

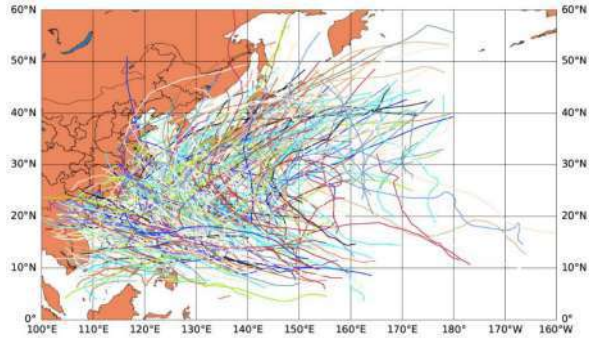
NC STATE UNIVERSITY

Computing the Hausdorff Distance

Parth Parikh


Masters Thesis Defense, 05/01/23

Applications of Hausdorff Distance



- Find similar cyclone trajectories from set of historic cyclone trajectories
- Analyze migration trajectories of different types of birds

Nutanong et al.



- Evaluation of 3D brain tumor segmentation (identifying tumor regions from MRI)

Khotanlou et al.

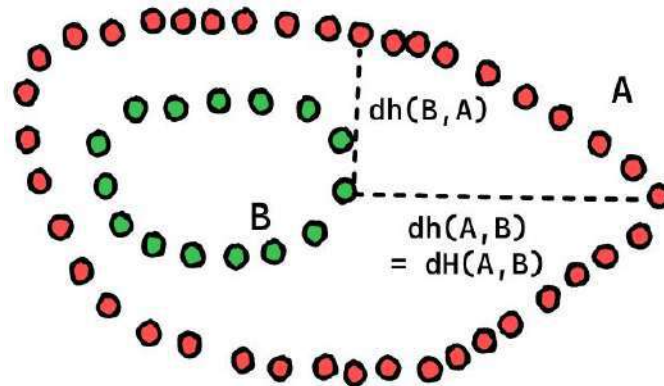
What is Hausdorff Distance?

- The **directed Hausdorff distance** between sets A and B :

$$\mathbf{d}_h(A, B) := \max_{a \in A} \mathbf{d}(a, B).$$

- The **undirected Hausdorff distance** is the larger of the directed Hausdorff distances:

$$\mathbf{d}_H(A, B) := \max\{\mathbf{d}_h(A, B), \mathbf{d}_h(B, A)\}.$$



Introduction and Contributions

- Survey of Hausdorff Distance
- Greedy Trees
- Hausdorff Distance using Greedy Trees

Papers

- Proximity Search in the Greedy Tree
 - *Symposium on Simplicity in Algorithms (SOSA), 2023*
- Linear-Time Approximate Hausdorff Distance
 - *The 30th Fall Workshop on Computational Geometry, 2022*
- Approximating the Directed Hausdorff Distance
 - Submitted to *35th Canadian Conference on Computational Geometry, 2023*
- Greedy Permutations and Finite Voronoi Diagrams
 - Submitted to Multimedia Exposition in Computational Geometry, SoCG 2023
- The Finite Voronoi Method
 - Prepared for resubmission

Survey of HD Techniques for Point Sets

Naive Algorithm for Directed Hausdorff Distance

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):  
  
  LB = 0  
  for a in A:  
    UB = float("inf")  
  
    for b in B:  
      UB = min(UB, d(a, b))  
  
  LB = max(LB, UB)  
  
  return LB
```

Survey of HD Techniques for Point Sets

- Previous surveys?
- Improvements on the naive Quadratic Algorithm
 - The Inner Loop is Nearest Neighbor Search
 - Nearest Neighbor Search is Overkill (The Early Break Heuristic)
 - Ordering Matters
 - There is Spatial Locality in the Searches
 - Preprocessing Allows for Efficient Branch and Bound Algorithms
 - Outer Loop Pruning
- Most heuristics were used for *constant factor* improvements

The Inner Loop is Nearest Neighbor Search

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):
```

```
  LB = 0
```

```
  for a in A:
```

```
    UB = float("inf")
```

```
    for b in B:
```

```
      UB = min(UB, d(a, b))
```

```
    LB = max(LB, UB)
```

```
  return LB
```



OBSERVATION:

inner loop
sequentially computes
the nearest neighbor distance
of each point in A

Nearest Neighbor Search is Overkill (The Early Break Heuristic)

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):
```

```
  LB = 0
```

```
  for a in A:
```

```
    UB = float("inf")
```

```
    for b in B:
```

```
      UB = min(UB, d(a, b))
```

```
    LB = max(LB, UB)
```

```
  return LB
```



OBSERVATION:

If
 $d(a, b) < LB$
in the inner loop,
we can break early!

Ordering Matters

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):
```

```
  LB = 0
```

```
  for a in A:
```

```
    UB = float("inf")
```

```
    for b in B:
```

```
      UB = min(UB, d(a, b))
```

```
  LB = max(LB, UB)
```

```
  return LB
```

OBSERVATION:

Traversal order important
for reducing computations
with spatial correlation

Use randomization or
greedy permutation

↓
Foreshadowing HD
using Greedy Trees

There is Spatial Locality in the Searches

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):
```

```
LB = 0
for a in A:
    UB = float("inf")
```

```
    for b in B:
        # Early Break
        if d(a, b) < UB:
```




OBSERVATION:
The location of
previous early break
can be stored.

Execution can continue from
where we previously left before

```
        break
        UB = min(UB, d(a, b))
```

```
    LB = max(LB, UB)
```

```
return LB
```



Morton Curves for mapping
multidimensional point sets to 1D

Preserves locality of data points

Preprocessing Allows for Efficient Branch and Bound Algorithm

IDEA #1:

1. Partition A into
(A1, A2, ..., At)

$dh(A, B) = \max(dh(A_i, B))$
for i from 1 to t

2. For each part in partition
compute the UB and LB to B

$A_i \subseteq S_i = \text{ball}(c, r)$

$d(c, B) + r \geq dh(A_i, B) \geq d(c, B)$

3. If the UB < LB,
we can prune the part!

```
procedure DIRECTED_HAUSDORFF_DISTANCE(A, B):
```

```
  LB = 0
```

```
  for a in A:  
    UB = float("inf")
```

```
    for b in B:  
      UB = min(UB, d(a, b))
```

```
  LB = max(LB, UB)
```

```
  return LB
```

IDEA #2:

1. Partition B into
(B1, B2, ..., Bt)

2. For each part in partition
compute the UB and LB to a

$B_i \subseteq S_i = \text{ball}(c, r)$

$d(a, c) \leq dh(a, B_i) \geq d(a, c) + r$

3. Pruning

Observation: Use both Idea #1 and #2

Outer Loop Pruning

Sample B' from B

```
graph TD; A[Sample B' from B] --> B[run DIRECTED_HAUSDORFF_DISTANCE(A, B')  
st. for each b' in the inner loop  
keep track of its NN seen so far]; B --> C[using LB, and NN distances]; C --> D[run DIRECTED_HAUSDORFF_DISTANCE(A, B)  
st.  
if at an iteration, b in B'  
AND  
NN(b, A) < LB, skip the iteration!];
```

run DIRECTED_HAUSDORFF_DISTANCE(A, B')
st. for each b' in the inner loop
keep track of its NN seen so far


using LB, and NN distances

run DIRECTED_HAUSDORFF_DISTANCE(A, B)
st.
if at an iteration, b in B'
AND
NN(b, A) < LB, skip the iteration!

Greedy Permutation and Greedy Trees

What is a Good Sample?

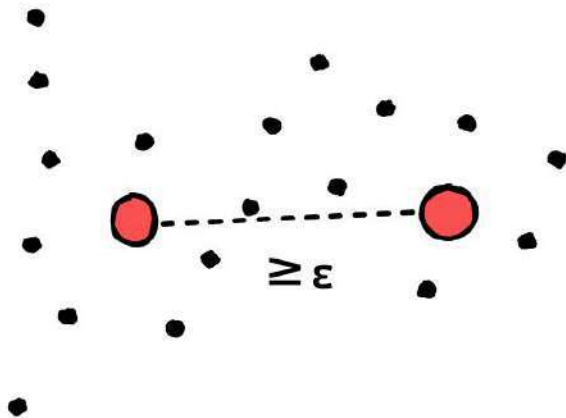
- A good sample can be observed using Hausdorff distance.



- We would like to capture the **geometry of the original set** with far fewer points
 - How?
 - Using two properties: **Packing** and **Covering**

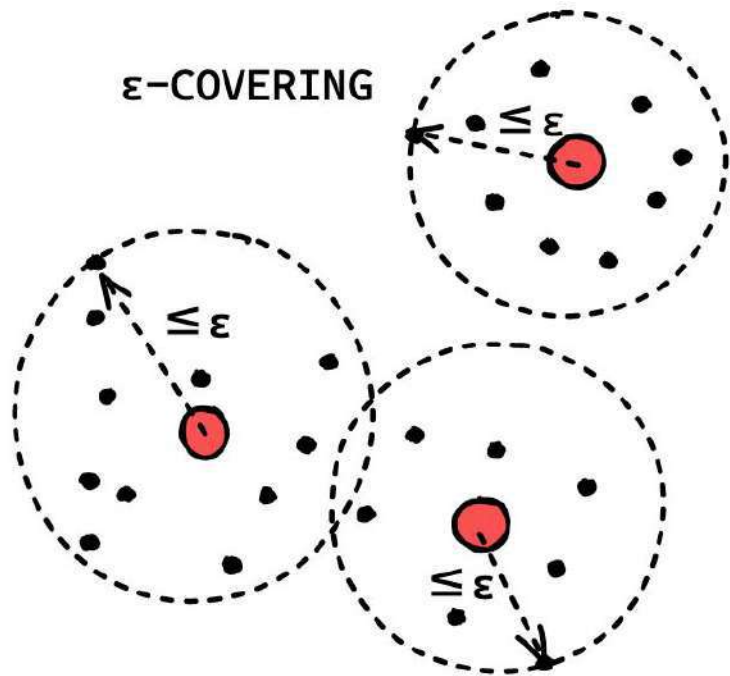
Packing, Covering, and ϵ -Net

ϵ -PACKING



All pairs of sample points are ϵ -apart


ϵ -COVERING



ϵ -radius balls centered at sample points that cover the set

Having seen that, we ask again, what is a Good Sample?

- Has an ε -net for some good ε
- If we know what ε is, we can construct it as:
- But, we might not know what the right ε is!
- **How to compute ε -net for every possible ε ?**
 - And do it all at once
 - Answer: ??



Intuition behind Greedy Permutation


- Key insight:
 - If we have a **really really big** ϵ , one point would be enough (as it would cover everybody!)
 - For **really really small** ϵ , need to include all the points in the ϵ -net
- How does it change from one point to all of the points?
 - Start shrinking down the radius!

Step by step workout of Greedy Permutation


Let the set of all points be P
Initial point 0 is in set P_0

Greedy Permutation: Every new point that we add is the
next farthest point from the
current set P to P_i


ϵ_1 is the
INSERTION
DISTANCE OF 1



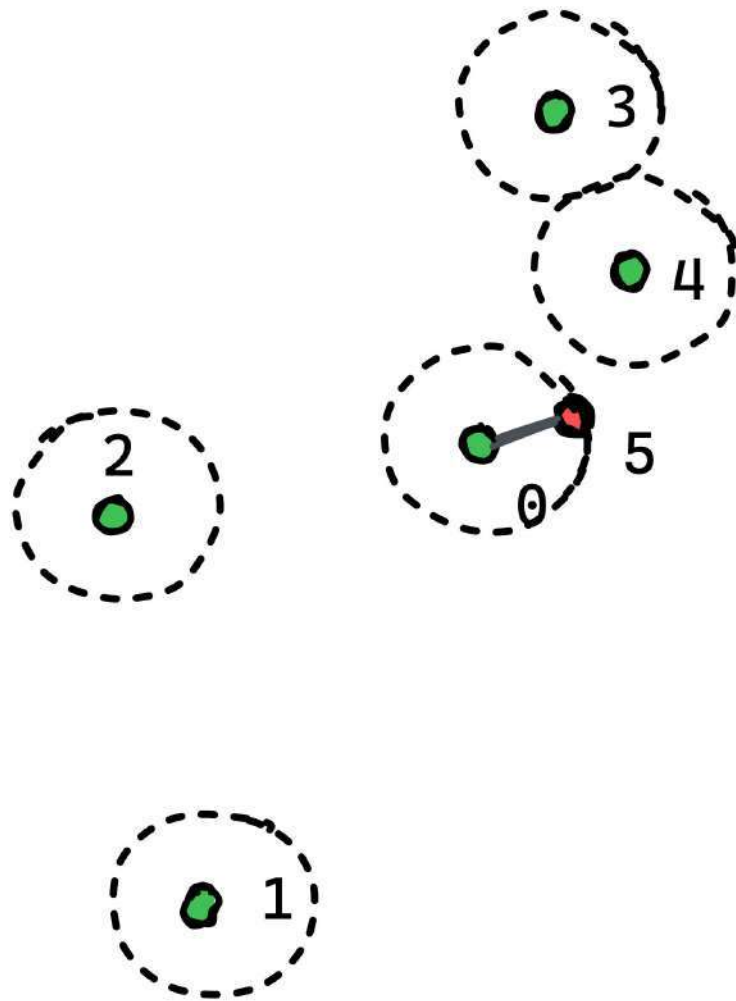
$$P_1 = P_0 \cup \{1\}$$




$$P3 = P2 \cup \{3\}$$



$$\epsilon_4 = \text{dirHD}(P, P_3)$$





NOTE:

1. It often suffices to have an α -APPROXIMATE GREEDY PERMUTATION.

2. Can be computed in $O(n \log n)$ time (Har Peled and Mendel)

3. We define: predecessor mapping T

It maps each point in P to its closest point in the prefix P_i







4. $\epsilon_1 \geq \epsilon_2 \geq \dots \geq \epsilon_5$

5. P_i is the ϵ_i -net

Greedy Trees


Construction of a Greedy Tree

- Greedy Tree? **Binary tree that uses Greedy Permutation to achieve packing guarantees**
- Given Greedy Permutation (G) and Predecessor mapping (\mathbb{T}) for point set \mathbb{P}
- Radius of Node?








Construction of a Greedy Tree

- Greedy Tree? **Binary tree that uses Greedy Permutation to achieve packing guarantees**
- Given Greedy Permutation (G) and Predecessor mapping (T) for point set P
- Incremental construction:
 - Start with a root node centered at p_0
 - Iterate through the points b in the permutation (starting at p_1)
 - Let $a = T(b)$ and let x be the unique leaf of G centered at a
 - Create new nodes centered at a and b and assign them to be the left and right children of x



a



What kind of a binary tree is Greedy Tree?

- **Ball Tree**
- A ball tree on a set \mathbb{A} is a binary tree defined by recursively partitioning \mathbb{A}
 - Each node of the tree stores a **center** and a **radius** that covers the points in its leaves

Motivation behind Greedy Tree

- Lack of strong theoretical guarantees for ball trees
 - Leading to complicated data structures such as cover trees and net-trees
- Greedy Tree *provides strong theoretical guarantees* for proximity search queries
 - **Constructed in $O(n \log n)$ time using Greedy Permutation**
- The radius of a node p is bounded

$$r_p \leq \frac{\epsilon_p}{\alpha - 1}.$$

- Packing guarantee:
 - Let X be a set of pairwise independent nodes from a Greedy Tree. Then the centers of X are:

$$\frac{(\alpha - 1)r}{\alpha} - \text{packed}$$

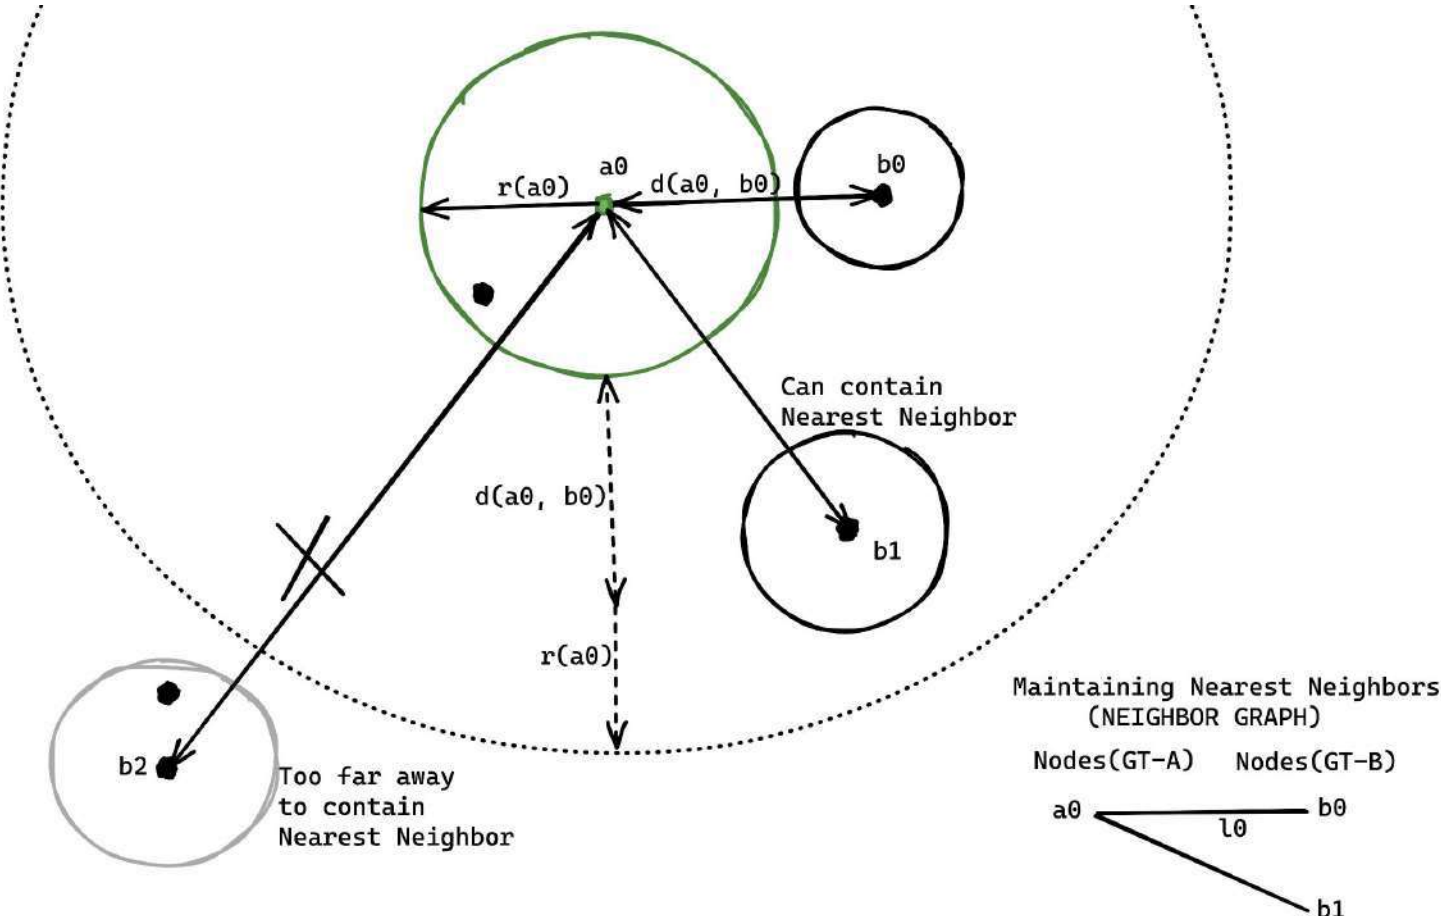
where r is the minimum radius of any parent of a node in X .

Hausdorff Distance using Greedy Trees

A High Level Glance


- What?
 - Computing $(1+\epsilon)$ -approximate **directed** Hausdorff Distance in linear time
 - with $O(n \log n)$ preprocessing time (for Greedy Permutation)
- Input
 - Pair of Greedy Trees G_A and G_B for sets A and B, respectively
 - List of their nodes in non-increasing order of radius
 - Iteration order of the algorithm
 - Approximation Factor ϵ

Intuition behind the Algorithm




Step by step workout of the Algorithm


Initializing the Neighbor Graph




Case: Node in GT-A




Case: Node in GT-B



Case: Node in GT-A



Case: Pruning




Pruning edges
 (a_0, b_0) and
 (a_2, b_0)

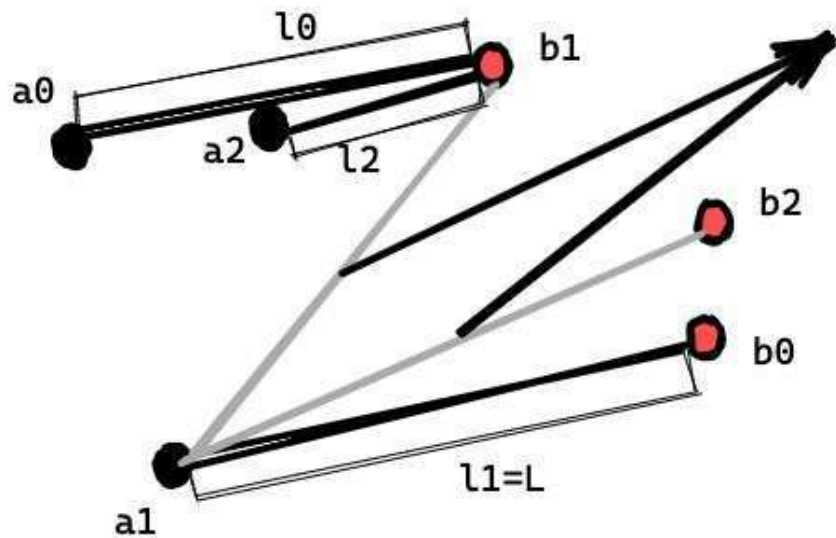
WHY? For (a_0, b_0) :

$$d(a_0, b_1) + r(a_0) < d(a_0, b_0) - r(a_0) - r(b_0)$$

Case: Node in GT-B



Case: Pruning

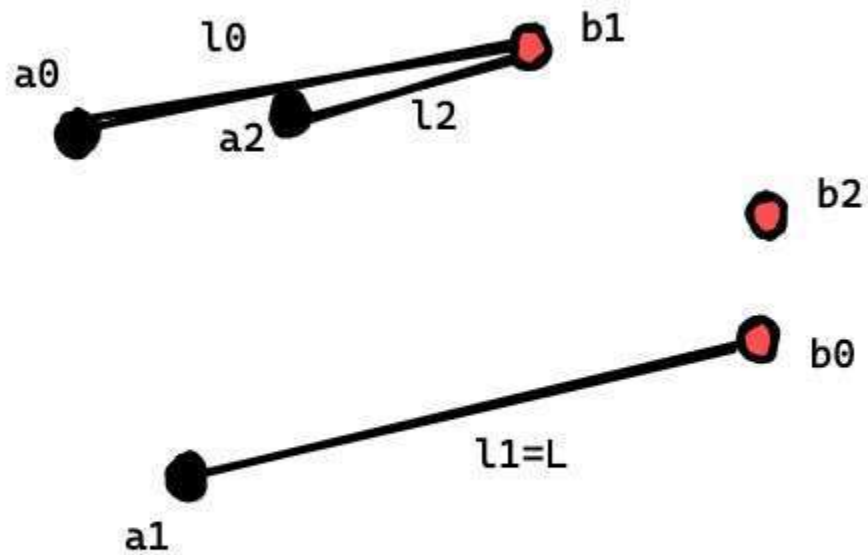


Pruning edges
 (a_1, b_1) and
 (a_1, b_2)

WHY? For (a_1, b_1) :

$$d(a_1, b_0) + r(a_1) < d(a_1, b_1) - r(a_1) - r(b_1)$$

Case: Hausdorff Distance



DIRECTED HAUSDORFF DISTANCE FROM A TO B
IS $d(a_1, b_0) = l_1 = L$

A better Stopping Criteria

- Key Idea: Stop once the unprocessed nodes have radii **too small** to significantly affect the LB, and return L
- What is “too small”?
 - If the radius r of the largest node yet to be processed is

$$r \leq \frac{\varepsilon}{2} L$$

Correctness

- We show that:
 - $L \leq \text{dirHD}(A, B) \leq L + 2r$
- Thus, when $r \leq \frac{\varepsilon}{2}L$
 - Directed HD is $(1+\varepsilon)$ -approximate

Conclusion

- Questions?
- Code is available <https://github.com/donsheehy/greedypermutation>
- *Thank you Don, Siddharth, Oliver, and Kirk*

Additional Resources

- Proximity Search in the Greedy Tree
 - *Symposium on Simplicity in Algorithms (SOSA), 2023*
- Linear-Time Approximate Hausdorff Distance
 - *The 30th Fall Workshop on Computational Geometry, 2022*
- Approximating the Directed Hausdorff Distance
 - Submitted to *35th Canadian Conference on Computational Geometry, 2023*
- Greedy Permutations and Finite Voronoi Diagrams
 - Submitted to Multimedia Exposition in Computational Geometry, SoCG 2023
- The Finite Voronoi Method
 - Prepared for resubmission