5339 - Algorithms design under a geometric lens
Spring 2014, CSE, OSU
Lecture 3: Random embeddings

Instructor: Anastasios Sidiropoulos

January 15, 2014

Limitations of embeddings

Embedding the n-cycle into a tree requires distortion $\Omega(n)$.

Limitations of embeddings

Embedding the n-cycle into a tree requires distortion $\Omega(n)$.

Can we embed the n-cycle in a *random* tree?

Random embeddings

Let (X, ρ) be a metric space.

Random embeddings

Let (X, ρ) be a metric space.
Let $\mathcal{M}$ be a family of metric spaces.

Random embeddings

Let (X, ρ) be a metric space.
Let $\mathcal{M}$ be a family of metric spaces.
A *random embedding* of (X, ρ) into $\mathcal{M}$ is a distribution $\mathcal{F}$ over pairs (f, M), where

- $M = (X', \rho')$ is a metric space in $\mathcal{M}$
- $f : X \to X'$
- For any $x, y \in X$, we have $\Pr[\rho'(f(x), f(y)) \geq \rho(x, y)] = 1$.
- For any $x, y \in X$, we have $\mathbf{E}[\rho'(f(x), f(y))] \leq \alpha \cdot \rho(x, y)$.

Random embeddings

Let (X, ρ) be a metric space.
Let $\mathcal{M}$ be a family of metric spaces.
A *random embedding* of (X, ρ) into $\mathcal{M}$ is a distribution $\mathcal{F}$ over pairs (f, M), where

- $M = (X', \rho')$ is a metric space in $\mathcal{M}$
- $f : X \to X'$
- For any $x, y \in X$, we have $\Pr[\rho'(f(x), f(y)) \geq \rho(x, y)] = 1$.
- For any $x, y \in X$, we have $\mathbf{E}[\rho'(f(x), f(y))] \leq \alpha \cdot \rho(x, y)$.

α : distortion

Examples

Random embedding of the $n \times n$ grid into a distribution over trees?

Random embeddings into trees

Theorem (Fakcharoenphol, Rao, Talwar '04)

Any n-point metric space admits a random embedding into a distribution over trees, with distortion $O(\log n)$.

From random partitions to random embeddings

Let (X, ρ) be a metric space.

From random partitions to random embeddings

Let (X, ρ) be a metric space.
Assume that for any $x, y \in X$, we have

$$1 \leq \rho(x, y) \leq \Delta.$$

From random partitions to random embeddings

Let (X, ρ) be a metric space.
Assume that for any $x, y \in X$, we have

$$1 \leq \rho(x, y) \leq \Delta.$$

For any $i \in \{0, \ldots, \log \Delta\}$, let $\mathcal{D}_i$ be a $(\beta, 2^i)$-Lipschitz partition of (X, ρ), for some $\beta = O(\log n)$.

From random partitions to random embeddings

Let (X, ρ) be a metric space.
Assume that for any $x, y \in X$, we have

$$1 \leq \rho(x, y) \leq \Delta.$$

For any $i \in \{0, \ldots, \log \Delta\}$, let $\mathcal{D}_i$ be a $(\beta, 2^i)$-Lipschitz partition of (X, ρ), for some $\beta = O(\log n)$.

For any i, sample a random partition $P_i \in \mathcal{D}_i$.

Building a hierarchical partition

Initially, all points are in the same "cluster".

Building a hierarchical partition

Initially, all points are in the same "cluster".

Consider all i from $\log \Delta$ to 0.

Building a hierarchical partition

Initially, all points are in the same "cluster".

Consider all i from $\log \Delta$ to 0.

For every current cluster C, *refine* C by intersecting it with P.

Building a hierarchical partition

Initially, all points are in the same "cluster".

Consider all i from $\log \Delta$ to 0.

For every current cluster C, *refine* C by intersecting it with P.

We obtain a family of partitions $C_{\log \Delta}, \ldots, C_0$, such that

- $C_{\log \Delta}$ contains a single cluster with all the points.
- C_i is a refinement of C_{i+1}.
- C_0 contains a singleton cluster for every point.

From hierarchical partitions to trees

Given $C_{\log \Delta}, \ldots, C_0$, we build a tree T as follows.

From hierarchical partitions to trees

Given $C_{\log \Delta}, \ldots, C_0$, we build a tree T as follows.

The root of T is the single cluster in $C_{\log \Delta}$.

From hierarchical partitions to trees

Given $C_{\log \Delta}, \dots, C_0$, we build a tree T as follows.

The root of T is the single cluster in $C_{\log \Delta}$.

The height-i nodes in T are the clusters in C_i.

From hierarchical partitions to trees

Given $C_{\log \Delta}, \ldots, C_0$, we build a tree T as follows.

The root of T is the single cluster in $C_{\log \Delta}$.

The height-i nodes in T are the clusters in C_i.

For any $i > 0$, every cluster A in C_i has as its children all the clusters A' in C_{i+1}, with $A' \subseteq A$.

From hierarchical partitions to trees

Given $C_{\log \Delta}, \ldots, C_0$, we build a tree T as follows.

The root of T is the single cluster in $C_{\log \Delta}$.

The height-i nodes in T are the clusters in C_i.

For any $i > 0$, every cluster A in C_i has as its children all the clusters A' in C_{i+1}, with $A' \subseteq A$.

The edges in T between a cluster A in C_i, and its children, have length 2^i.

The embedding

We map every point $x \in X$ to the leaf of T corresponding to the singleton cluster in C_0 containing x.

Distortion analysis

Let $x, y \in X$.

Distortion analysis

Let $x, y \in X$.

The probability that x and y are separated in P_i is at most

$$\Pr_{P_i \in \mathcal{D}_i}[P_i(x) \neq P_i(y)] \leq O(\log n) \cdot \frac{\rho(x, y)}{2^i}$$

Distortion analysis

Let $x, y \in X$.

The probability that x and y are separated in P_i is at most

$$\Pr_{P_i \in \mathcal{D}_i}[P_i(x) \neq P_i(y)] \leq O(\log n) \cdot \frac{\rho(x,y)}{2^i}$$

Let $\mathcal{E}_i$ be the random event that i is the maximum integer such that $P_i(x) \neq P_i(y)$.

Distortion analysis

Let $x, y \in X$.

The probability that x and y are separated in P_i is at most

$$\Pr_{P_i \in \mathcal{D}_i}[P_i(x) \neq P_i(y)] \leq O(\log n) \cdot \frac{\rho(x,y)}{2^i}$$

Let $\mathcal{E}_i$ be the random event that i is the maximum integer such that $P_i(x) \neq P_i(y)$.

Conditioned on $\mathcal{E}_i$, we have $d_T(f(x), f(y)) = O(2^i)$.

Distortion analysis (cont.)

We have

$$\mathbf{E}[d_T(f(x), f(y))] \leq \sum_{i=0}^{\log \Delta} \Pr[\mathcal{E}_i] \cdot O(2^i)$$
$$\leq \sum_{i=0}^{\log \Delta} O(\log n) \cdot \frac{\rho(x,y)}{2^i} \cdot O(2^i)$$
$$= O(\log n \cdot \log \Delta \cdot \rho(x,y))$$

Therefore, the distortion is $O(\log n \cdot \log \Delta)$.

Applications of random embeddings

Let V be a set, $\mathcal{I} \subset \mathbb{R}_+^{V \times V}$ a set of non-negative vectors corresponding all feasible solutions for a minimization problem, and $c \in \mathbb{R}_+^{V \times V}$.

Applications of random embeddings

Let V be a set, $\mathcal{I} \subset \mathbb{R}_+^{V \times V}$ a set of non-negative vectors corresponding all feasible solutions for a minimization problem, and $c \in \mathbb{R}_+^{V \times V}$.
In the *linear minimization problem* $(\mathcal{I}, c)$ we are given a graph G with vertex set V, and want to find some $s \in \mathcal{I}$, minimizing

$$\sum_{\{u,v\} \in V \times V} c_{u,v} \cdot s_{u,v} \cdot d(u,v)$$

Applications of random embeddings

Let V be a set, $\mathcal{I} \subset \mathbb{R}_+^{V \times V}$ a set of non-negative vectors corresponding all feasible solutions for a minimization problem, and $c \in \mathbb{R}_+^{V \times V}$.
In the *linear minimization problem* $(\mathcal{I}, c)$ we are given a graph G with vertex set V, and want to find some $s \in \mathcal{I}$, minimizing

$$\sum_{\{u,v\} \in V \times V} c_{u,v} \cdot s_{u,v} \cdot d(u,v)$$

Captures MST, TSP, Facility-Location, k-Server, Bi-Chromatic Matching, etc.

Applications (cont.)

Theorem

For any a linear minimization problem Π, if there exists a polynomial-time α-approximation algorithm for Π on trees, then there exists a randomized polynomial-time $O(\alpha \cdot \log n)$-approximation algorithm for Π on arbitrary graphs.

Applications (cont.)

Theorem

For any a linear minimization problem Π, if there exists a polynomial-time α-approximation algorithm for Π on trees, then there exists a randomized polynomial-time $O(\alpha \cdot \log n)$-approximation algorithm for Π on arbitrary graphs.

Proof.

Sampling a random embedding into a tree T with distortion $O(\log n)$, solve Π on T, and finally pull the solution back to the original graph G. The guarantee on the resulting approximation factor follows by the definition of distortion, and linearity of expectation. □