Bounded Geometries, Fractals, and Low-Distortion Embeddings[1]

Anupam Gupta, Robert Krauthgamer and James R. Lee

February 27, 2007

presented by Brian McFee for CSE254

Problem Setup

- Given a tree metric, we'll use path partitions to embed the vertices into ℓ_p for $1 \leq p \leq \infty$.
- We'll bound the distortion of the embedding in terms of the doubling constant (independent of n).
- The main result: If a tree metric has a bounded doubling dimension, it can be embedded with *constant* distortion. This can be extended to a constant number of dimensions.

Notation

- Let $T = (V, E)$ be a rooted tree, and let ρ be the shortest-path metric induced by the edge weights of T. Let r be the root of the tree, and $P(x, y)$ is the shortest path in T from x to y. Assume all edge weights are at least 1.
- Let $f : V \to \mathbf{R}_+^d$ be the embedding function.
- Let $\chi : E \to \mathcal{C}$ be a coloring of the tree edges.
- A *monotone path* is a simple path from a vertex to one of its descendants. We will only consider colorings χ such that every color class forms a monotone path in T.

The Embedding

- Given a coloring χ, our embedding will assign a dimension to each color class $c \in \mathcal{C}$:

$$f(v)_c = \sum_{e \in P(r,v) : \chi(e) = c} w(e)$$

$$f(v) = (f(v)_c : \forall c \in \mathcal{C}).$$

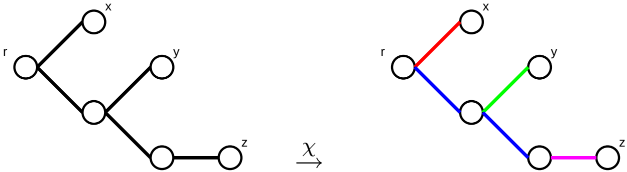

$f(r) = (0,0,0,0) \quad f(x) = (1,0,0,0)$

$f(y) = (0,1,1,0) \quad f(z) = (0,2,0,1)$

α-Good Colorings and Distortion Bound

- For $0 < \alpha \leq 1$, a coloring χ is called α-Good if for every vertex $v \in V$, and each ancestor u of v, there is some monochromatic segment of $P(u, v)$ with length $\geq \alpha\rho(u, v)$.
- High-level: $P(u, v)$ doesn't change colors too many times.

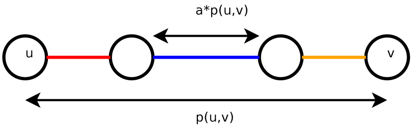

Lemma 1: A tree metric with an α-Good coloring can be embedded into $\ell_p^{|\mathcal{C}|}$ with distortion $\frac{2^{1-1/p}}{\alpha}$.

Ancestor Contraction

Consider the case where y is an ancestor of x in T. Then for all colors $c \in \mathcal{C}$, $f(x)_c \geq f(y)_c$. Then we can bound contraction as follows:

$$\begin{aligned} \|f(x) - f(y)\|_p^p &= \sum_{c \in \mathcal{C}} (f(x)_c - f(y)_c)^p \\ &\geq (\alpha\rho(x, y))^p \,, \end{aligned}$$

by α-Goodness of χ.

$$\Rightarrow \qquad \alpha\rho(x, y) \leq \|f(x) - f(y)\|_p .$$

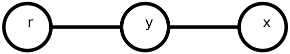

Ancestor Expansion

- Note that if y is an ancestor of x, the embedding cannot expand distances.
- If $p = 1$, then

$$\|f(x) - f(y)\|_1 = \sum_{e \in P(x,y)} w(e) = \rho(x, y).$$

- For $p > 1$,

$$\|f(x) - f(y)\|_p \leq \|f(x) - f(y)\|_1.$$

$$\Rightarrow \quad \|f(x) - f(y)\|_p \leq \rho(x, y).$$

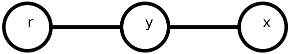

Non-Ancestor Expansion

- For a node y outside of x's ancestor set, and some $a \in V$,

$$\|f(x) - f(y)\|_p \leq \|f(x) - f(a)\|_p + \|f(a) - f(y)\|_p.$$

- Let a be the least common ancestor of x and y. Then

$$\begin{aligned}\|f(x) - f(y)\|_p &\leq \|f(x) - f(a)\|_p + \|f(a) - f(y)\|_p \\ &\leq \|f(x) - f(a)\|_1 + \|f(a) - f(y)\|_1 \\ &= \rho(x, a) + \rho(a, y) = \rho(x, y).\end{aligned}$$

$$\Rightarrow \qquad \|f(x) - f(y)\|_p \leq \rho(x, y).$$

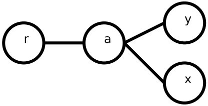

Non-Ancestor Contraction

To bound contraction between x and y, again let $a = LCA(x, y)$.

$$\begin{aligned}
\|f(x) - f(y)\|_p^p &= \|(f(x) - f(a)) - (f(y) - f(a))\|_p^p \\
&= \sum_{c \in P(a,x)} f'(x)_c^p + \sum_{c \in P(a,y)} f'(y)_c^p \\
&\geq \alpha^p(\rho(x,a)^p + \rho(a,y)^p) \\
&= 2\alpha^p \left(\frac{\rho(x,a)^p + \rho(a,y)^p}{2} \right) \\
&\geq 2\alpha^p \left(\frac{\rho(x,a) + \rho(a,y)}{2} \right)^p = \left(\frac{\alpha}{2^{1-1/p}} \rho(x,y) \right)^p
\end{aligned}$$

$$\Rightarrow \quad \frac{\alpha}{2^{1-1/p}} \rho(x,y) \leq \|f(x) - f(y)\|_p \leq \rho(x,y).$$

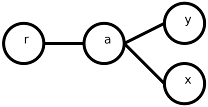

Okay, but so what?

- Given χ, we can compute α and find the distortion of the embedding.
- But how do we find χ? How do we know that α will be sufficiently big?

Doubling Dimension and ϵ-nets

We'll need some more definitions...

- $Y \subseteq X$ is an ϵ-net of X if the following hold:

$$\forall x \neq y \in Y, \rho(x, y) \geq \epsilon,$$

  and

$$\forall x \in X, \exists y \in Y : \rho(x, y) < \epsilon.$$

- The doubling dimension $\dim(X, \rho) = \log_2 \lambda$, where λ is the smallest integer such that every ball $B(x, 2\epsilon)$ can be covered by at most λ balls $B(x_i, \epsilon)$. λ is the doubling constant.

The Coloring Algorithm

$\chi = \text{Color}(T)$

1. set $k = \lceil \log_2 \text{diam}(T) \rceil$ and $Y_0 = \emptyset$.
2. for $i = 1, \ldots, k$ do
3. set $\epsilon_i = 2^{k-i}$ and let $Y_i \supseteq Y_{i-1}$ be an ϵ_i-net of T.
4. for each $y \in Y_i \setminus Y_{i-1}$ do
5. color all uncolored edges in $P(r, y)$ with a new color c_y.
6. color each remaining uncolored edge e with a new color c_e.

Coloring Example

$k = \lceil \log_2 \operatorname{diam}(T) \rceil = 3.$

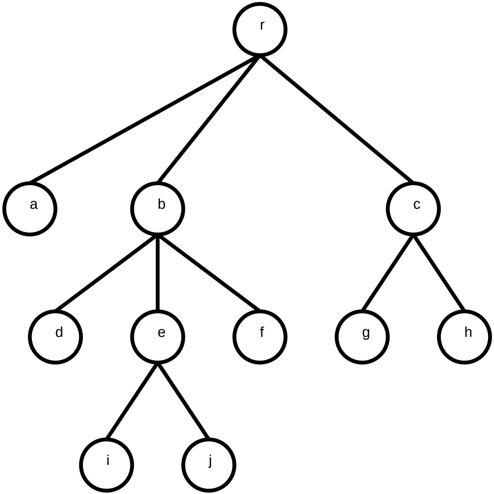

Coloring Example

$\epsilon_1 = 2^{3-1} = 4,\ Y_1 = \{b\}$

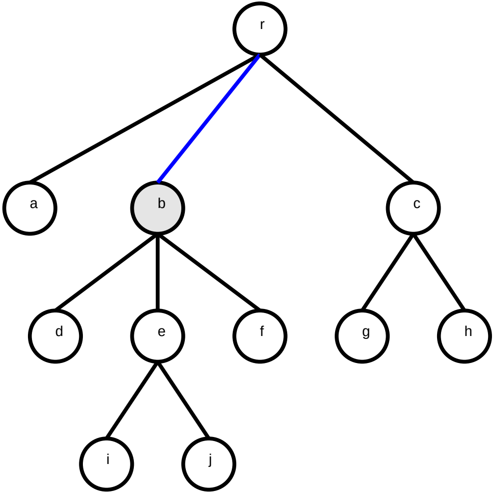

Coloring Example

$$\epsilon_2 = 2^{3-2} = 2,\ Y_2 = \{b, c, j\}$$

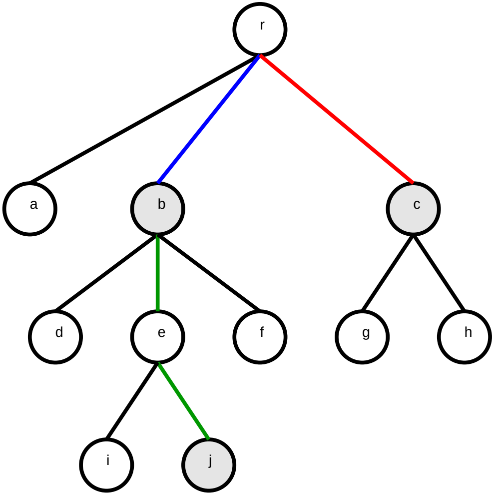

Coloring Example

$\epsilon_3 = 2^{3-3} = 1,\ Y_2 = \{b, c, j, a, e\}$

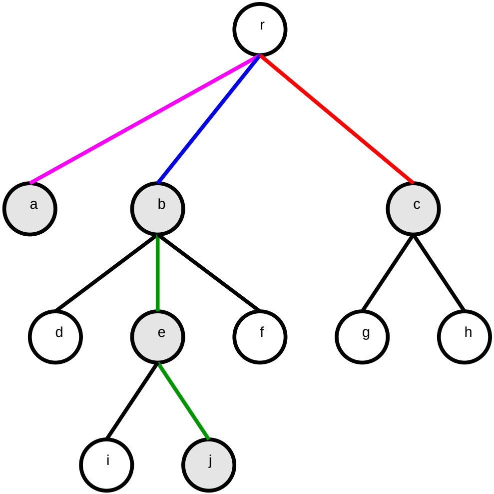

Coloring Example

And the rest...

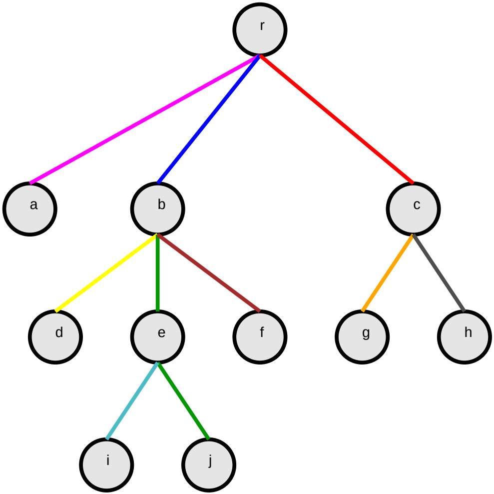

γ-Bad k-Combs

- Define a k-comb as a tree metric with the following properties:
  - The tree contains some vertex c called the center,
  - and the tree contains k edge-disjoint paths $P_1, \dots, P_k$ called hairs.
- A k-comb is γ-Bad if for some value $L \geq max_i \rho(c, P_i)$, the length of each hair is in $[L, \gamma L)$.
- High-level: each hair has more length than it has distance to the center, but not too much more.
- **Claim**: A tree with doubling constant $\lambda \leq k^{1/(\gamma+3)}$ does not contain a 2^γ-bad k-comb (proof on the discussion board).

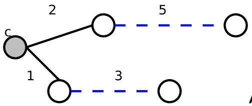

A 3-bad 2-comb

α-Goodness vs. γ-Badness

Theorem: Let $0 < \alpha \leq 1/60$. Then every tree metric has either an α-good coloring, or a 4-bad $(1/40\alpha)$-comb sub-metric.

α-Goodness or 4-bad $(1/40\alpha)$-comb

Proof Outline:

- Assume χ is not α-good. Then there is some vertex v with ancestor u such that no monochromatic portion of $P(u, v)$ has length more than $\alpha\rho(u, v)$.
- Let $u = w_0, \ldots, w_J = v$ be the color transition vertices of $P(u, v)$ according to χ.
- High-level: We'll build the comb with u as the center, and use paths from some w_i to ϵ-net points for hairs.

4-bad $(1/40\alpha)$-comb: $y_j \in Y_m$

- Let $P(w_{j-1}, w_j)$ be a maximal monochromatic path in $P(u, v)$, and let y_j be the net point that colored it. Let $D = \rho(u, v)$.
- Let m be the step in the algorithm where $D/4 < \epsilon_m \leq D/2$. If $\rho(w_j, v) \geq D/2$, then $y_j \in Y_m$:

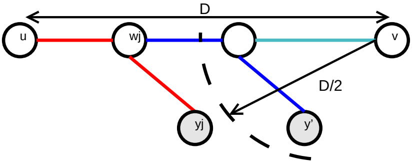

- Since Y_m is an $\epsilon_m \leq D/2$-net, $B(v, D/2)$ contains some net point y' which must be a descendant of w_j.
- Then $P(w_{j-1}, w_j)$ must already be colored before y' was added, so $y_j \in Y_m$. Otherwise would violate the assumption that $P(w_{j-1}, w_j)$ was maximal: the blue edges would be red.

4-bad $(1/40\alpha)$-comb: Hairs are long enough

- Consider each w_j where $\rho(u, w_j) \leq D/10$. Then $\rho(w_j, v) > D/2$, so each corresponding $y_j \in Y_m$.
- Since Y_m is an ϵ_m-net, $\rho(y_j, y_{j+1}) \geq \epsilon_m > D/4$. Expanding the path out,

$$\begin{aligned}
D/4 < \rho(y_j, y_{j+1}) &= \rho(y_j, w_j) + \rho(w_j, w_{j+1}) + \rho(w_{j+1}, y_{j+1}) \\
&< \rho(y_j, w_j) + \alpha D + \rho(w_{j+1}, y_{j+1}) \\
&< \rho(y_j, w_j) + \frac{D}{60} + \rho(w_{j+1}, y_{j+1}) \\
\frac{14}{60} D &< \rho(y_j, w_j) + \rho(w_{j+1}, y_{j+1}).
\end{aligned}$$

- So either $\rho(w_j, y_j) \geq D/10$ or $\rho(w_{j+1}, y_{j+1}) \geq D/10$.

α-Goodness vs. γ-Badness: The Comb

- For each j, one of $P(w_j, y_j)$ or $P(w_{j+1}, y_{j+1})$ has length at least $D/10$: these will be our hairs.
- Each hair is distinctly colored, so they're edge-disjoint.
- Each hair is at most $D/10$ away from u. We can upper-bound the length of each hair by pruning past $\approx 4(D/10)$.
- There are at least $\frac{1}{2}\left(\frac{D/10}{\alpha D} - 3\right) \geq \frac{1}{40\alpha}$ hairs.

So either χ is α-good, or it contains a 4-bad $1/(40\alpha)$-comb.

Putting It All Together

- Let λ be the doubling constant of T. For $\gamma = 2$, T does not contain a 4-bad λ^6-comb.
- Let $\alpha = 1/(40\lambda^6)$. Since $\lambda \geq 2$, $\alpha \leq 1/60$, so χ is α-good.
- Plugging α back into our distortion lemma gives distortion

$$\frac{2^{1-1/p}}{\frac{1}{40\lambda^6}} \leq 80\lambda^6 = \mathcal{O}(\lambda^6)$$

Conclusion

- We've shown that any tree metric can be embedded into (high-dimensional) ℓ_p with distortion at most $\mathcal{O}(\lambda^6)$. If λ is bounded, then the distortion is constant.
- It can also be shown that we can reduce the dimensionality by reusing colors, $|\mathcal{C}| = \lambda^{\mathcal{O}(\log \lambda)}$.

References

Anupam Gupta, Robert Krauthgamer, and James R. Lee.
Bounded geometries, fractals, and low-distortion embeddings.
In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 534, Washington, DC, USA, 2003. IEEE Computer Society.