Metrics

A *metric space* $\mathsf{M} = (V, d)$ consists of a set of points V and a function $d : V \times V \to \mathbb{R}_{\geq 0}$ satisfying the properties:

- (Symmetry) $d(x, y) = d(y, x)$ for all $x, y \in V$.
- (Triangle inequality) $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in V$.

The definition above is often called a *semi-metric*, and a metric space is also required to satisfy the property that $d(x, y) = 0 \iff x = y$. However, we will blur the distinction between semi-metrics and metrics.

Diameter. The *diameter* of a metric is $\max_{x,y \in V} d(x, y)$.

Ball. The *ball* $\mathbf{B}(x, r) := \{y \in V \mid d(x, y) \leq r$. The *open ball* $\mathbf{B}^\circ(x, r) := \{y \in V \mid d(x, y) < r$.

r-net. A set of points $N \subseteq V$ which is:

- (r-packing) $d(x, y) \geq r$ for all $x, y \in N$, and
- (r-covering) for $x \in V$, there exists $y \in N$ such that $d(x, y) \leq r$.

One can build an r-net using a simple greedy algorithm.

Distortion

Distortion. Given metrics $\mathsf{M} = (V, d)$ and $\mathsf{M}' = (V', d')$, and a map $f : \mathsf{M} \to \mathsf{M}'$,

- the *expansion* of f is $\max_{x,y \in V} \frac{d'(f(x), f(y))}{d(x,y)}$.
- the *contraction* of f is $\max_{x,y \in V} \frac{d(x,y)}{d'(f(x), f(y))}$.
- the *distortion* of f is

$$\textit{expansion} \times \textit{contraction} = \max_{x,y \in V} \frac{d'(f(x), f(y))}{d(x, y)} \times \max_{x,y \in V} \frac{d(x, y)}{d'(f(x), f(y))}.$$

If $\textit{distortion}(f) \leq D$, we write this as $\mathsf{M} \stackrel{D}{\hookrightarrow} \mathsf{M}'$. When we write $\mathsf{M}_1 \stackrel{\geq D}{\hookrightarrow} \mathsf{M}_2$, this is a *lower bound* statement: *every map* $f : V \to V'$ has distortion at least D.

This naturally extends to the case when $\mathcal{G}$ is a *family of metrics or graphs*, then $\mathsf{M} \stackrel{D}{\hookrightarrow} \mathcal{G}$ implies that *there exists* $\mathsf{M}' \in \mathcal{G}$ such that $\mathsf{M} \stackrel{D}{\hookrightarrow} \mathsf{M}'$; similarly, $\mathsf{M} \stackrel{\geq D}{\hookrightarrow} \mathcal{G}$ implies that *for all* $\mathsf{M}' \in \mathcal{G}$, it holds that $\mathsf{M} \stackrel{D}{\hookrightarrow} \mathsf{M}'$. If $\mathsf{M} \stackrel{1}{\hookrightarrow} \mathsf{M}'$, then we say that M *isometrically embeds* into M'; or just that M *embeds into* M'.

Metric Families

ℓ_p Spaces. For $1 \leq p < \infty$, the metric space ℓ_p consists of all infinite sequences $x = (x_i)_{i \leq 0}$ in $\mathbb{R}^{\mathbb{N}}$ for which $\sum_i |x_i|^p$ is finite; the distance is given by $|x - y|_p := (\sum_i |x_i - y_i|^p)^{1/p}$. The space ℓ_∞ is the set of bounded infinite sequences x, with the distance $|x - y|_\infty = \max_i(|x_i - y_i|)$. Often we will deal with ℓ_p^m for some finite m, when these sequences x just represent points in m-dimensional space $\mathbb{R}^m$; the distances are defined in the same way as above.

We say a metric M is an ℓ_p-metric (or it belongs to ℓ_p) if there is an isometric embedding of M into ℓ_p.

Tree-Metric. A metric $\mathsf{M} = (V, d)$ is a *tree metric* if there exists a tree $T = (V \cup S, E)$ with edge-weights, such that the shortest-path distance in T according to these edge-weights (denoted by d_T) agrees with d on all pairs in $V \times V$—in other words, $d_T(x, y) = d(x, y)$ for all $x, y \in V$.

Given a class $\mathcal{G}$ of graphs, one can define a $\mathcal{G}$*-metric* in the same way as above. E.g., we will often talk about *planar graph metrics.*

k**-HST.** A k*-Hierarchical well-Separated Tree* is rooted (weighted) tree with the following properties: (a) it is a balanced tree—all the leaves are at the same depth, (b) given any node x in the tree, all the children edges of x have the same length l_x, and the length of the edge from x to its parent node p_x (if any) has length $k \times l_x$. Hence, if the length of the root's children edges is L, and the height of the tree is h, then the edge lengths on any root-leaf path are $(L, L/k, L/k^2, \ldots, L/k^{h-1})$.

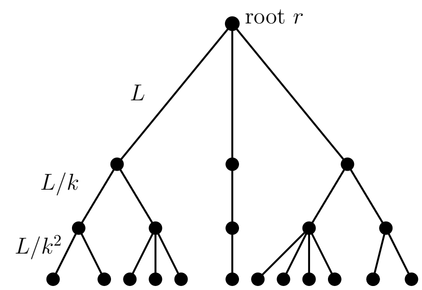

Figure 1: A k-HST with height 3.

Distributions over trees. Given a metric space $\mathsf{M} = (V, \delta)$ on $|V| = n$ points, let $\mathcal{T}$ be the set of trees $T = (V, E_T)$ on the with vertex set V with edge lengths $\ell : E_T \to \mathbb{R}$ such that each edge $e = \{u, v\} \in E_T$ has length $\ell(\{u, v\}) \geq \delta(u, v)$—i.e., trees whose distances *dominate* those in M.

A probability distribution $\mathcal{D}$ on this set of "dominating" trees $\mathcal{T}$ is said to α*-approximate* the metric $\mathcal{M}$ if for every $u, v \in V$,

$$\mathbf{E}_{T \leftarrow \mathcal{D}}[d_T(u, v)] \leq \alpha \cdot \delta(u, v). \tag{1}$$

I.e., for any two points, the expected distance in a random tree (drawn from this distribution) is at most α times what it was in (V, δ).

Other Useful Definitions

Padded Decomposition. A metric $\mathsf{M} = (V, d)$ is said to admit an α*-padded decomposition* if there exists a randomized procedure that takes as input a parameter $\Delta > 0$, and outputs a (random) partition $V_1, V_2, \ldots, V_k$ of the set V with the following properties:

- each set V_i has diameter at most Δ,
- for any $\rho > 0$, the probability $\Pr[\mathbf{B}(x, \rho) \text{ split by partitioning}] \leq \alpha \cdot \frac{\rho}{\Delta}$.

Note that this probability is taken over the randomness of the padded decomposition procedure. (The ball $\mathbf{B}(x, \rho)$ is *split by* the partitioning if it is not contained within any single "cluster" V_i.)

Tree Cover. Given a metric $\mathsf{M} = (V, d)$, an (α, k)-tree cover is a collection of trees $\mathcal{T} = \{T_1, T_2, \ldots, T_k\}$ such that for any pair of nodes $x, y \in V$, there exists a tree $T_j \in \mathcal{T}$ with

$$d(x, y) \le d_{T_j}(x, y) \le \alpha \cdot d(x, y).$$

Neighborhood Cover. Given a metric $\mathsf{M} = (V, d)$, an (α, r, t)-neigborhood cover is a collection $\mathcal{S} = \{S_1, S_2, \ldots\}$ of subsets $S_i \subseteq V$ of points such that (a) for each point $x \in V$, there is a subset S_j that contains the r-ball $B(x, r) = \{x' \in V \mid d(x, x') \le r\}$, (b) each point $x \in V$ is contained in at most t of the subsets in $\mathcal{S}$, and (c) each subset S_i has diameter at most $O(\alpha r)$.

Graphs

Outerplanar Graphs. These are planar graph such that there exists a face containing all the vertices; often this face is drawn as the outer face, hence the name. Equivalently, these are the graphs that exclude $K_{2,3}$ and K_4 as minors.

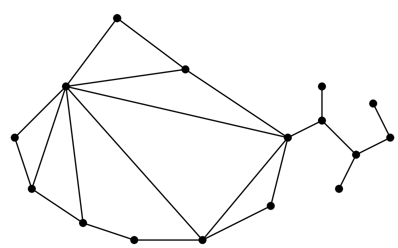

Figure 2: An outerplanar graph.

Series-Parallel Graphs. An (s, t)-series-parallel graph G is either (a) a single edge (s, t), or (b) the graph obtained by taking an (s_1, t_1)-series-parallel graph and an (s_2, t_2)-series-parallel graph and identifying $s_1 = s_2 = s$ and $t_1 = t_2 = t$ (this is called a *parallel composition*, or (c) the graph obtained by taking an (s_1, t_1)-series-parallel graph and an (s_2, t_2)-series-parallel graph and identifying $t_1 = s_2$ and setting $s_1 = s$ and $t_2 = t$ (this is called a *series composition*. A series-parallel graph G is a graph that contains vertices s and t such that G is an (s, t)-series-parallel graph.

Equivalently, take any planar graph that excludes K_4 as a minor: each 2-node-connected component of this is a series-parallel graph.

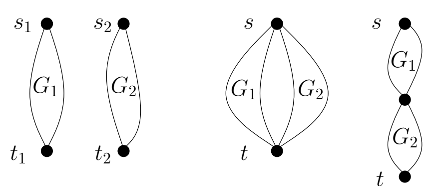

Figure 3: Parallel and Series compositions.

Expander Graphs. A (d, α)-*expander graph* on n vertices is a d-regular graph $G_n = (V_n, E_n)$ such that for every set $S \subseteq V_n$ with $|S| \le n/2$, the number of edges in ∂S (i.e., with one endpoint in S and the other in $V \setminus S$ is at least $\alpha|S|$.

We are interested in families of graphs (for infinitely many values of n) where both the degree d and the "expansion parameter" α are constants (independent of the size n). In this case, we just refer to the graphs as constant-degree expander graphs.

One can show (by a probabilistic construction) that there exist constant degree expander graphs; explicit constructions are known as well. For more details, see the survey by Linial, Hoory and Wigderson.

Diamond Graphs. Let the graph G_1 be a single edge, and for each $i \geq 1$, let G_i be obtained by taking G_{i-1} and replacing each edge by an (s,t)-series-parallel graphs consisting of two paths of length 2 (see figure below).