

A Fuzzy Spatial Relationship Graph for Point Clouds Using Bounding Boxes

Andrew R. Buck^a, Derek T. Anderson^a, James M. Keller^a,
Robert H. Luke III^b, and Grant Scott^a

^aElectrical Engineering and Computer Science (EECS) Department,
University of Missouri, Columbia, MO, USA

^bUS Army DEVCOM C5ISR Center, Fort Belvoir, VA, USA

FUZZ-IEEE 2021

Presented by Andrew Buck




University of Missouri





Motivation

- ▶ 3D scenes have a lot of information!
 - ▶ How should it be represented?
 - ▶ What are the relationships between objects?
- ▶ Point clouds give raw 3D data
 - ▶ Easy to acquire as raw data from LIDAR or depth camera
 - ▶ Files can be huge! Voxels can help sometimes...
 - ▶ How to represent the important aspects of the scene?
- ▶ Semantic segmentation can identify individual objects
 - ▶ How to store and query spatial configurations?
 - ▶ Use bounding boxes to represent important objects
 - ▶ Compute relationships between bounding boxes



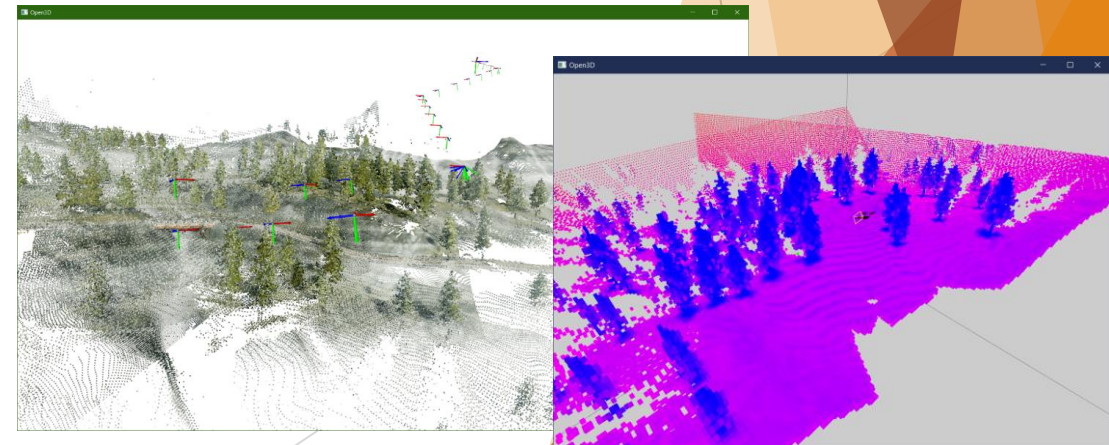
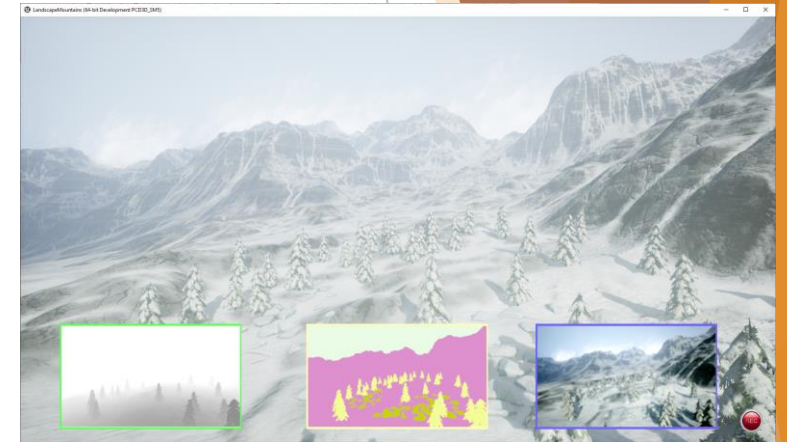
<http://www.semantic3d.net/>

<https://github.com/microsoft/AirSim>




Applications

- ▶ Ultimately, we would like AI systems to have an interpretable understanding of their environment
 - ▶ This can help design and communicate intended behaviors
 - ▶ Make an AI agent act more like a human
- ▶ Unmanned Aerial Vehicles (UAVs)
 - ▶ Small embedded systems need to respond in real-time
 - ▶ Require minimal overhead and processing
- ▶ Human robot interaction
 - ▶ Use natural language to communicate
 - ▶ Mobile computing devices (AR/VR headsets) with limited streaming bandwidth





Benchmark Datasets

- ▶ We assume that a semantic segmentation of the scene can be acquired
 - ▶ Many recent works for labeling and segmenting point clouds
- ▶ Focus here on ground truth, hand-labeled benchmark datasets
 - ▶ Each object instance is given a unique ID
 - ▶ Easy to find individual objects or categories
- ▶ Looking at the NPM3D benchmark suite
 - ▶ <https://npm3d.fr/paris-lille-3d>
- ▶ Outdoor street scene (static)
- ▶ Chose this dataset because it has ground truth segmentation
 - ▶ Objects are identified by class and an instance ID
 - ▶ 50 classes organized in a hierarchal ontology



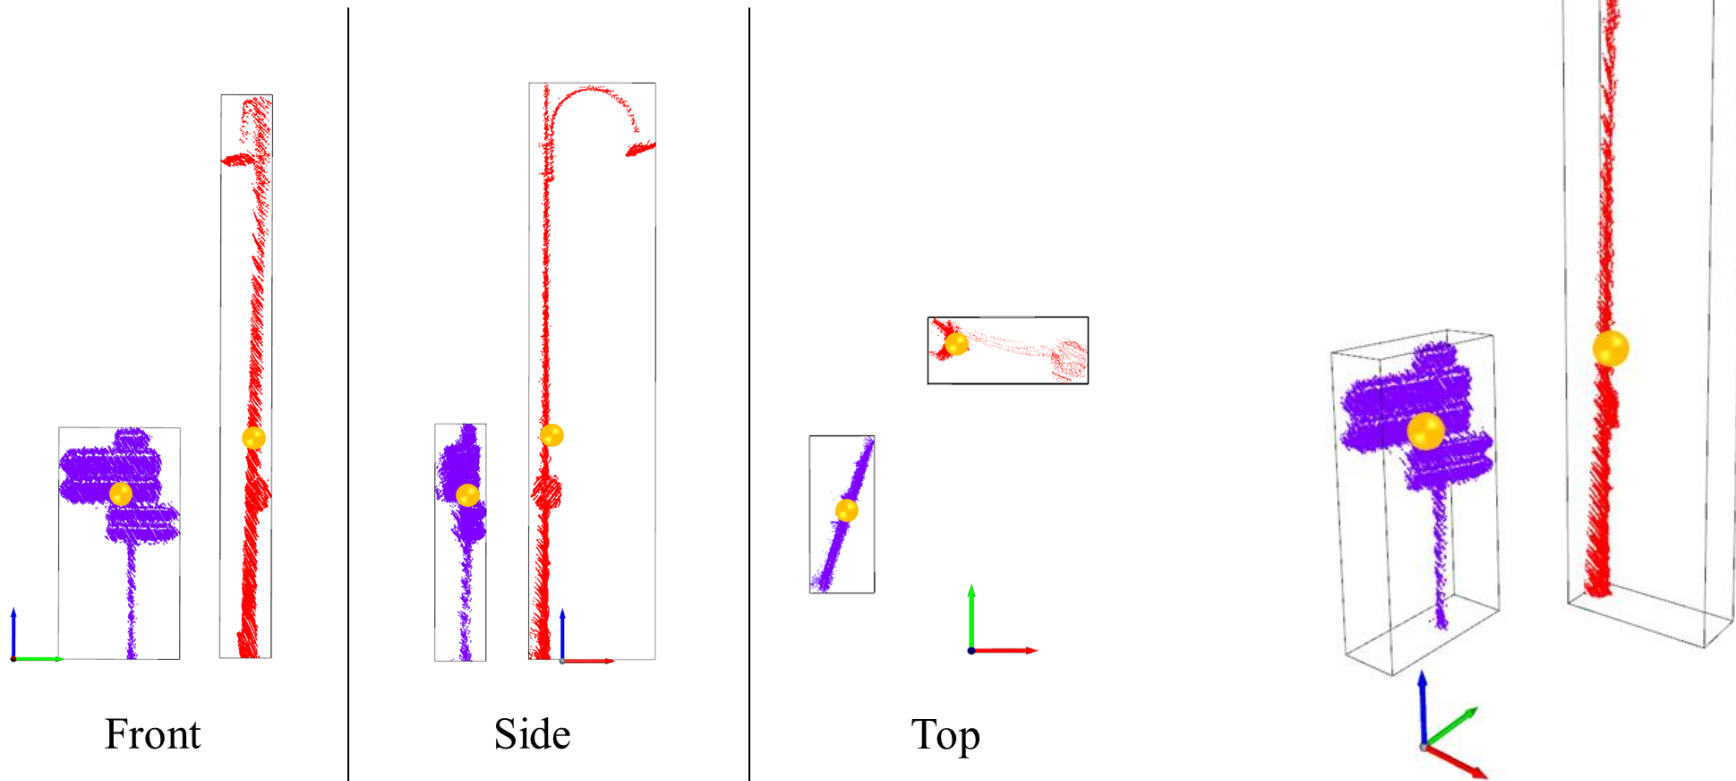
Example Scene

- ▶ We chose to look at an example region with ~10,000,000 points and ~100 labeled objects
- ▶ How to compute the spatial relationships between objects?
- ▶ Each object can be shown with an axis-aligned bounding box (trivial to compute)



Bounding Box Representation

- ▶ Consider the relationship between these two objects.
- ▶ We can easily compute the bounding boxes and centroids.



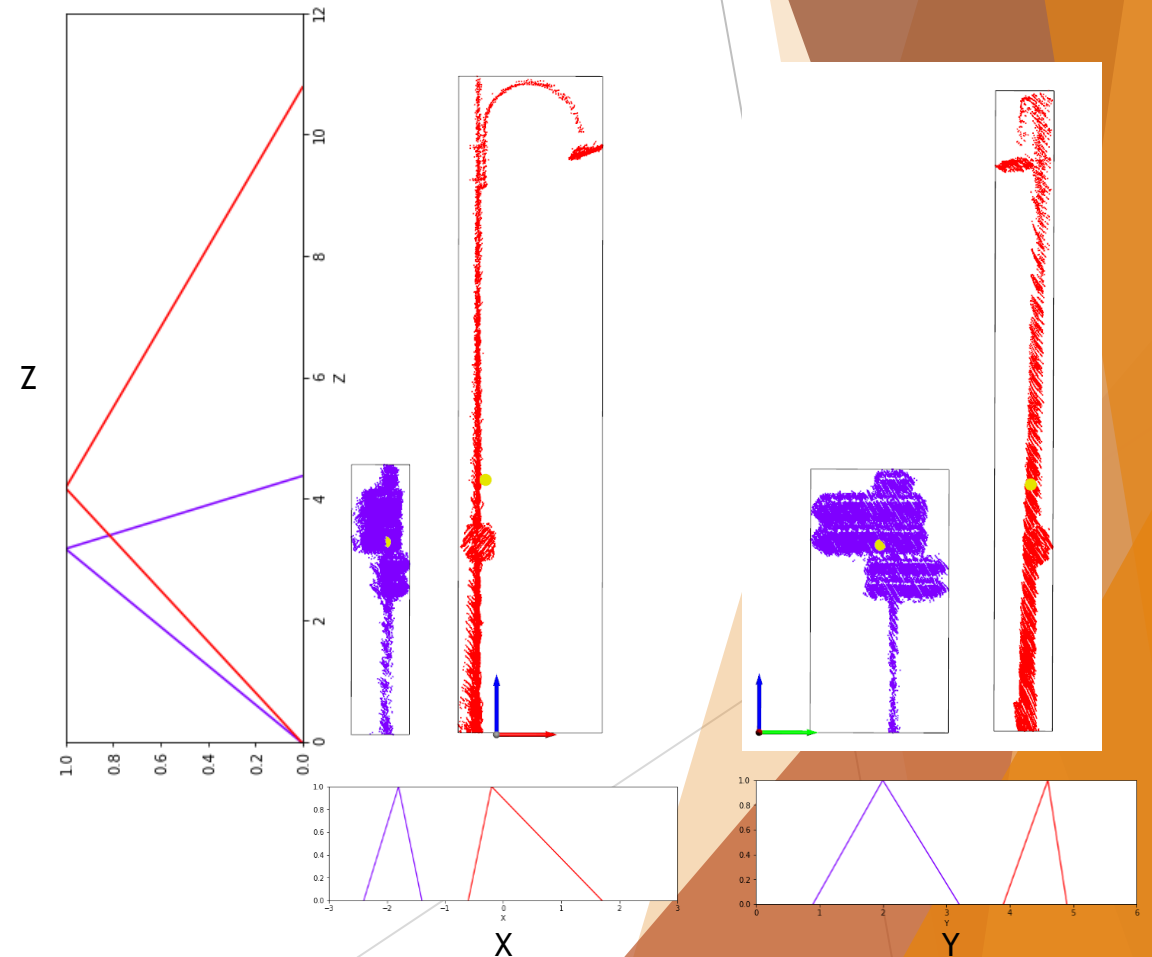
Triangular Fuzzy Numbers

- ▶ Along each dimension, we define a triangular fuzzy number (TFN) from the minimum and maximum extents of the bounding box and the object centroid.
- ▶ Store 9 values for each object

$$A = \text{Tri}(a_1, a_2, a_3), \quad B = \text{Tri}(b_1, b_2, b_3)$$

Object Bounds and Centroid: [min, centroid, max]

	X	Y	Z
Signpost	[-2.4, -1.8, -1.4]	[0.9, 2.0, 3.2]	[0, 3.2, 4.4]
Light pole	[-0.6, -0.2, 1.7]	[3.9, 4.6, 4.9]	[0, 4.2, 10.8]



Bounding Box Distance

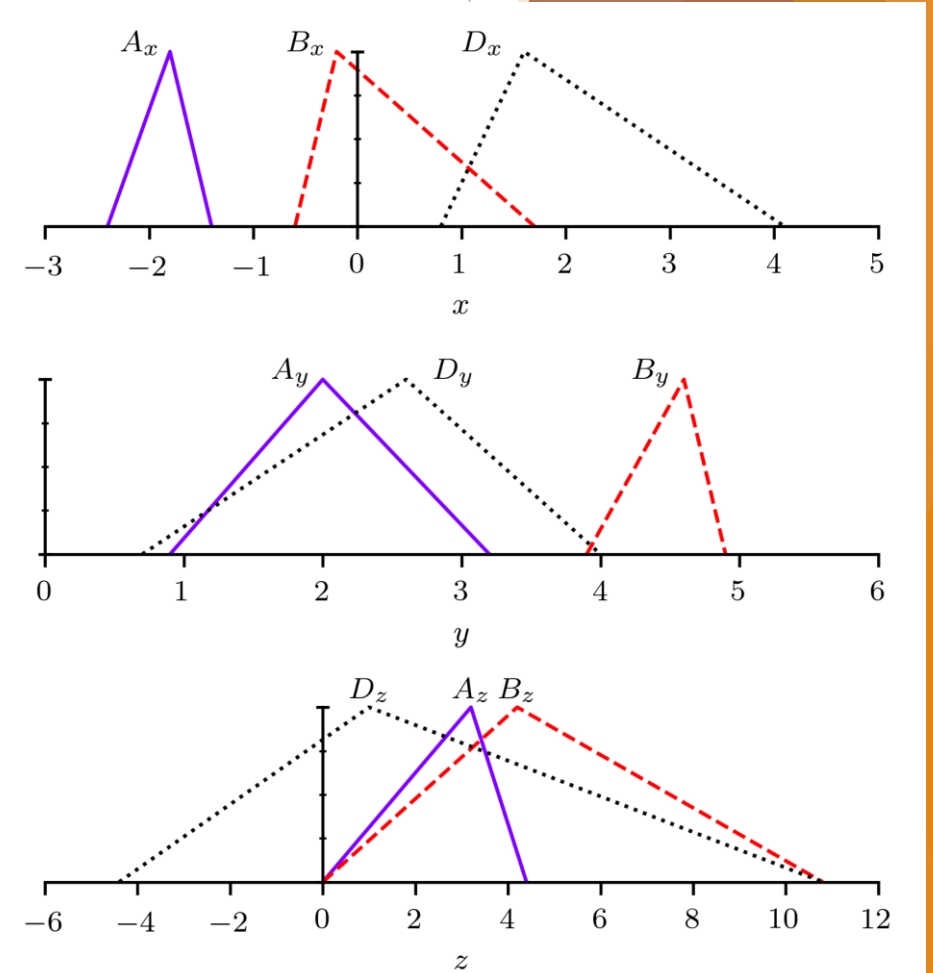
- ▶ Using fuzzy arithmetic, the difference between the two objects is computed along each dimension as a new TFN.
- ▶ This represents the minimum, maximum, and average distance between objects A and B in each dimension.
- ▶ (can be negative)

$$A - B = \text{Tri}(a_1 - b_3, a_2 - b_2, a_3 - b_1)$$

$$\begin{aligned} A_x &= \text{Tri}(-2.4, -1.8, -1.4) \\ A_y &= \text{Tri}(0.9, 2.0, 3.2) \\ A_z &= \text{Tri}(0.0, 3.2, 4.4) \end{aligned}$$

$$\begin{aligned} B_x &= \text{Tri}(-0.6, -0.2, 1.7) \\ B_y &= \text{Tri}(3.9, 4.6, 4.9) \\ B_z &= \text{Tri}(0.0, 4.2, 10.8) \end{aligned}$$

$$\begin{aligned} D_x &= B_x - A_x = \text{Tri}(0.8, 1.6, 4.1) \\ D_y &= B_y - A_y = \text{Tri}(0.7, 2.6, 4.0) \\ D_z &= B_z - A_z = \text{Tri}(-4.4, 1.0, 10.8) \end{aligned}$$



Overall Distance

- ▶ The overall distance is computed as the Euclidean norm of the differences along each axis using fuzzy arithmetic.

$$D_x = \text{Tri}(0.8, 1.6, 4.1)$$

$$D_y = \text{Tri}(0.7, 2.6, 4.0)$$

$$D_z = \text{Tri}(-4.4, 1.0, 10.8)$$

$$D_x^2 = \text{Tri}(0.64, 2.56, 16.81)$$

$$D_y^2 = \text{Tri}(0.49, 6.76, 16.0)$$

$$D_z^2 = \text{Tri}(0.0, 1.0, 116.64)$$

$$D_x^2 + D_y^2 + D_z^2 = \text{Tri}(1.13, 10.32, 149.45)$$



$$D_{AB} = \sqrt{D_x^2 + D_y^2 + D_z^2} = \text{Tri}(1.06, 3.21, 12.22)$$

$$A^2 = \text{Tri}(a_{\min}, a_2^2, \max\{a_1^2, a_3^2\}),$$

$$a_{\min} = \begin{cases} \min\{0, a_1^2, a_3^2\}, & \text{if } a_1 \leq 0 \leq a_3 \\ \min\{a_1^2, a_3^2\}, & \text{otherwise} \end{cases}$$



$$A + B = \text{Tri}(a_1 + b_1, a_2 + b_2, a_3 + b_3)$$

$$\sqrt{A} = \text{Tri}(\sqrt{a_1}, \sqrt{a_2}, \sqrt{a_3}), \quad 0 \leq a_1 \leq a_2 \leq a_3$$



Spatial Relationship Graph

- ▶ Knowing the distances between objects lets us define a spatial relationship graph over a scene to show how objects are connected.
- ▶ We'll add an edge between two nodes (objects) if the distance between them is less than some threshold d .
- ▶ So, we need a way to determine if a triangular fuzzy number represents a distance that is less than d .



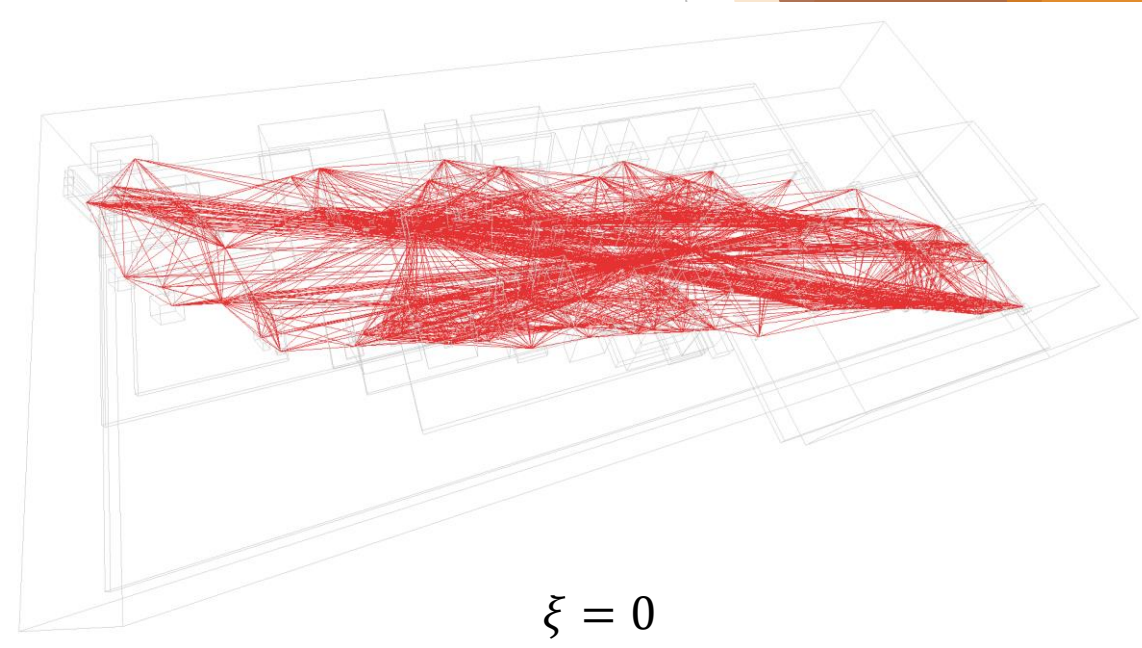
TFN Defuzzification

- ▶ Given a triangular fuzzy number $X = \text{Tri}(a, b, c)$, we can defuzzify to a crisp value using an optimism/pessimism parameter $\xi \in [0,1]$.

$$\Gamma(X|\xi) = \begin{cases} a + 2\xi(b - a), & \xi \leq 0.5 \\ b + 2(\xi - 0.5)(c - b), & \xi > 0.5 \end{cases}$$

- ▶ This gives a way to select the minimum ($\xi = 0$), maximum ($\xi = 1$), or average ($\xi = 0.5$) values of the TFN.
- ▶ When ξ is high, it's like complete linkage clustering.
- ▶ When ξ is low, it's like single linkage clustering.

$d = 10$ meters




Distance Queries

- ▶ Suppose we want to find all objects that are a certain distance away from a reference object.
- ▶ We define the query distance as a TFN $Q = \text{Tri}(q_1, q_2, q_3)$.
- ▶ The similarity between two TFNs can be computed as the maximum of their intersecting points.

$$S(A, B) = \max_{x \in \mathbb{R}} \{ \min(\mu_A(x), \mu_B(x)) \}$$


- ▶ The distance D_{AB} between objects A and B can be compared with the query distance Q to give the distance similarity $s_{\text{dist}} \in [0, 1]$.



Example Distance Queries



$$Q = \text{Tri}(0, 0, 1)$$




$$Q = \text{Tri}(5, 10, 15)$$



$$Q = \text{Tri}(20, 30, 40)$$



$$Q = \text{Tri}(50, 75, 100)$$



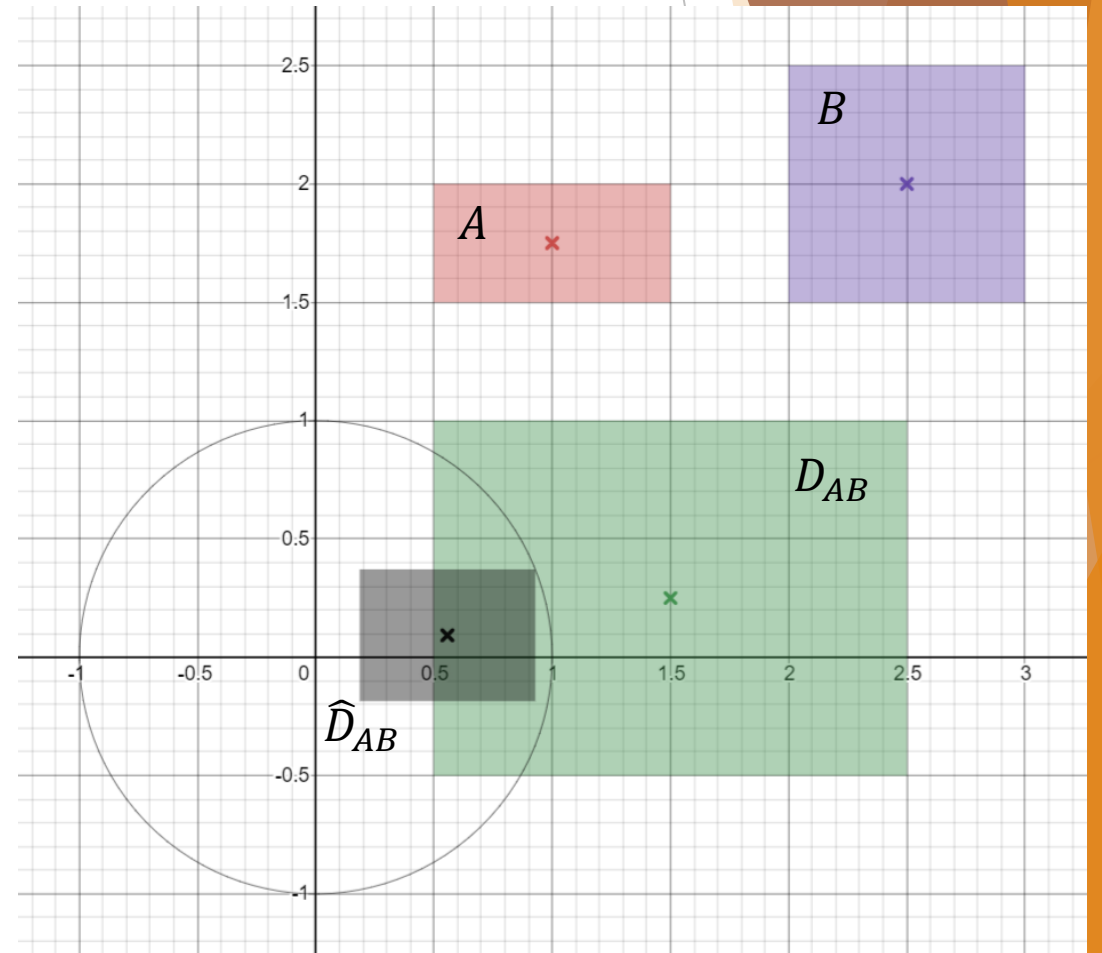
Looking at the heat maps for different distance queries from a person in the scene, outlined with a red box.

Normalized Direction

- ▶ The axis-aligned difference TFNs D_x , D_y , and D_z encode both the relative *distance* and *direction* between two objects.
- ▶ We can use this to compute how much support there is for the statement “Object B is in direction \hat{u} from Object A ,” where \hat{u} is a unit vector pointing in the direction of interest.
- ▶ First, we need a normalized difference vector, $\hat{D} = [\hat{D}_x, \hat{D}_y, \hat{D}_z]$, where \hat{D}_x , \hat{D}_y , and \hat{D}_z are normalized versions of the computed difference TFNs D_x , D_y , and D_z .

$$\hat{D}_k = \alpha D_k, \quad \text{s.t.} \quad \max_{k \in \{x, y, z\}} \hat{D}_k = 1$$

- ▶ Consider a 2-dimensional example...




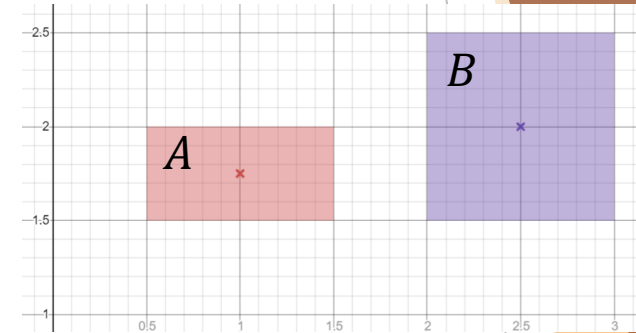
Directional Similarity

- ▶ Given a reference direction \hat{u} , the directional similarity to the normalized difference TFN \hat{D} is the dot product.

$$S_{\text{dir}} = \hat{D} \cdot \hat{u} = \hat{D}_x u_x + \hat{D}_y u_y + \hat{D}_z u_z$$


- ▶ S_{dir} is a TFN bounded in the range $[-1, 1]$.
 - ▶ See example...
- ▶ To reduce the directional similarity to a scalar value (like distance), we can use the defuzzification parameter ξ and clamp to positive values.

$$s_{\text{dir}} = \max\{0, \Gamma(S_{\text{dir}}|\xi)\}$$




Example Directional Queries

Strictly to the north



$$\xi = 0$$
$$\hat{\mathbf{u}} = [0, 1, 0]$$

Mainly to the southeast




$$\xi = 0.5$$
$$\hat{\mathbf{u}} = [0.707, -0.707, 0]$$

Generally to the northeast




$$\xi = 1$$
$$\hat{\mathbf{u}} = [0.707, 0.707, 0]$$

Generally above



$$\xi = 1$$
$$\hat{\mathbf{u}} = [0, 0, 1]$$



Looking at the heat maps for different directional queries from a person in the scene, outlined with a red box.

A Multi-Criteria Framework


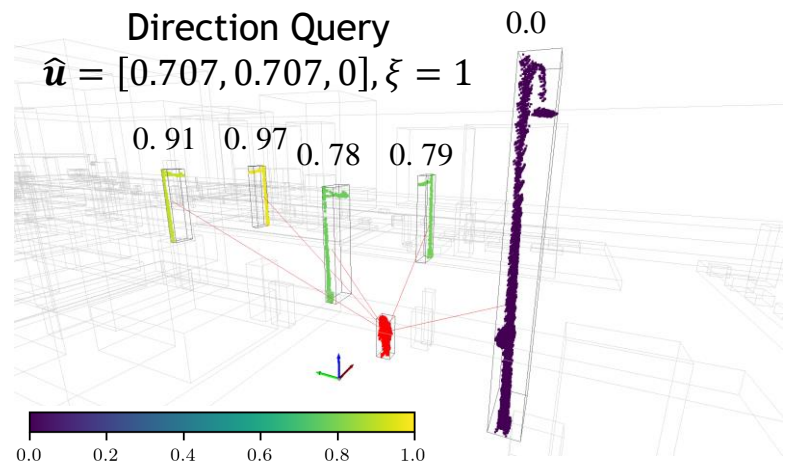
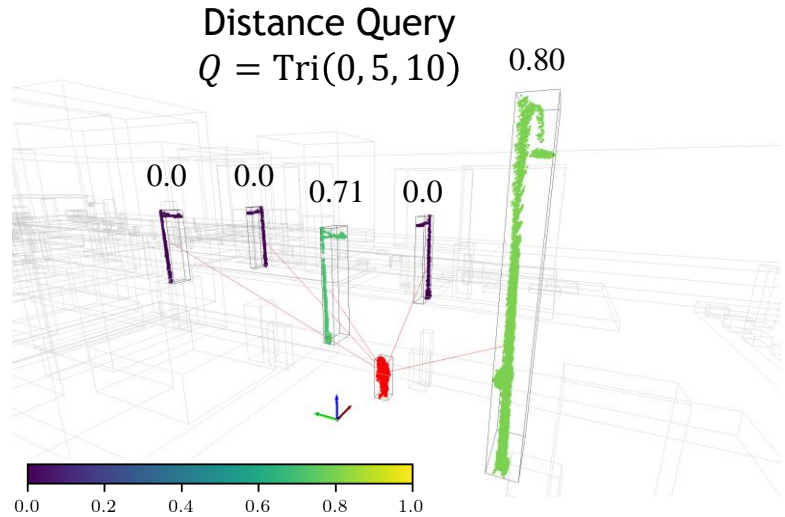
- ▶ Distance and direction are two features that can be used to search for objects in a scene.
- ▶ Other features might include class type, number of neighbors of a certain type, location in world space, etc.
- ▶ Our criteria for object selection is represented by a normalized feature vector $\mathbf{s} = [s_1, \dots, s_n]$, where each $s_i \in [0, 1]$ represents the degree to which an object satisfies a particular set of criteria.
- ▶ The multidimensional feature vector can be mapped to a single value with a scalarization function $g_\theta(\mathbf{s})$, where θ represents the parameterization.

$$g_{\text{avg}}(\mathbf{s}) = \frac{1}{n} \sum_{i=1}^n s_i$$

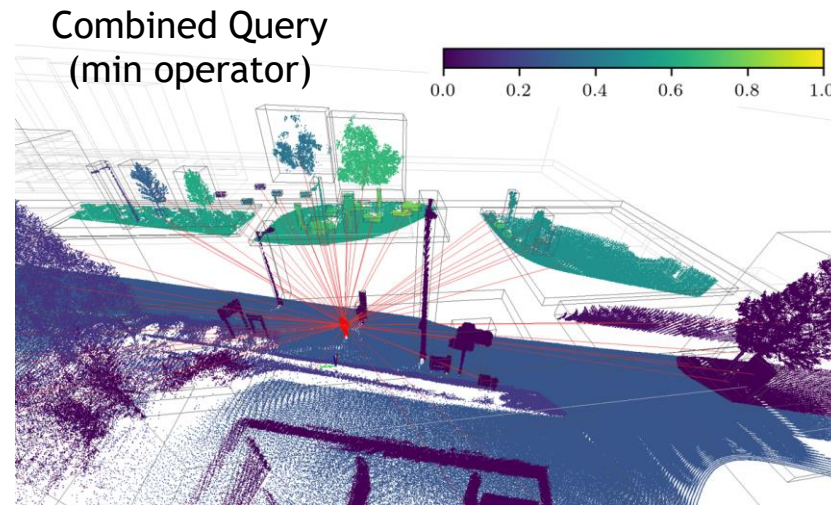

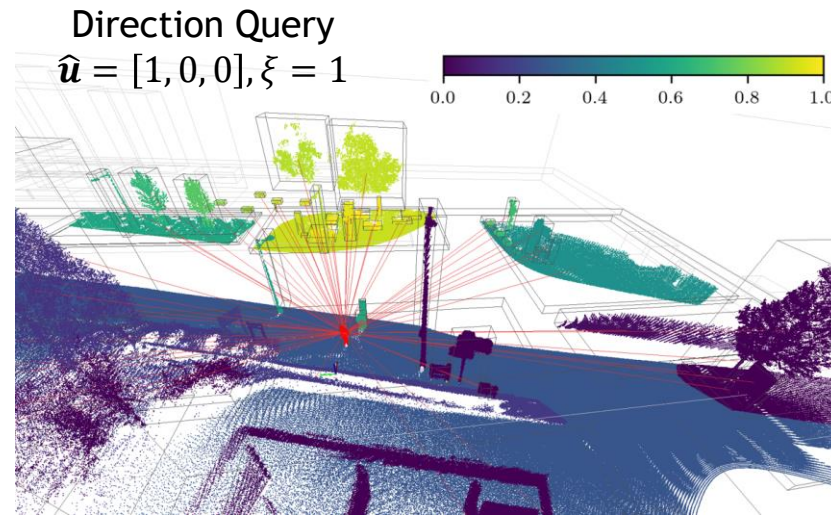
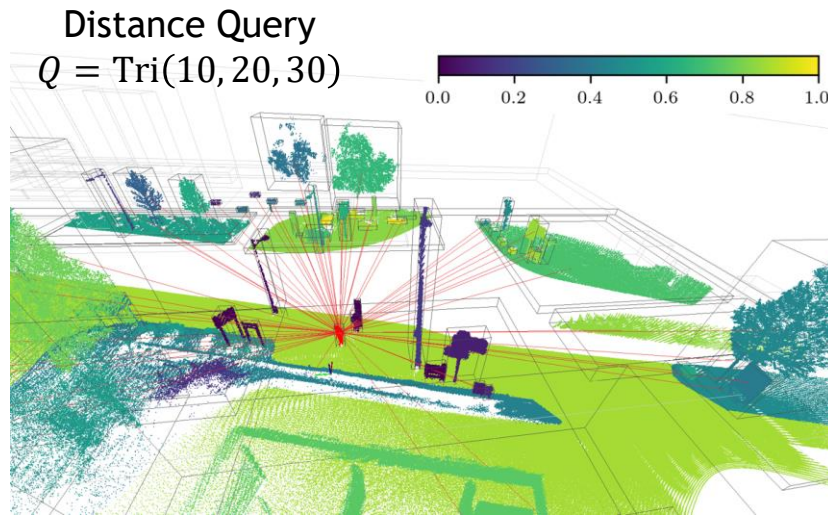
$$g_{\text{min}}(\mathbf{s}) = \min_i s_i$$

Example: Multi-Criteria Selection

- ▶ Consider a person (red) looking to identify a lamp post that is “near” and to the “front-left”.



Example: Choosing an Exploration Target



Consider an agent in the environment choosing its next exploration target.

Want to select an object to interrogate that is

- Nearby
- In the forward direction
- Near the edge of what's already been explored

Conclusions

- ▶ We've shown a way to compute distance and directional spatial relationships between objects in a 3D scene represented with axis-aligned bounding boxes.
- ▶ These features can be used to construct a spatially attributed graph of the environment and search for objects using multiple criteria.
- ▶ Simplifying the representation to bounding boxes instead of full point clouds helps achieve real-time performance on embedded hardware.
- ▶ Future directions:
 - ▶ Handle dynamic environments and an incremental/updating graph
 - ▶ Integrate with semantic segmentation algorithms for point clouds
 - ▶ Use a hierarchical representation to represent large compound objects
 - ▶ E.g., buildings with windows and doorways, roads and intersections, etc.

