Massachusetts Institute of Technology — Lecturer: Michel X. Goemans
18.409: Topics in TCS: Embeddings of Finite Metric Spaces
September 6, 2006 — Scribe: Punyashloka Biswal

Lecture 1

- Instructor: Michel X. Goemans (`goemans@math.mit.edu`, x3-2668, 2-351).
- Website at `http://www-math.mit.edu/~goemans/18409.html`.

Definition 1 (Metric, Semimetric)**.** *An ordered pair* (X, d_X) *consisting of a set* X *and a map* $d_X : X \times X \to [0, \infty)$ *is said to be a* metric space *if* d_X *satisfies the following conditions:*

1. $d_X(x, y) = 0 \Leftrightarrow x = y$,
2. $d_X(x, y) = d_X(y, x)$ *for all* $x, y \in X$,
3. $d_X(x, y) + d_X(y, z) \leq d_X(x, z)$ *for all* $x, y, z \in X$.

If d_X *satisfies (2) and (3) but only* $d_X(x, x) = 0$ *for all* $x \in X$ *instead of (1), then it is called a* semimetric *(or* pseudometric*).*

In this course, we shall usually be concerned with *finite* metric spaces. In this case, we write $|X| = n$.

Normed Spaces. We shall be especially interested in *normed* metric spaces, which are defined by $d(x, y) = f(x - y)$ for some suitable norm f. For example, consider the norm ℓ_p^d ($p \geq 1$, d an integer) which is defined on $X = \mathbb{R}^d$ as

$$\|x\|_p = \Big(\sum_{i=1}^{d} |x_i|^p\Big)^{1/p}$$

and its associated metric $d_X(x, y) = \|x - y\|_p$. When $p = 2$, we recover the usual Euclidean metric. When $p = \infty$, we have $\|x\|_\infty = \max_i |x_i|$. For finite p, there is an infinite-dimensional analog of this metric, known as ℓ_p:

$$\|x\|_p = \Big(\sum_{i \geq 1} |x_i|^p\Big)^{1/p}.$$

Application: Sparsest Cut. Let $G = (V, E)$ be a finite graph. We wish to find the sparsest cut $S \subseteq G$, i.e.,

$$\arg \min_{\emptyset \neq S \subsetneq V} \frac{|\delta(S, V \setminus S)|}{|S|\,|V \setminus S|},$$

where $\delta(A, B) = \{(u, v) \in E \mid u \in A \text{ and } v \in B\}$. The continuous analog of this notion yields an isoperimetric inequality, which is fundamental in geometry.

While this problem is known to be NP-hard, techniques based on metric embeddings let us find approximations to the sparsest cut. In particular, we have the following theorem (due to Bourgain):

Theorem 1 (Bourgain)**.** *If* $|X| = n$*, then* X *can be embedded in* ℓ_2 *with distortion* $O(\log n)$*.*

This leads automatically to a $O(\log n)$-approximation algorithm for the sparsest cut problem, and better embeddings for metrics of negative type (which we shall define later in the course) lead to further improvements.

Definition 2 (D-embedding, distortion, isometry). *Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f\colon X \to Y$ is said to be a D-*embedding *if there exists $r > 0$ such that for all $x, y \in X$,*

$$r d_X(x, y) \le d_Y(x, y) \le r D d_X(x, y).$$

(The r in the definition makes sure that a scaled copy of a metric is considered undistorted with respect to the original.) We define the distortion *of f to be the infimum of all D such that f is a D-embedding. If f has distortion 1, then X is said to be* isometrically embeddedable *in Y.*

Definition 3 (Lipschitz constant). *The Lipschitz constant of $f\colon X \to Y$ is defined as*

$$\|f\|_{Lip} = \sup_{x,y} \frac{d_Y(f(x), f(y))}{d_X(x, y)}$$

It can be shown that the distortion of f equals $\|f\|_{\mathrm{Lip}} \left\|f^{-1}\right\|_{\mathrm{Lip}}$, and mappings f with finite distortion are called *bi-Lipschitz*.

Theorem 2. *Let $c_p(X)$ be the infimum of the distortions of all embeddings of a space X into ℓ_p. If (X, d_X) is finite, then $c_\infty(X) = 1$. In words, a finite metric space can always be isometrically embedded into ℓ_∞.*

Proof. Define an embedding $f\colon X \to \mathbb{R}^{|X|}$ by

$$x \mapsto f(x) = (\ldots, \underbrace{d_X(x, z)}_{z\text{-th coordinate}}, \ldots).$$

(This is a special case of a Frechet embedding, which we shall use extensively in the remainder of this course.) Then $\|f(y) - f(x)\|_1 = \max_{z \in X} |d_X(y, z) - d_X(x, z)|$. By the triangle inequality, $|d_X(y, z) - d_X(x, z)| \le d_X(y, z)$ for each z, so that $\|f(y) - f(x)\|_1 \le d_X(y, z)$. But when $z = x$, we have $|d_X(y, z) - d_X(x, z)| = |d_X(x, y)| \le \max_{z \in X} |d_X(y, z) - d_X(x, z)|$. The result follows. ☐

Note that the above construction used $p = \infty, d = |X|$. Reducing the dimension usually forces a larger distortion.

Lemma 3. *If (X, d_X) embeds isometrically into ℓ_2, then it embeds isometrically into $\ell_2^{|X|-1}$.*

Proof. Let f be the original embedding, and fix an arbitrary point $v \in X$. Then X embeds isometrically into the subspace spanned by the vectors $(f(x) - f(v))_{x \ne v}$, which has dimension at most $|X| - 1$. ☐

Lemma 4. *If (X, d_X) embeds isometrically into ℓ_1, then it embeds isometrically into ℓ_1^N, where $N = \binom{n}{2}$.*

Proof. Every semimetric on X can be represented completely as a N-dimensional vector $d = (d_{ij})_{i<j\le n}$ consisting of all the pairwise distances. Define CUT_n to be the set of all semimetrics that can be embedded isometrically in ℓ_1 (see Deza and Laurent, *Geometry of Cuts and Metrics*). CUT_n is called the *cut cone*. It is a *cone* because it has the following properties:

1. When $x \in \mathrm{CUT}_n$, then $\alpha x \in \mathrm{CUT}_n$ for all $\alpha \geq 0$.

2. If $x, y \in \mathrm{CUT}_n$, then $x + y \in \mathrm{CUT}_n$.

Consider the extreme rays of this cone. (The *extreme rays* of a cone C are vectors $x \in C$ such that if $x = y + z$ and $y, z \in C$, then $y = \alpha x$ and $z = \beta x$ for some $\alpha, \beta \in \mathbb{R}$.) The extreme rays of CUT_n must correspond to embeddings into ℓ_1^1, because otherwise they could be further decomposed into one-dimensional embeddings. In fact, we can say something stronger: every one-dimensional embedding can be decomposed into *cut metrics* of the form

$$d(x, y) = \begin{cases} 0 & x, y \in S \text{ or } x, y \notin S \\ \alpha & x \in S \text{ and } y \notin S. \end{cases}$$

(What we have just shown justifies why CUT_n is called the cut cone: any element in CUT_n (the ℓ_1-embeddable metrics) can be decomposed into a conic combination of cut metrics.) Now, we can use Caratheodory's theorem to conclude that every element $d \in \mathrm{CUT}_n$ can be written in the form $d = \sum_{i=1}^{N} d_i$ where the d_i are extreme rays and we have used the fact that there we do not need more extreme rays than the dimension of the space. Finally, use one dimension for each extreme ray to construct an ℓ_1^N-embedding for X. $\square$

The above result can be generalized to the case when $p \neq 1$. The proof carries through with minor modifications. Instead of CUT_n, one needs to consider the cone $\mathrm{NOR}_p = \{(d_{ij}^p) : d$ is a semi-metric embeddable into $l_p\}$.

Theorem 5. *Let $D_n(X, Y)$ be the distortion needed to embed a metric subspace of X of size n into Y. Then*

1. $D_n(\ell_2, \ell_1) = 1$.

2. $D_n(\ell_2, \ell_p) = 1$ for any $p > 1$.

Proof. Consider the case of l_1. Let X be an n-element subset of $\mathbb{R}^d$ equipped with the Euclidean metric d_X. Pick a random direction (a point r on the unit sphere S_{d-1}) and look at the projections $\tau(x) = r \cdot x$ for every $x \in X$. The quantity $\mathbf{E}[|\tau(x) - \tau(y)|]$ is linear in $x - y$ and spherically symmetric. Therefore, it must have the form $\alpha \|x - y\|_2$ for some constant α.

Take a dense approximation $\{r_i\}_{1 \leq i \leq M}$ to S_{d-1} and define the map $x \mapsto f(x) = (x \cdot r_1, \ldots, x \cdot r_M)$. Then $\alpha \|x - y\|_2 \leq \|f(x) - f(y)\|_1 \leq \alpha(1 + \epsilon) |x - y|_2$, and we can make ϵ as small as we want by taking a sufficiently dense approximation. By the previous theorem, we can always reduce our embedding to N dimensions. Let $\xi = (f(x))_{x \in X} \in \mathbb{R}^{Nn}$ be the image of all the points of X under this map. Because we are considering a finite metric space and we know that the embedded distances are within a small constant of the original ones, all the ξs must lie within a bounded region of $\mathbb{R}^{Nn}$ (we can always fix one of the $f(x)$'s to the origin). If we consider the values of ξ that arise from $\epsilon = 1, 1/2, 1/3, \ldots$, then they all lie within this compact region. By compactness, this sequence must have a convergent subsequence. Then the limit of this sequence is the desired isometric embedding.

This proof can also be generalized to the case $1 < p$, by considering $\mathbf{E}[|\tau(x) - \tau(y)|^p]$ and first showing that this is equal to $\alpha_p \|x - y\|_2^p$ for some $\alpha_p > 0$. $\square$