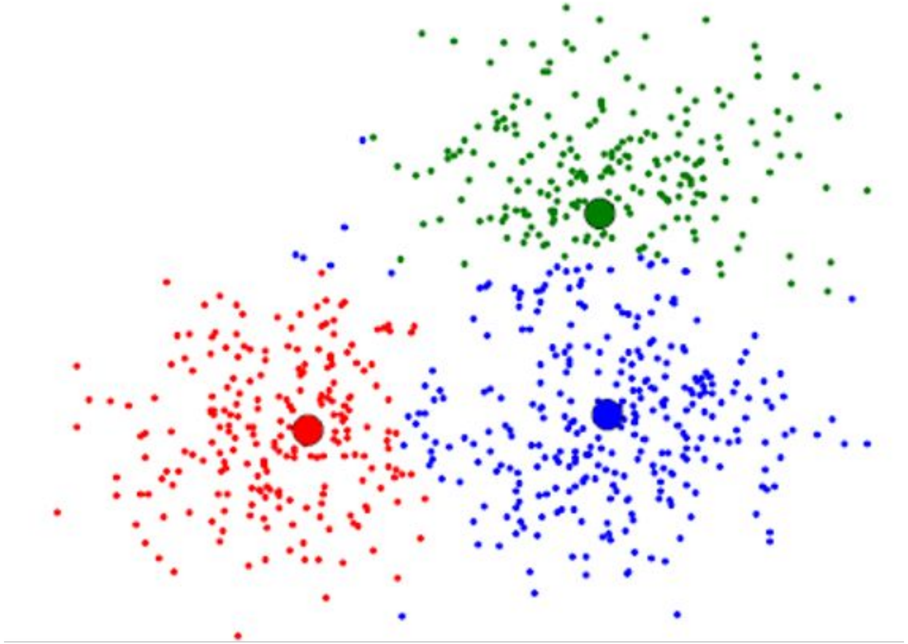

Clustering Applications

- Facility Locations

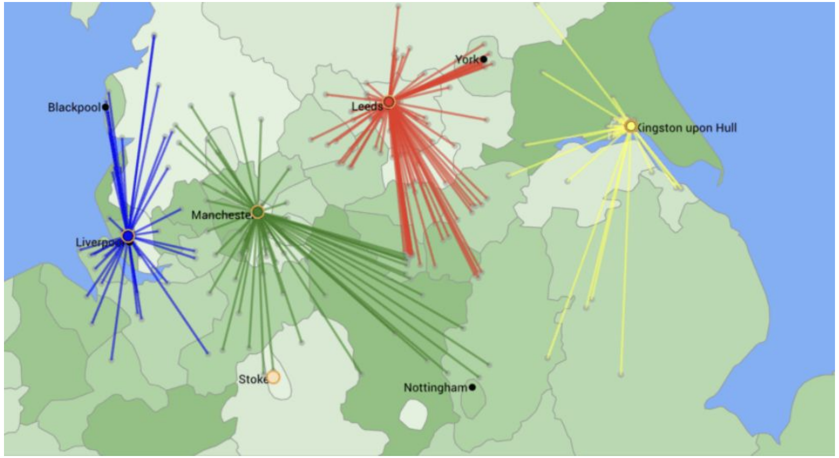

Clustering Applications

- Text Summarization

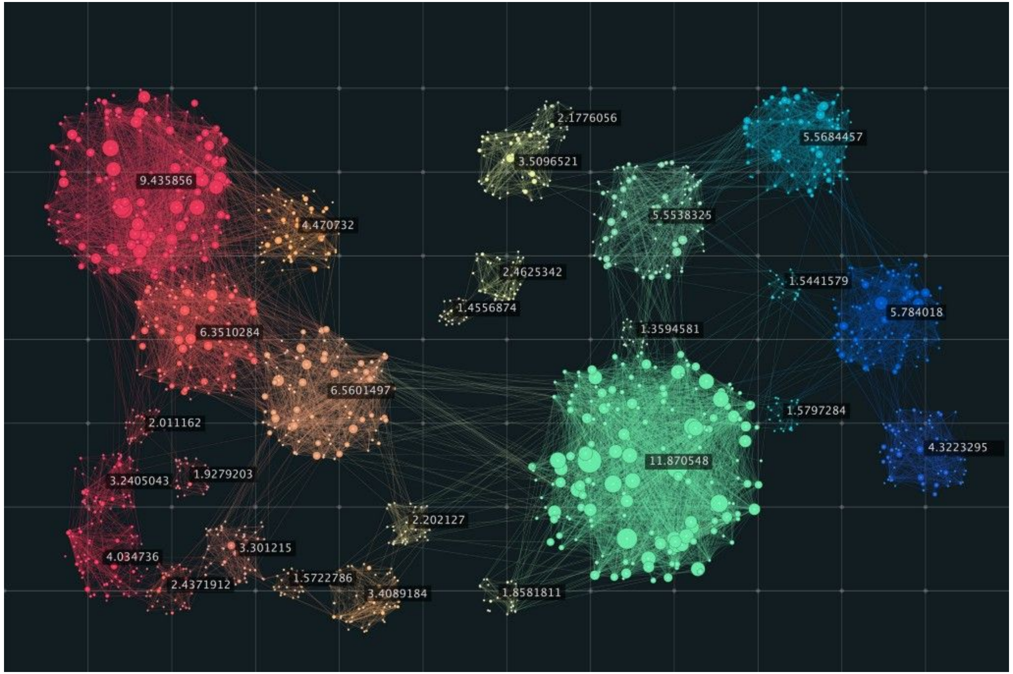

Clustering Applications

- Partition
    - Loan applications

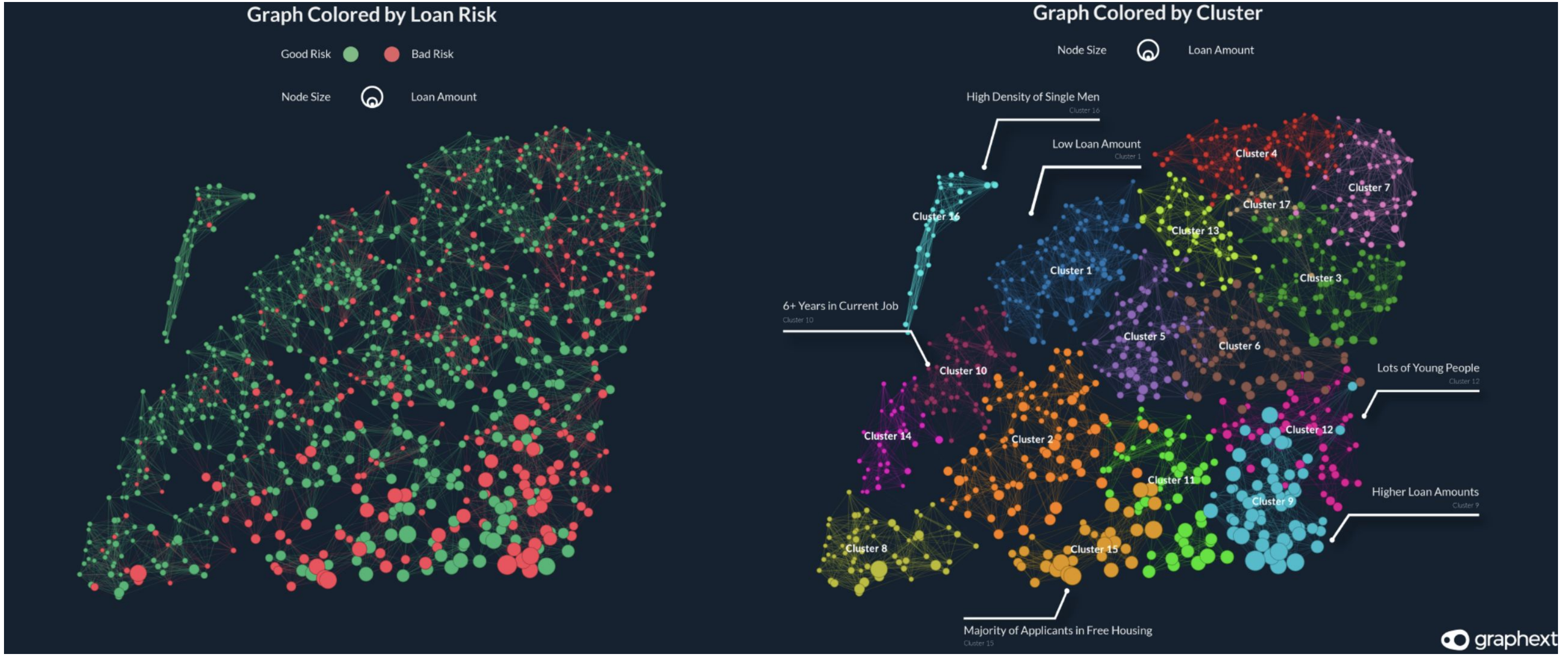

Centroid-Based Clusterings

Centroid-Based Clusterings

- k-means

Centroid-Based Clusterings

- k-means
- k-medians

Centroid-Based Clusterings

- k-means
- k-medians
- k-center

Centroid-Based Clusterings

- k-means
- k-medians
- k-center

$$\text{k-means} : \min \sum_{x \in \mathcal{D}} d(x, C(x))^2$$

$$\text{k-medians} : \min \sum_{x \in \mathcal{D}} d(x, C(x))$$

$$\text{k-center} : \min \max_{x \in \mathcal{D}} d(x, C(x))$$

Centroid-Based Clusterings

- k-means
- k-medians
- k-center

Global objective function

$$\text{k-means} : \min \sum_{x \in \mathcal{D}} d(x, C(x))^2$$

$$\text{k-medians} : \min \sum_{x \in \mathcal{D}} d(x, C(x))$$

$$\text{k-center} : \min \max_{x \in \mathcal{D}} d(x, C(x))$$

Centroid-Based Clusterings

- k-means
- k-medians
- k-center
- **Unfair treatments**

Global objective function

$$\text{k-means} : \min \sum_{x \in \mathcal{D}} d(x, C(x))^2$$

$$\text{k-medians} : \min \sum_{x \in \mathcal{D}} d(x, C(x))$$

$$\text{k-center} : \min \max_{x \in \mathcal{D}} d(x, C(x))$$

Challenge

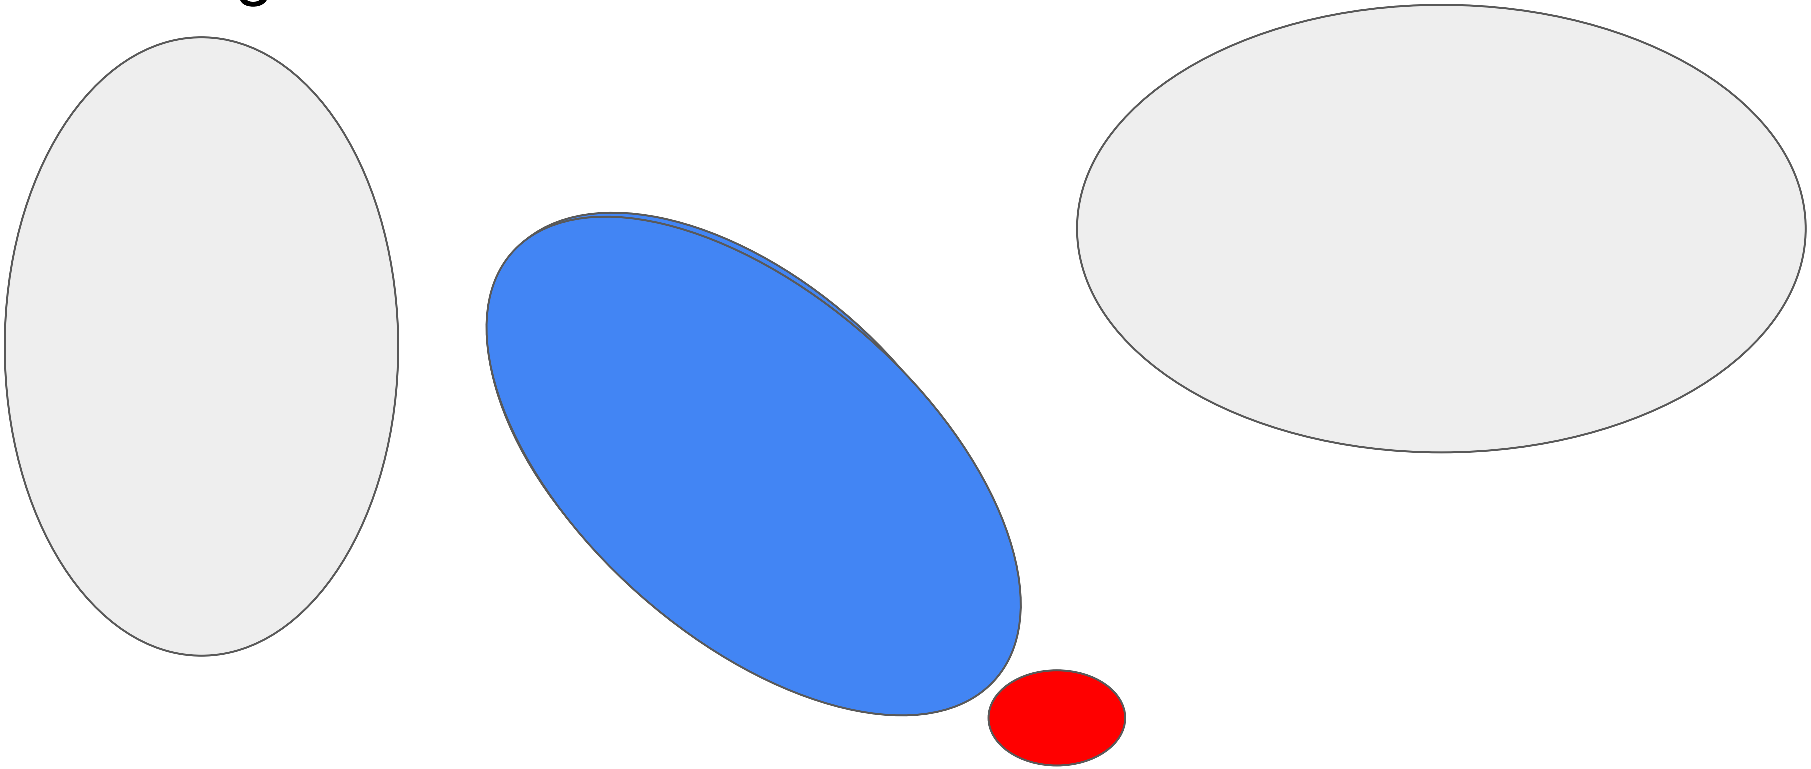

Challenge

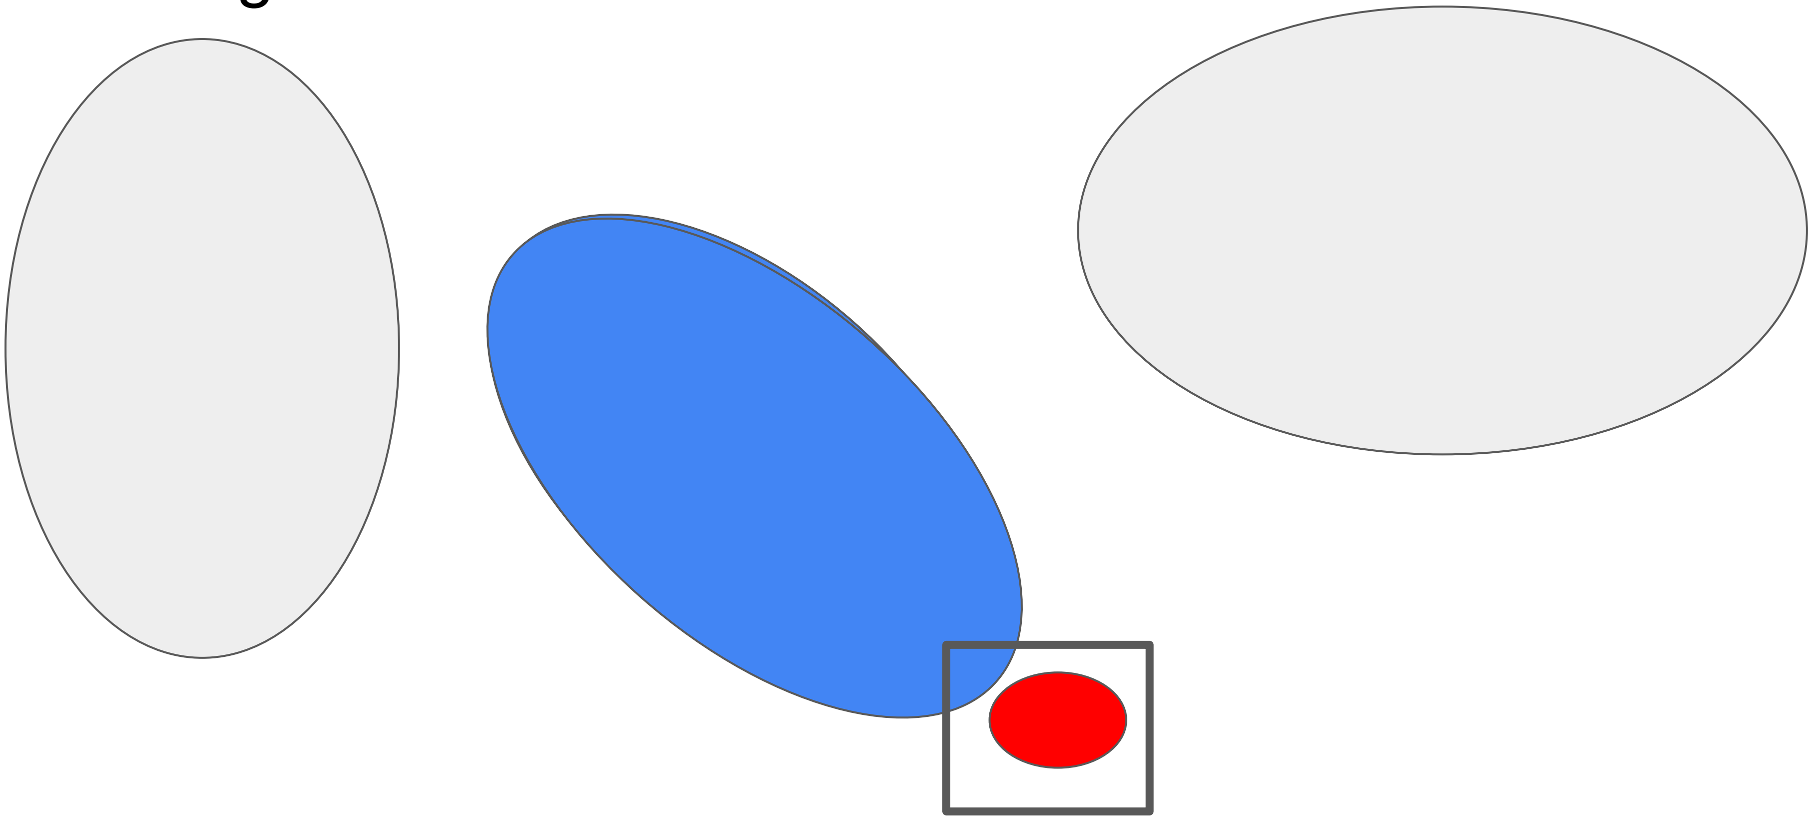

Challenge

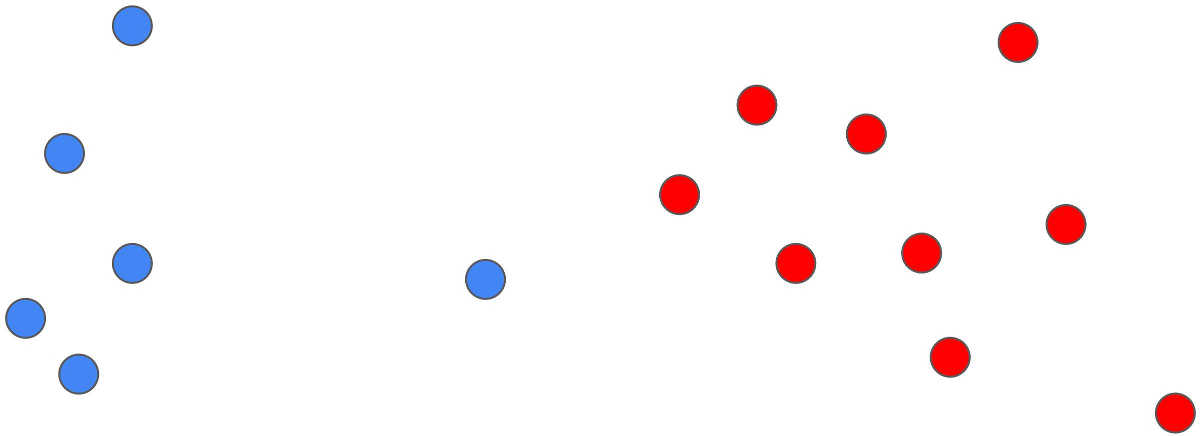

Challenge

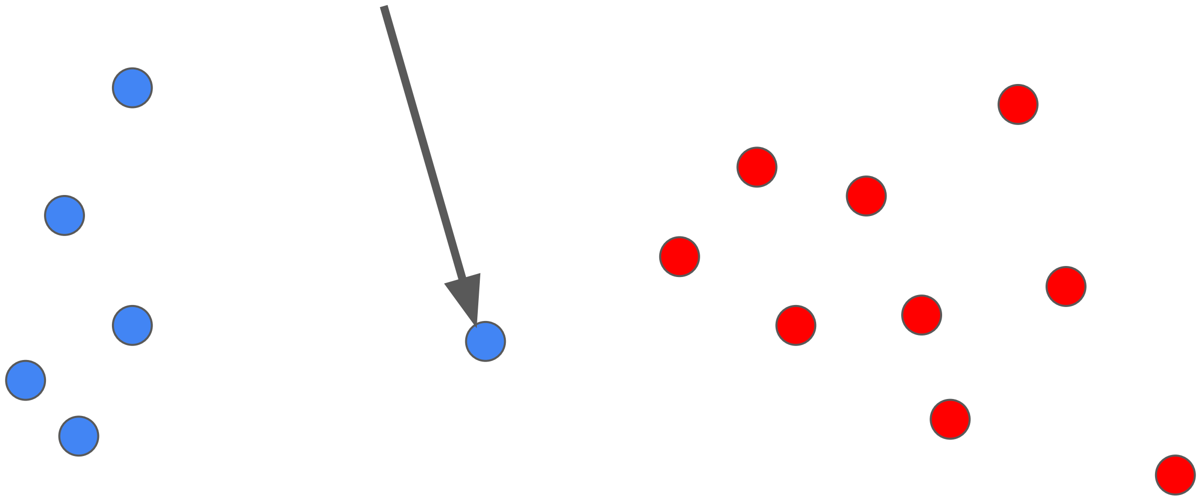

Challenge

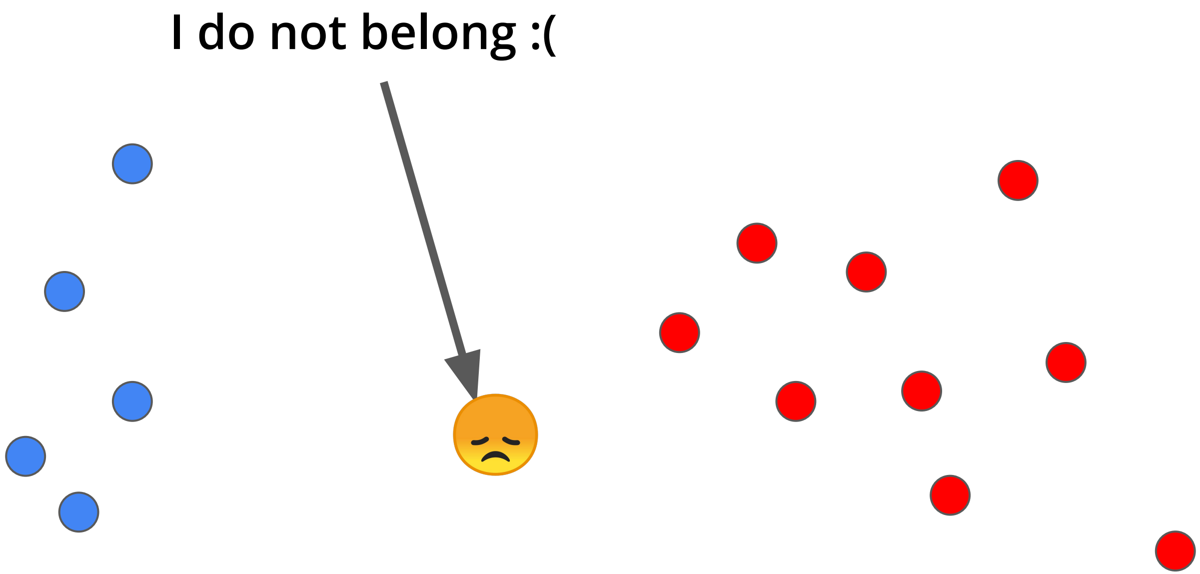

Challenge

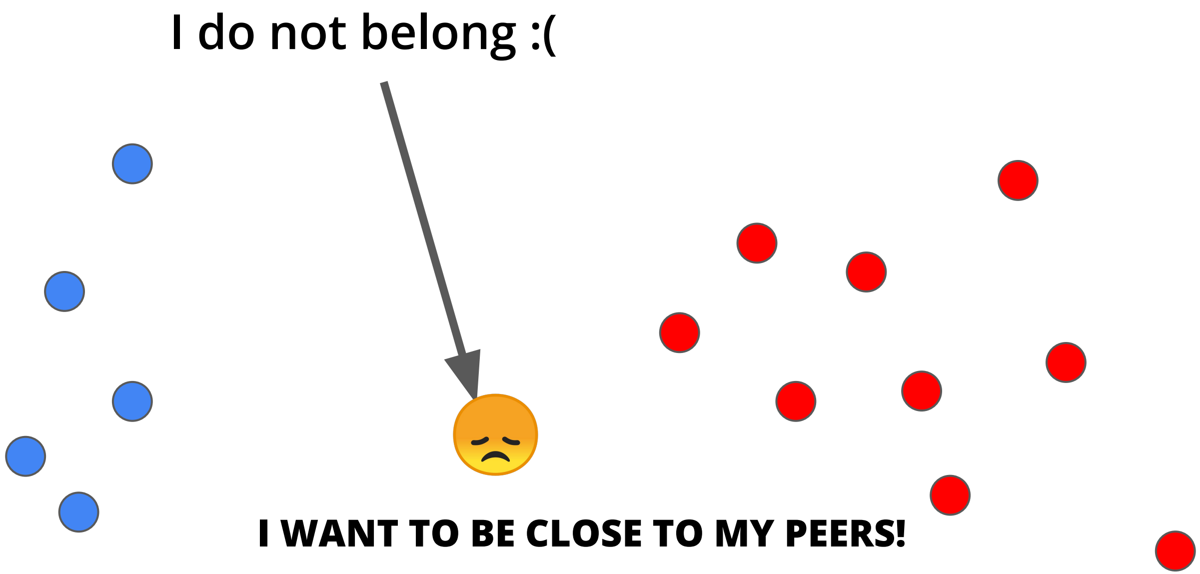

Fairness Notions

Fairness Notions

- Dwork et. al. (ITCS’12)
  - Similar individuals should be treated similarly

Fairness Notions

- Dwork et. al. (ITCS'12)
  - Similar individuals should be treated similarly
- Jung, Kannan, and Lutz (FORC'20)
  - Center in neighborhood

Fairness Notions

- Dwork et. al. (ITCS’12)
  - Similar individuals should be treated similarly
- Jung, Kannan, and Lutz (FORC’20)
  - Center in neighborhood
- Kleindessner, Awasthi, and Morgenstern (ICML’19)
  - Different demographics are well represented

Fairness Notions

- Dwork et. al. (ITCS'12)
  - Similar individuals should be treated similarly
- Jung, Kannan, and Lutz (FORC'20)
  - Center in neighborhood
- Kleindessner, Awasthi, and Morgenstern (ICML'19)
  - Different demographics are well represented
- Chen et. al. (ICML'19)
  - No large portion of points can form another cluster

Fairness Notions

- Dwork et. al. (ITCS'12)
    - Similar individuals should be treated similarly
- Jung, Kannan, and Lutz (FORC'20)
    - Center in neighborhood
- Kleindessner, Awasthi, and Morgenstern (ICML'19)
    - Different demographics are well represented
- Chen et. al. (ICML'19)
    - No large portion of points can form another cluster
- Many others (MV, ICML'20), (NC, NeurIPS'21), (VY, AISTATS'22) …

Individual Preference Stability

Individual Preference Stability

x is IP-stable iff: $$\frac{1}{|C(x)|-1}\sum_{y\in C(x)} d(x,y) \leq \frac{1}{|C_i|}\sum_{y\in C_i} d(x,y)$$

The clustering $\mathcal{C}$ is IP-stable if every $x \in \mathcal{D}$ is IP-stable

Individual Preference Stability

x is IP-stable iff: $$\frac{1}{|C(x)|-1}\sum_{y\in C(x)} d(x,y) \leq \frac{1}{|C_i|}\sum_{y\in C_i} d(x,y)$$

The clustering $\mathcal{C}$ is IP-stable if every $x \in \mathcal{D}$ is IP-stable

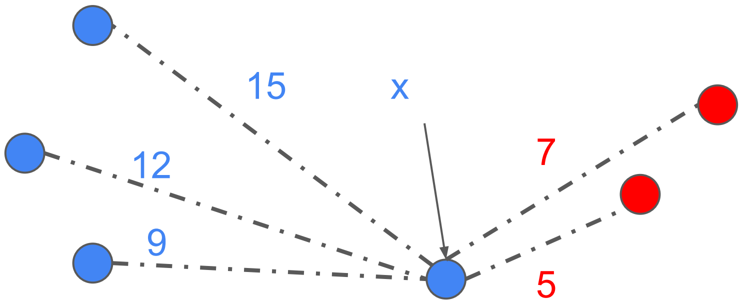

Individual Preference Stability

x is IP-stable iff: $$\frac{1}{|C(x)|-1}\sum_{y\in C(x)} d(x,y) \leq \frac{1}{|C_i|}\sum_{y\in C_i} d(x,y)$$

The clustering $\mathcal{C}$ is IP-stable if every $x \in \mathcal{D}$ is IP-stable

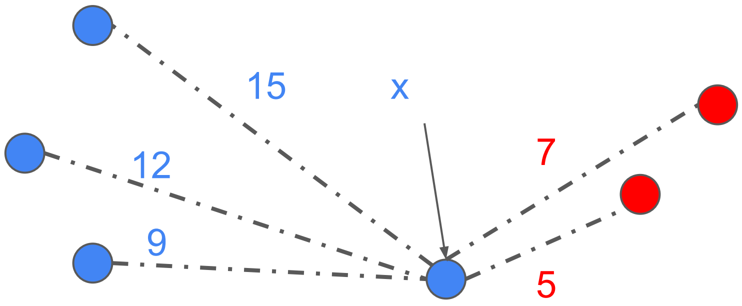

$$d(x, Blue) = \frac{9+12+15}{3} = 12$$

$$d(x, Red) = \frac{5+7}{2} = 6$$

Individual Preference Stability

x is IP-stable iff: $$\frac{1}{|C(x)|-1}\sum_{y\in C(x)} d(x,y) \leq \frac{1}{|C_i|}\sum_{y\in C_i} d(x,y)$$

The clustering $\mathcal{C}$ is IP-stable if every $x \in \mathcal{D}$ is IP-stable

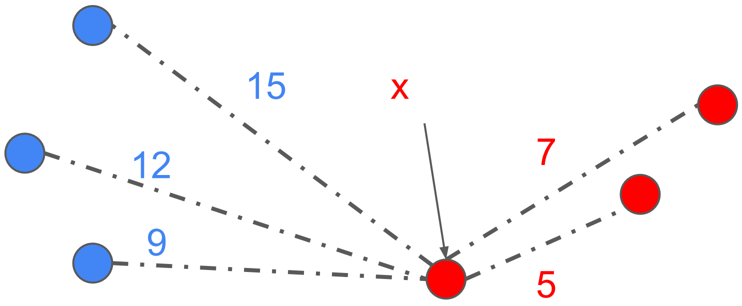

$$d(x, Blue) = \frac{9+12+15}{3} = 12$$

$$d(x, Red) = \frac{5+7}{2} = 6$$

Individual Preference Stability

x is IP-stable iff: $$\frac{1}{|C(x)|-1}\sum_{y\in C(x)} d(x,y) \leq \frac{1}{|C_i|}\sum_{y\in C_i} d(x,y)$$

x is happy if x is closer to people in the same cluster

The clustering $\mathcal{C}$ is IP-stable if every $x \in \mathcal{D}$ is IP-stable

Clustering $\mathcal{C} = (C_1, \ldots, C_k)$ is a happy clustering if everyone is happy

Problem Statement

- Given data points **D** and an integer **k**
- Can we find a happy clustering of **k** clusters?

Problem Statement

- Given data points **D** and an integer **k**
- Can we find a happy clustering of **k** clusters?
- No global objective involved (yet)

Results

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

Results

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - **Impossibility**
  - **NP-hardness**
- Experiments

Results

- Positive Results
  - **1D**
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

Results

- Positive Results
  - 1D
  - Trees for 2 clusters
  - **General Metric Space**
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

Can everyone be happy?

Can everyone be happy?

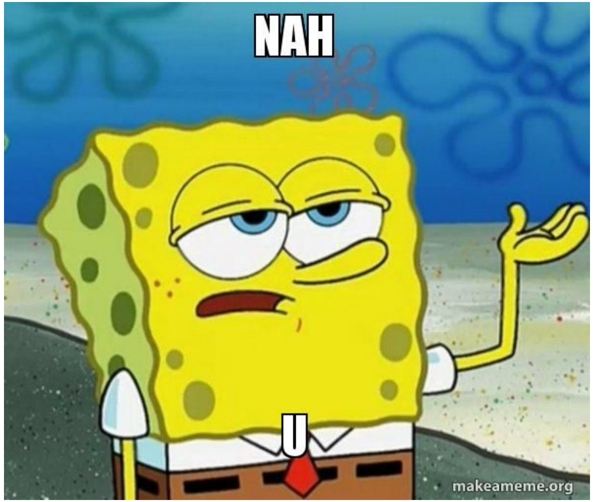

Can everyone be happy?

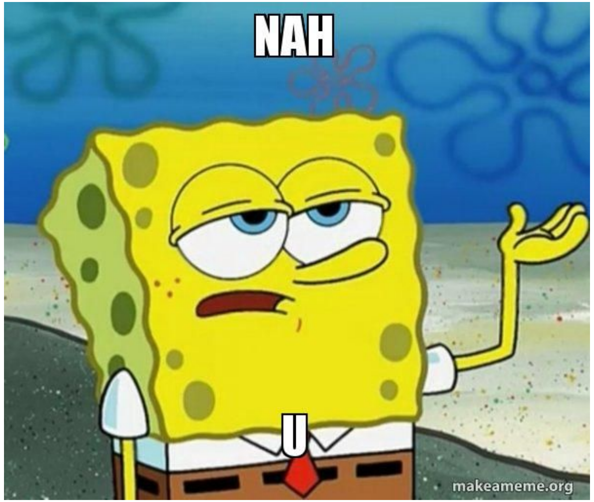

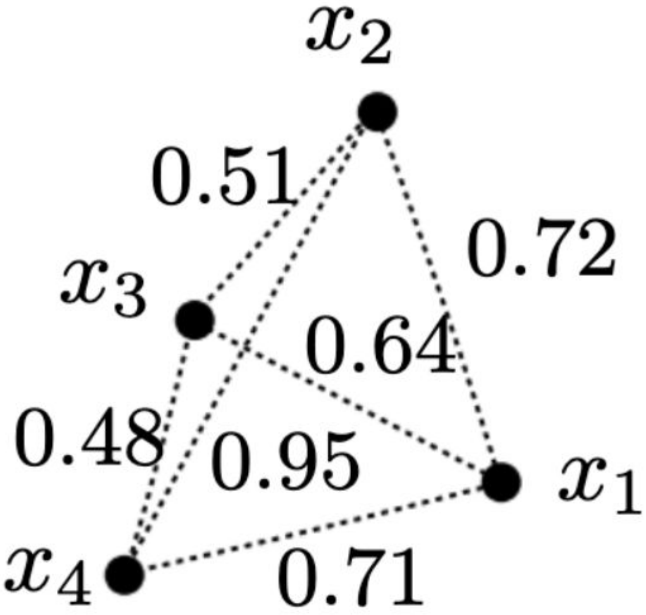

Can everyone be happy?

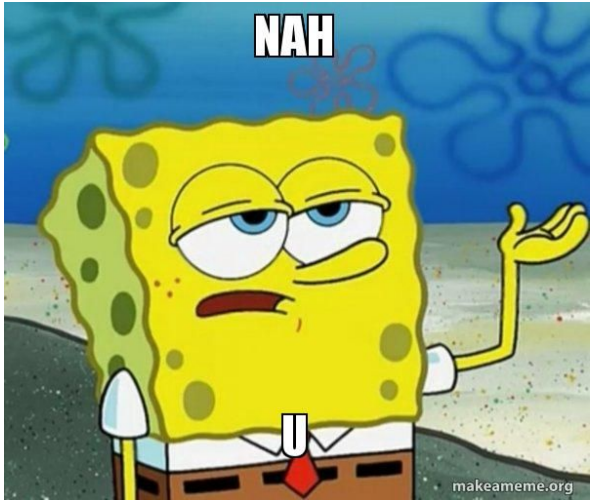

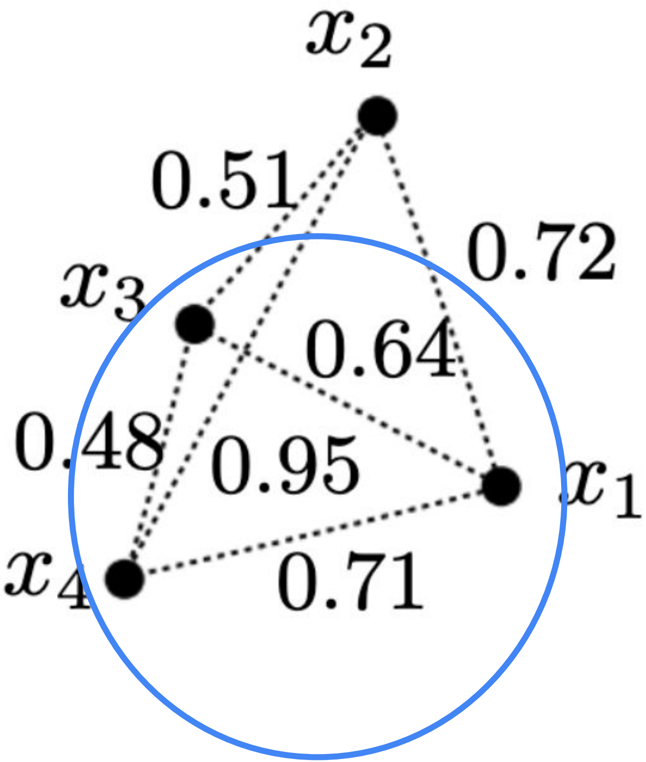

Can everyone be happy?

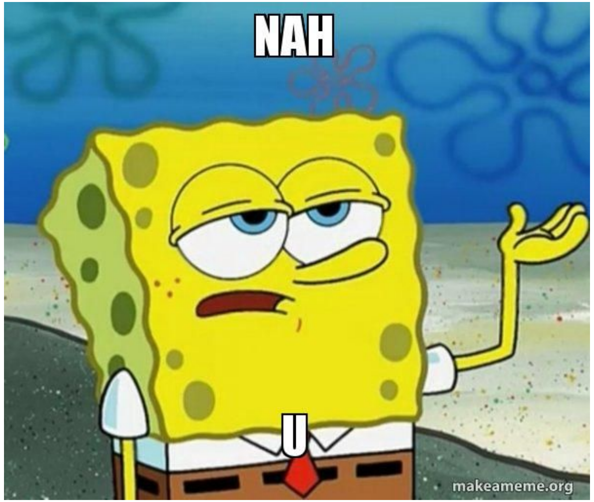

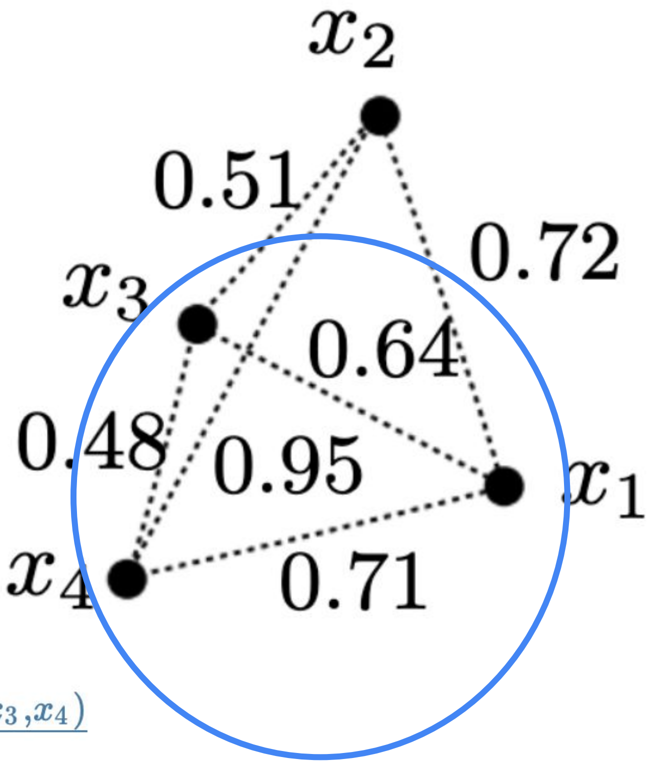

$d(x_2, x_3) = 0.51 < 0.56 = \frac{d(x_3, x_1) + d(x_3, x_4)}{2}$

Can everyone be happy?

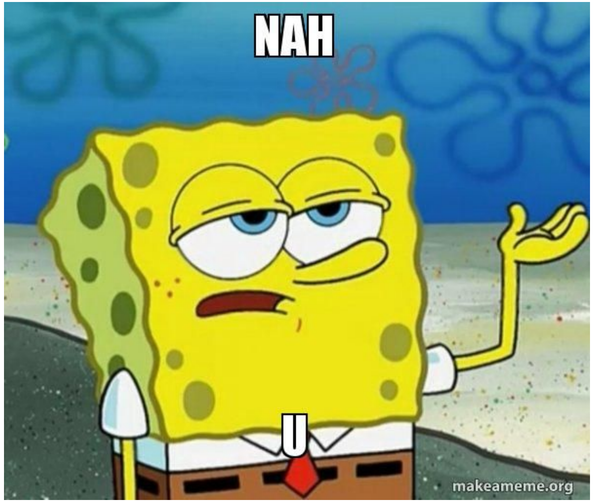

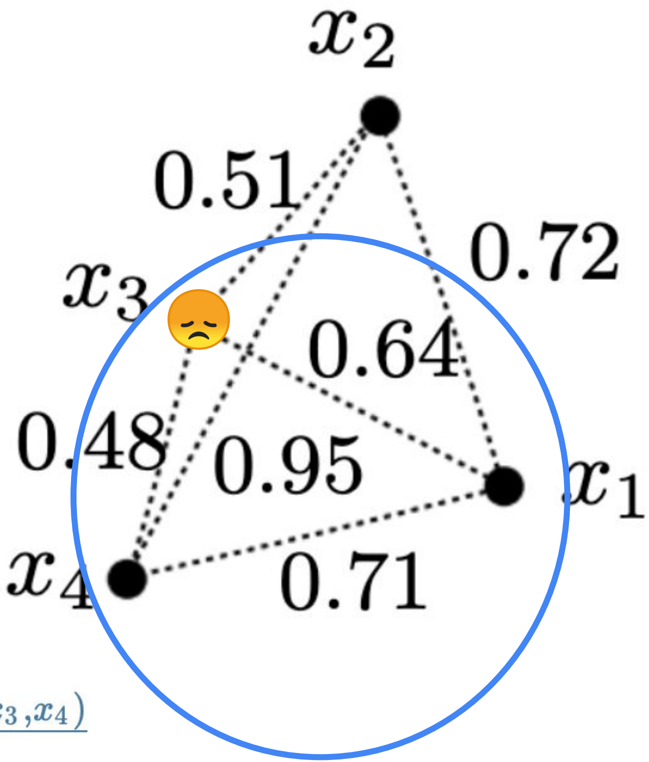

$$d(x_2, x_3) = 0.51 < 0.56 = \frac{d(x_3, x_1) + d(x_3, x_4)}{2}$$

Deciding if a happy clustering exist?

Deciding if a happy clustering exist?

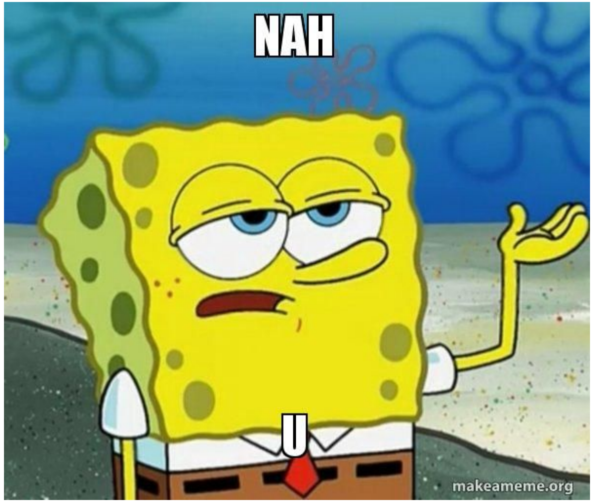

Deciding if a happy clustering exist?

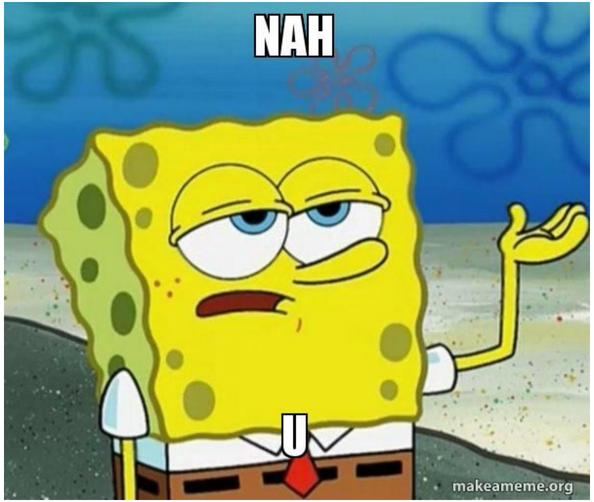

3-SAT $\rightarrow$ IP-Stability

1-dimension

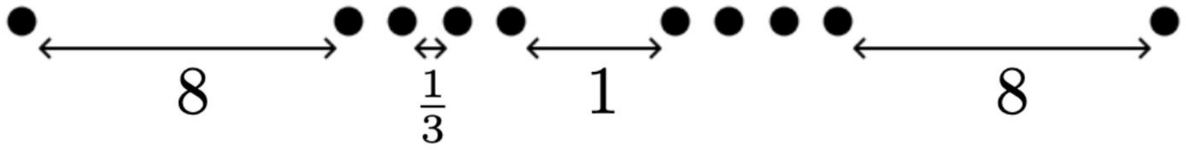

1-dimension

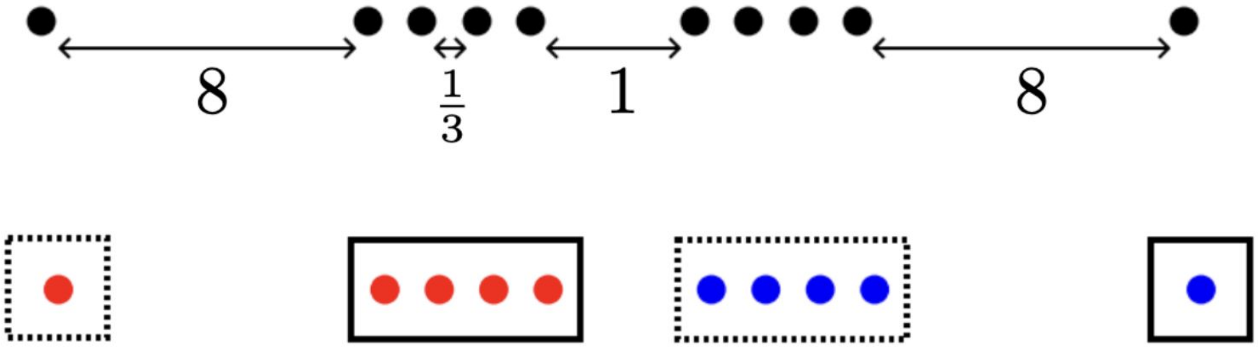

1-dimension

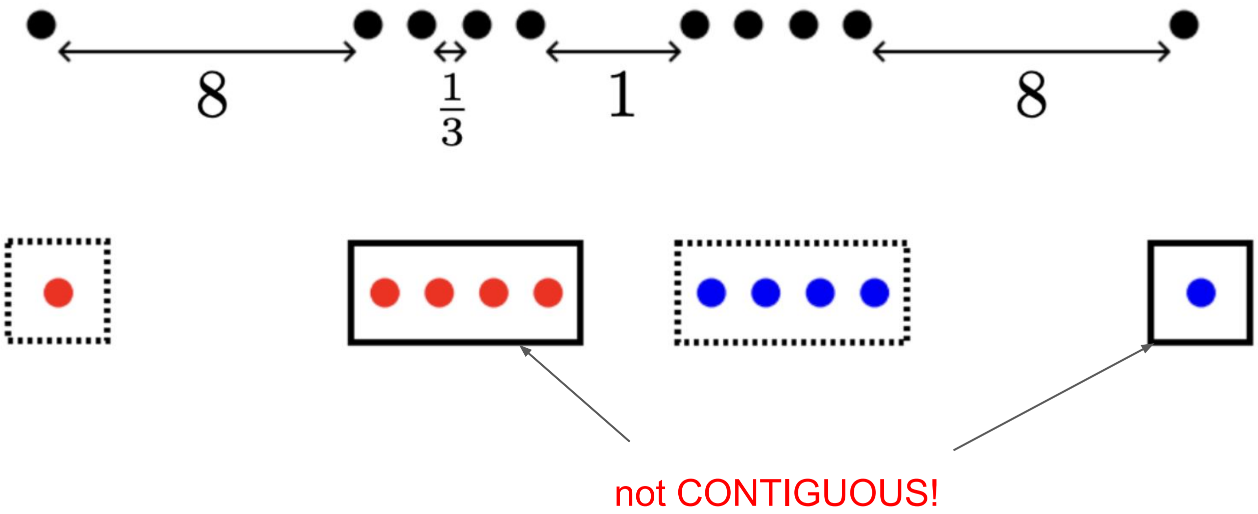

1-dimension

1-dimension

1-dimension

1-dimension

1-dimension

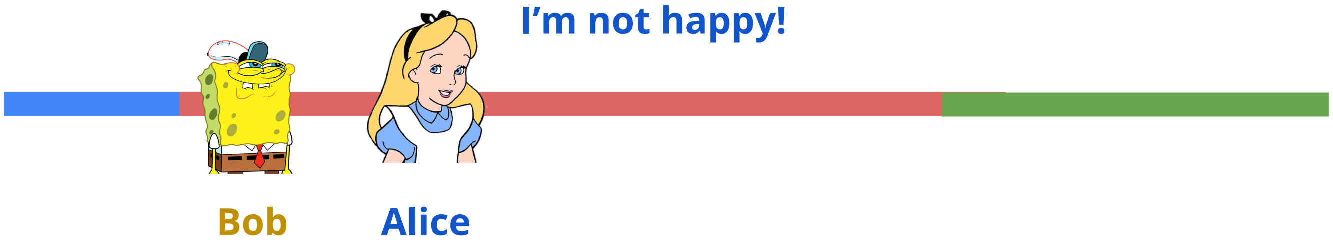

1-dimension

Alice is not happy implies that Bob is also not happy!

1-dimension

Alice is not happy implies that Bob is also not happy!

We want folks at the **border** to be **happy**

1-dimension

1-dimension

1-dimension

1-dimension

1-dimension

1-dimension

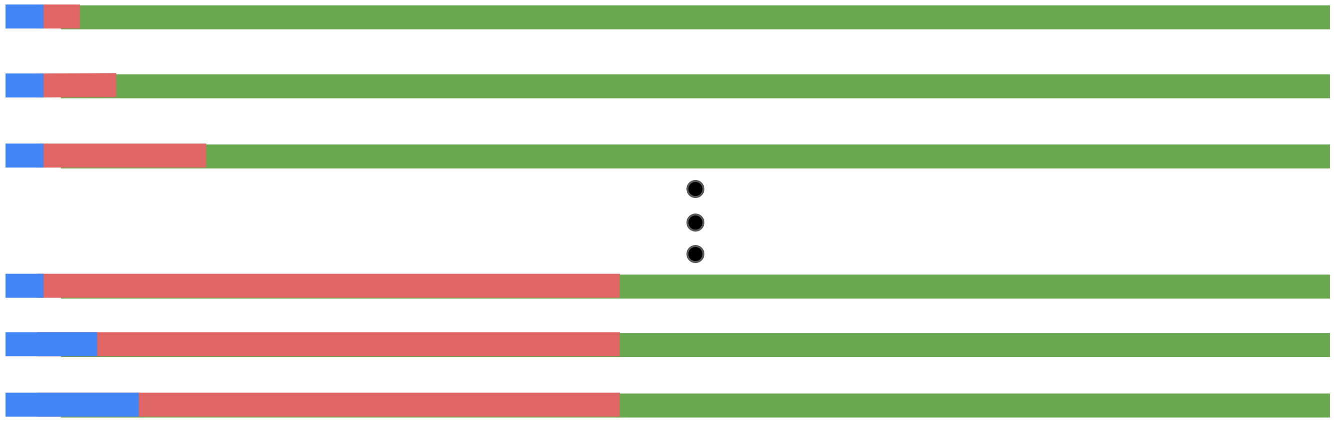

1-dimension

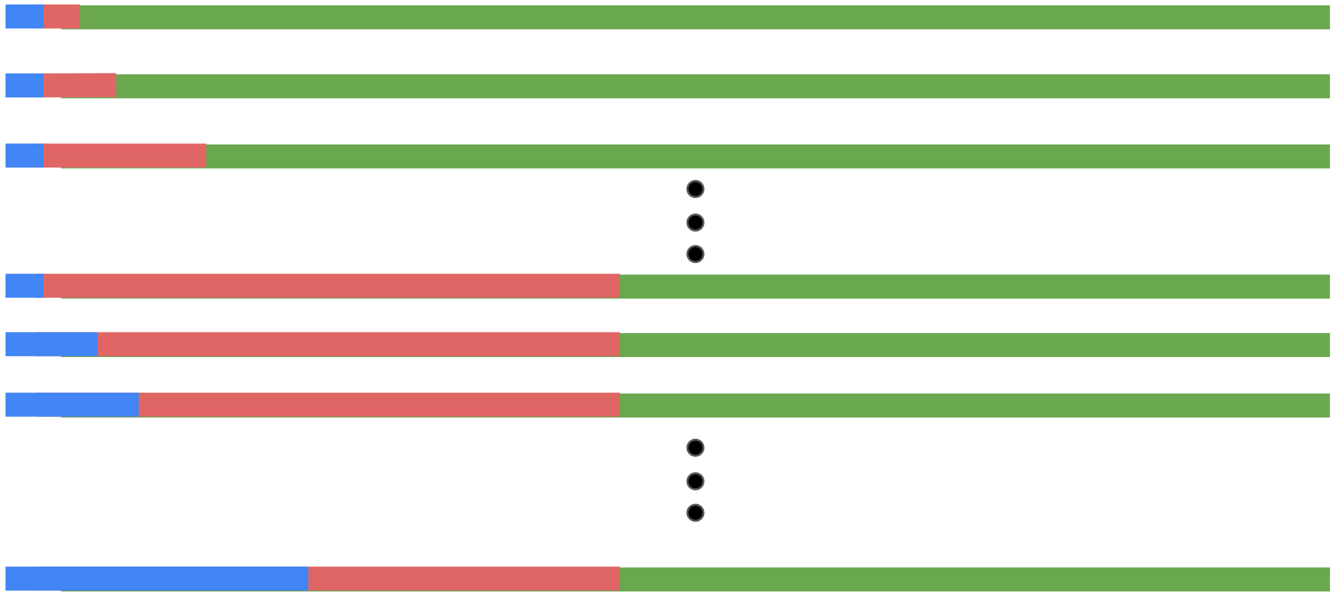

1-dimension

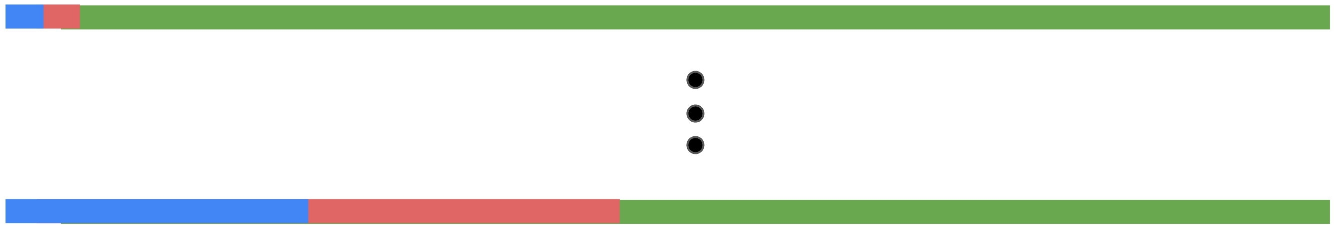

Finding a happy clustering can be done in $\mathcal{O}(kn)$

General Metric Space

General Metric Space

- Distances respect triangle inequality

General Metric Space

- Distances respect triangle inequality

$$d(x,z) \leq d(x,y) + d(y,z)$$

General Metric Space

- Distances respect triangle inequality
  - e.g., Euclidean space

$$d(x, z) \leq d(x, y) + d(y, z)$$

General Metric Space

- Distances respect triangle inequality
  - e.g., Euclidean space
- Finding a happy clustering is hard

$$d(x, z) \leq d(x, y) + d(y, z)$$

General Metric Space

- Distances respect triangle inequality
  - e.g., Euclidean space
- Finding a happy clustering is hard
- Need to relax the definition!

$$d(x, z) \leq d(x, y) + d(y, z)$$

x is **almost** happy if x's cluster is not too far away $d(x, C(x)) \leq t \cdot d(x, C')$

General Metric Space

- Distances respect triangle inequality
  - e.g., Euclidean space
- Finding a happy clustering is hard
- Need to relax the definition!

$$d(x, z) \leq d(x, y) + d(y, z)$$

x is **almost** happy if x's cluster is not too far away $d(x, C(x)) \leq t \cdot d(x, C')$

"t-approximately IP-stable"

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \le \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \le \mathcal{O}(\log n)d(x,y)$$

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

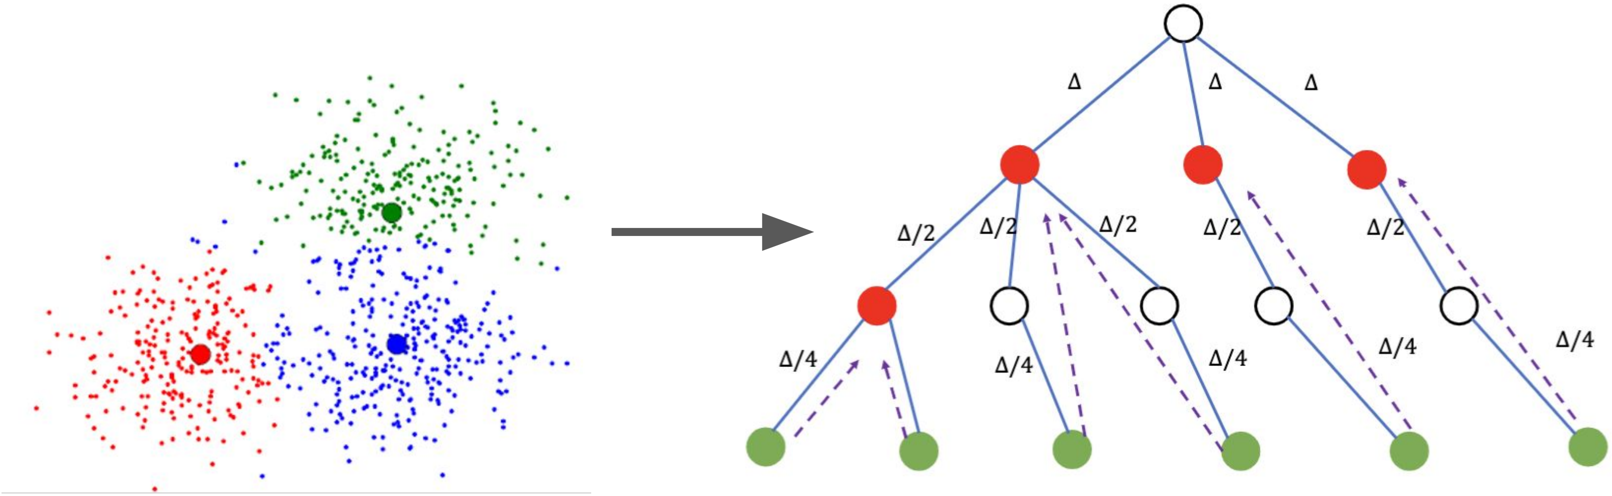

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T \sim \mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n) d(x,y)$$

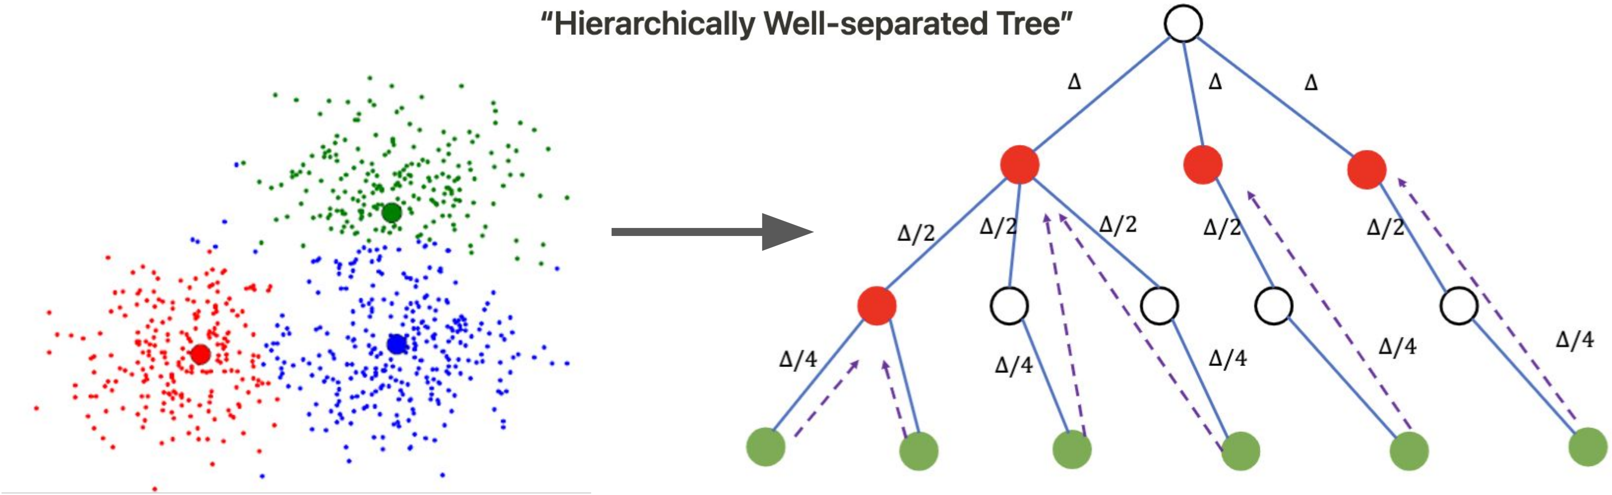

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

"Hierarchically Well-separated Tree"

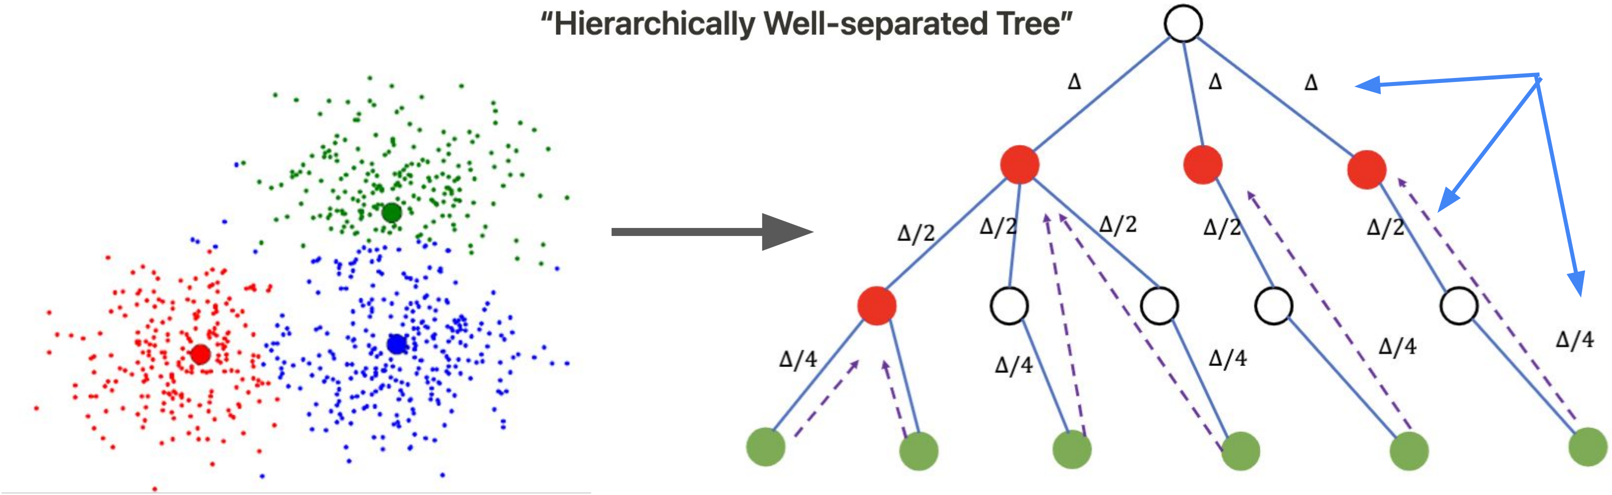

General Metric: Algorithm

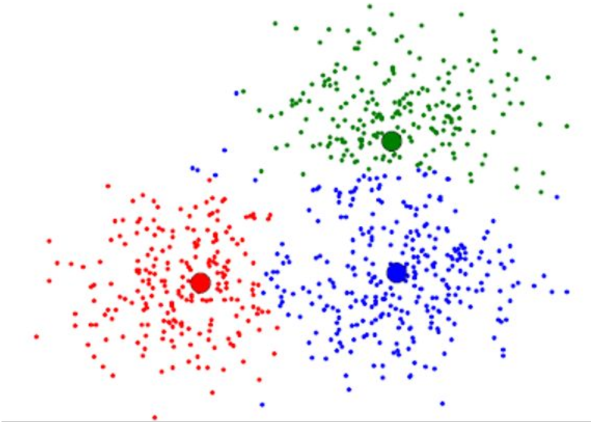

General Metric: Algorithm

- Embed the input space

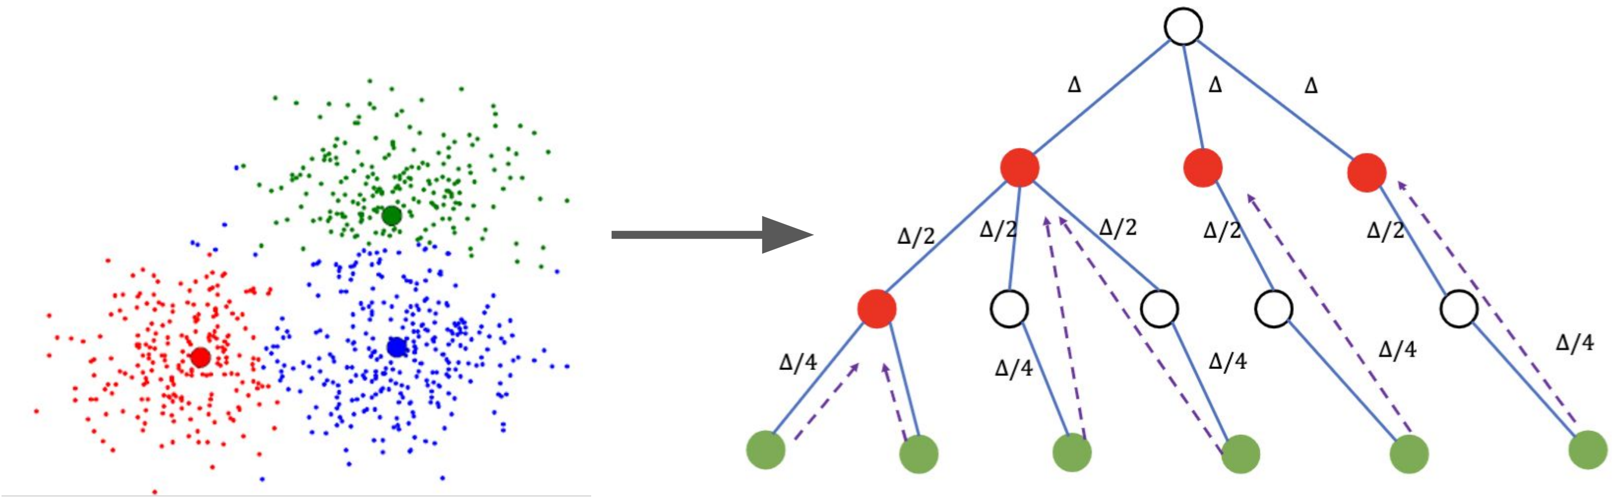

General Metric: Algorithm

- Embed the input space
- Find a clustering from the tree

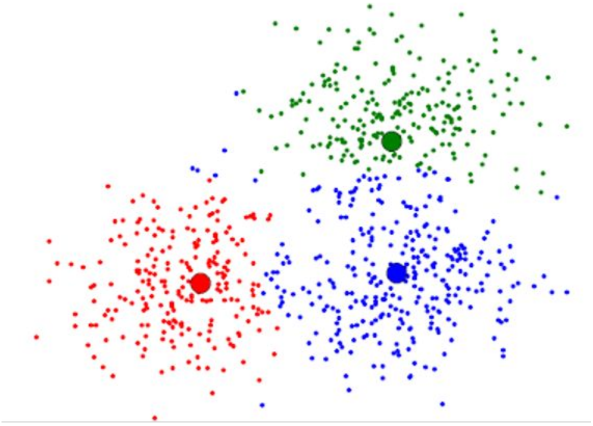

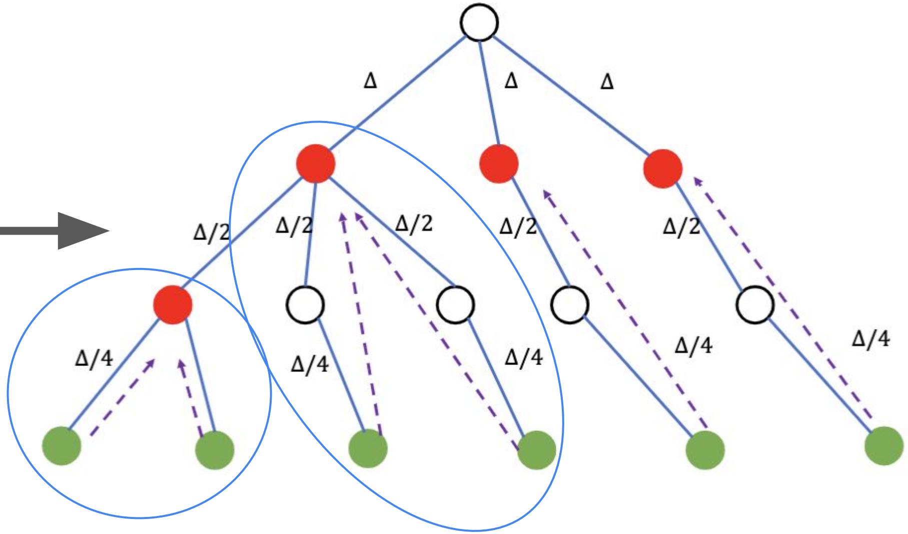

General Metric: Algorithm

- Embed the input space
- Find a clustering from the tree
- Output the corresponding clustering!

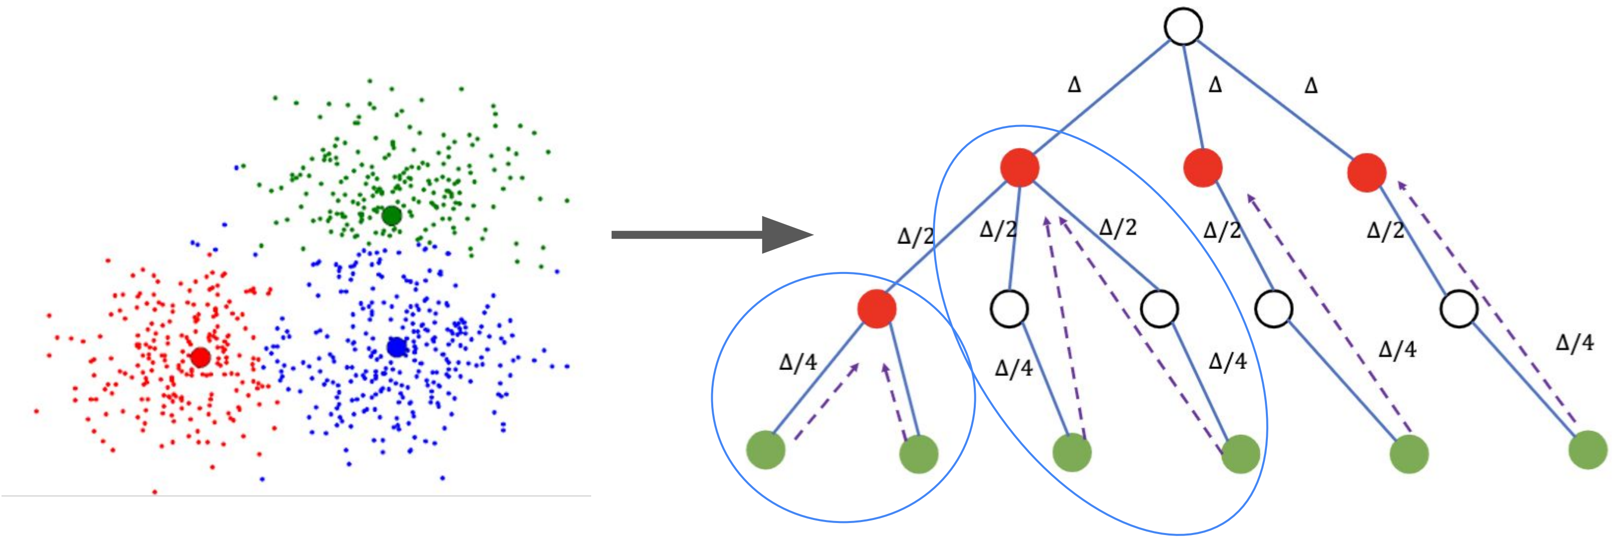

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

The expectation is over a distribution of trees :(

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n) d(x,y)$$

- In STOC'21 Haeupler, Hershkowitz, and Zuzic gave a worst-case bound
    - But ϵ points are excluded

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

- In STOC'21 Haeupler, Hershkowitz, and Zuzic gave a worst-case bound
  - But ϵ points are excluded

$$d(x,y) \leq d_T(x,y) \leq \mathcal{O}(\frac{\log^2 n}{\epsilon})d(x,y)$$

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

- In STOC'21 Haeupler, Hershkowitz, and Zuzic gave a worst-case bound
  - But ϵ points are excluded

$$d(x,y) \leq d_T(x,y) \leq \mathcal{O}(\frac{\log^2 n}{\epsilon})d(x,y)$$

We can make 99% of people almost happy!

Tree Embedding

- From **a metric space** to **a tree**
- FRT Tree (Fakcharoenphol Rao Talwar, 04)

$$d(x,y) \leq \mathbb{E}_{T\sim\mathcal{T}}\left[d_T(x,y)\right] \leq \mathcal{O}(\log n)d(x,y)$$

- In STOC'21 Haeupler, Hershkowitz, and Zuzic gave a worst-case bound
    - But ϵ points are excluded

$$d(x,y) \leq d_T(x,y) \leq \mathcal{O}(\frac{\log^2 n}{\epsilon})d(x,y)$$

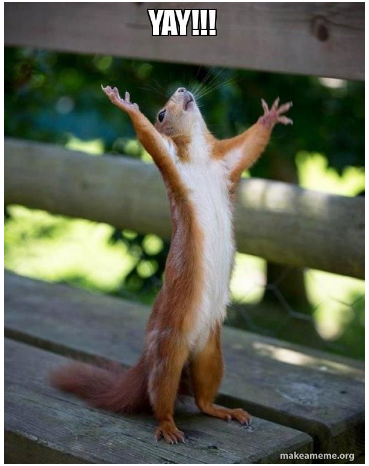

We can make 99% of people almost happy!

Experiments

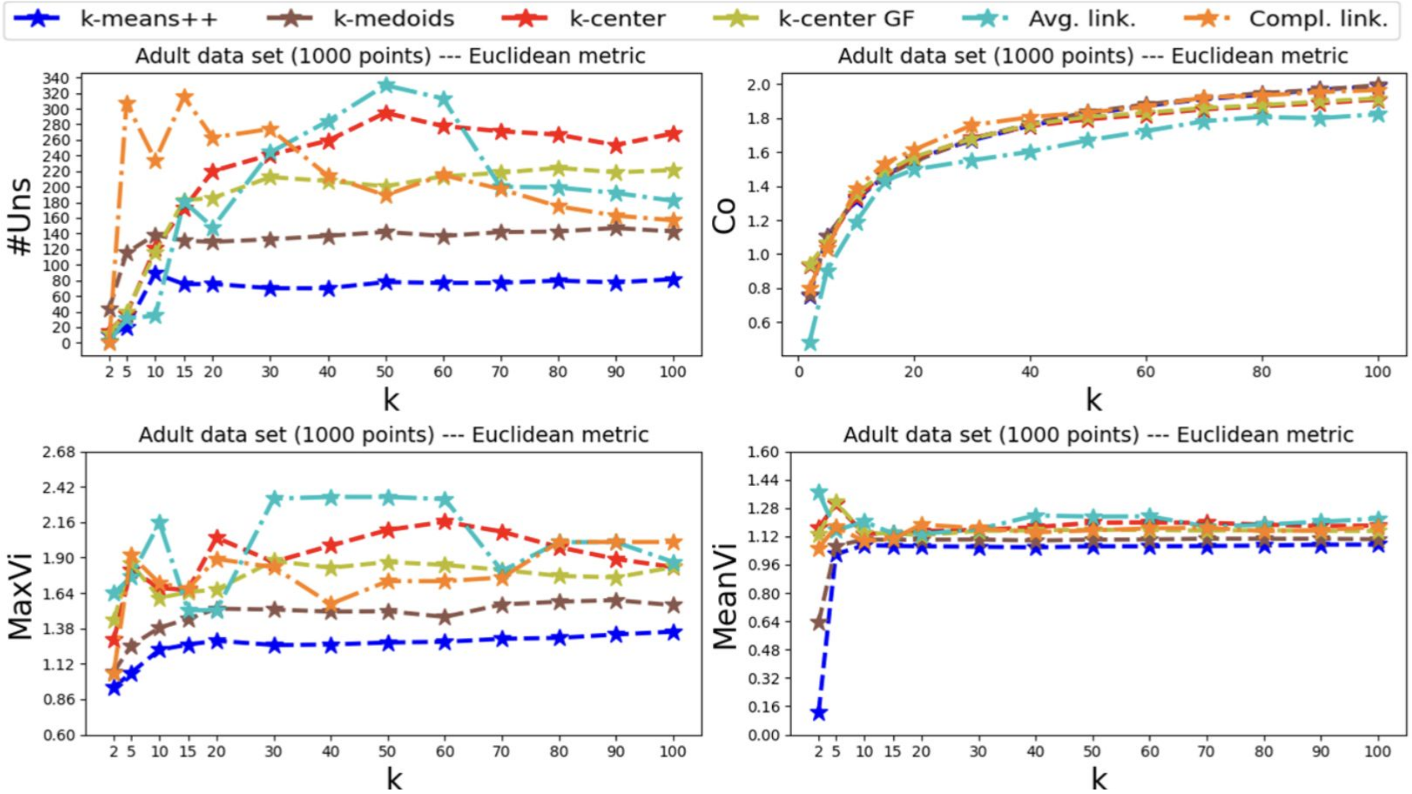

Summary

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space

Summary

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness

Summary

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

Summary

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

weighted instance?

k > 2?

better guarantee?

Summary

- Positive Results
  - 1D
  - Trees for 2 clusters
  - General Metric Space
- Negative Results
  - Impossibility
  - NP-hardness
- Experiments

weighted instance?

k > 2?

better guarantee?

hardness of approximation?

Summary

- Positive Results
  - 1D — weighted instance?
  - Trees for 2 clusters — $k > 2$?
  - General Metric Space — better guarantee?
- Negative Results
  - Impossibility — hardness of approximation?
  - NP-hardness
- Experiments — guarantee for k-means++?

Thank you!