


# Tetrahedron-Tetrahedron Intersection and Volume Computation Using Neural Networks



Erendiro Pedro<sup>1</sup>, Hermann Meißenhelger<sup>2</sup>, Gabriel Zachmann<sup>2</sup>

<sup>1</sup>Institute of Engineering, Polytechnic of Porto, Portugal


<sup>2</sup>University of Bremen, Germany

**EUROGRAPHICS  
AACHEN 2026**



*Eurographics 2026, 4-8 May, Aachen, Germany*



- Narrow-phase collision detection is a bottleneck in physics simulation
  - Tetrahedral meshes are common in FEM and deformable simulation
  - May require millions of pairwise checks per frame
- Beyond yes/no collision, many applications need an estimate of intersection extent
- Intersection volume is the “most complicated yet accurate method” for intersection extent [FisherLin01]



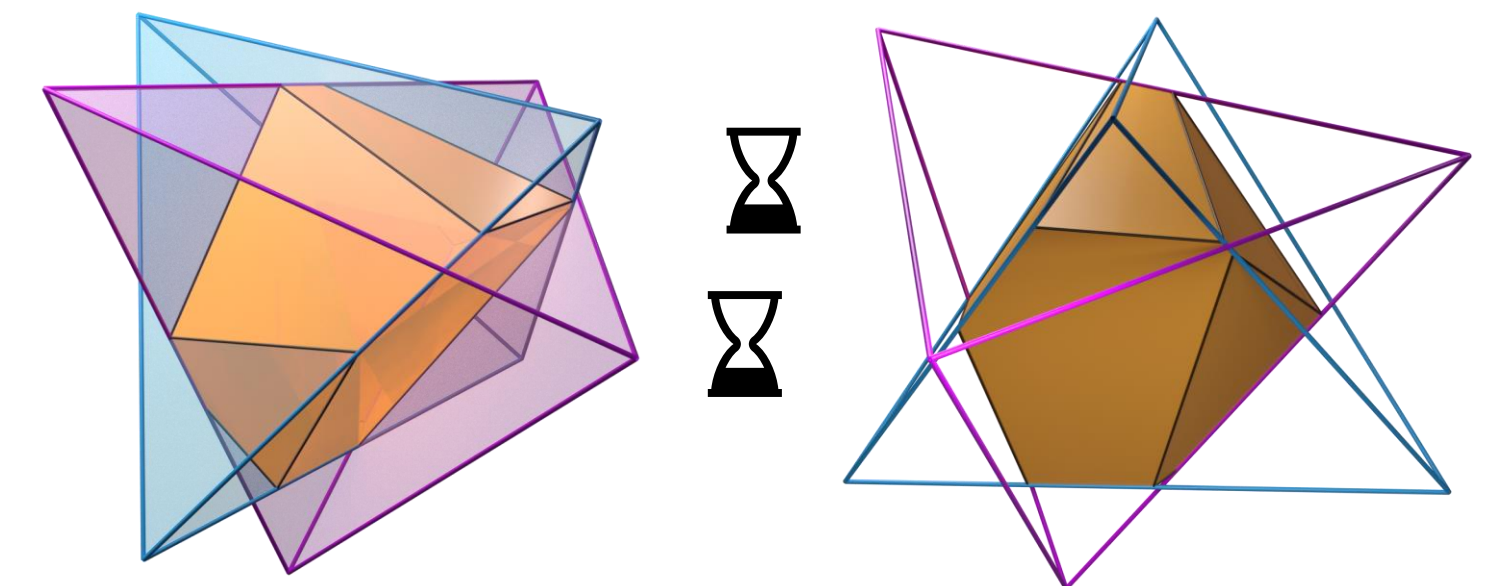
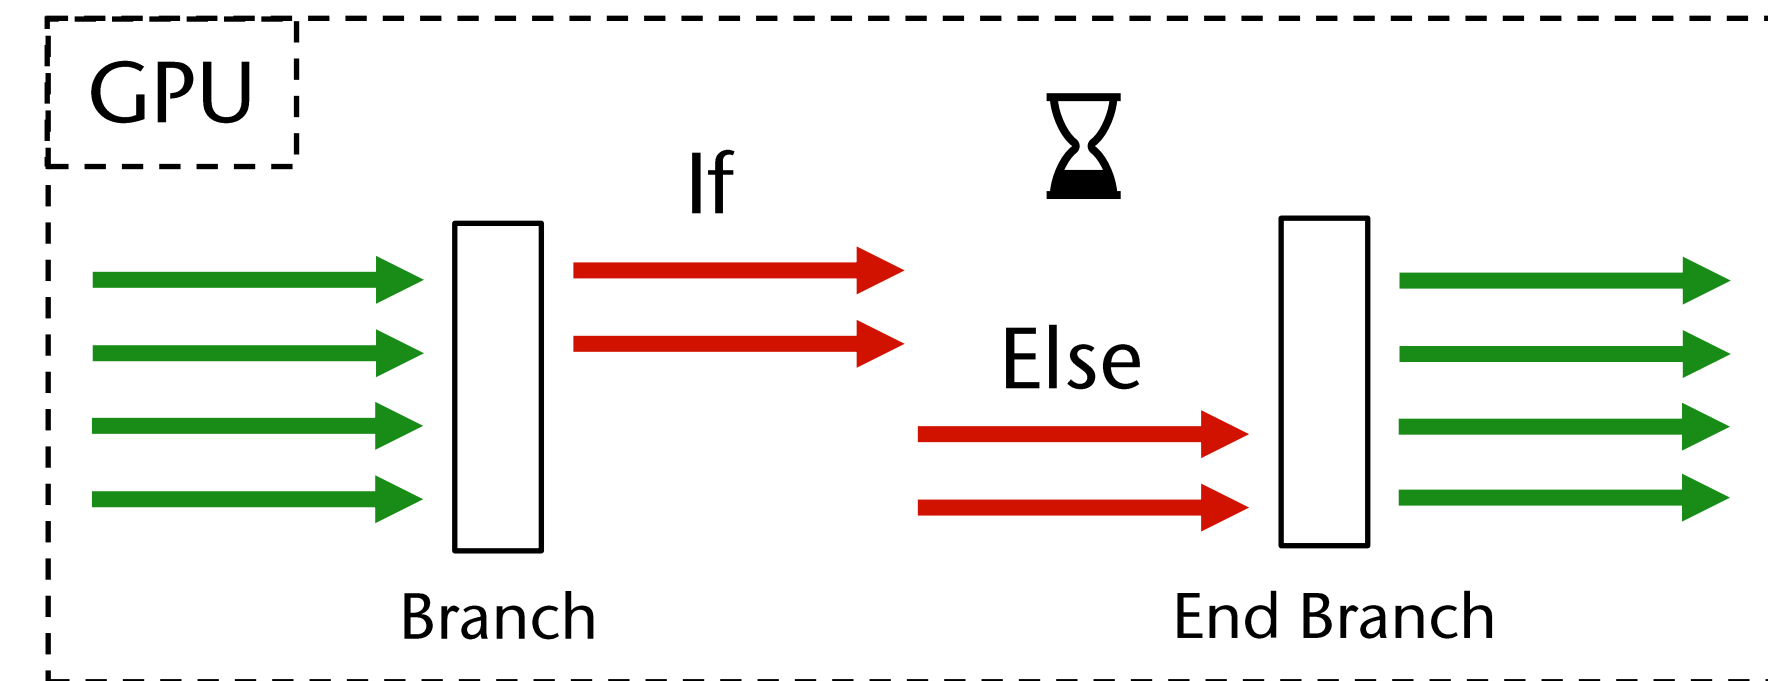
Armadillo models, each with 371,000 tetrahedral elements [Bender17]



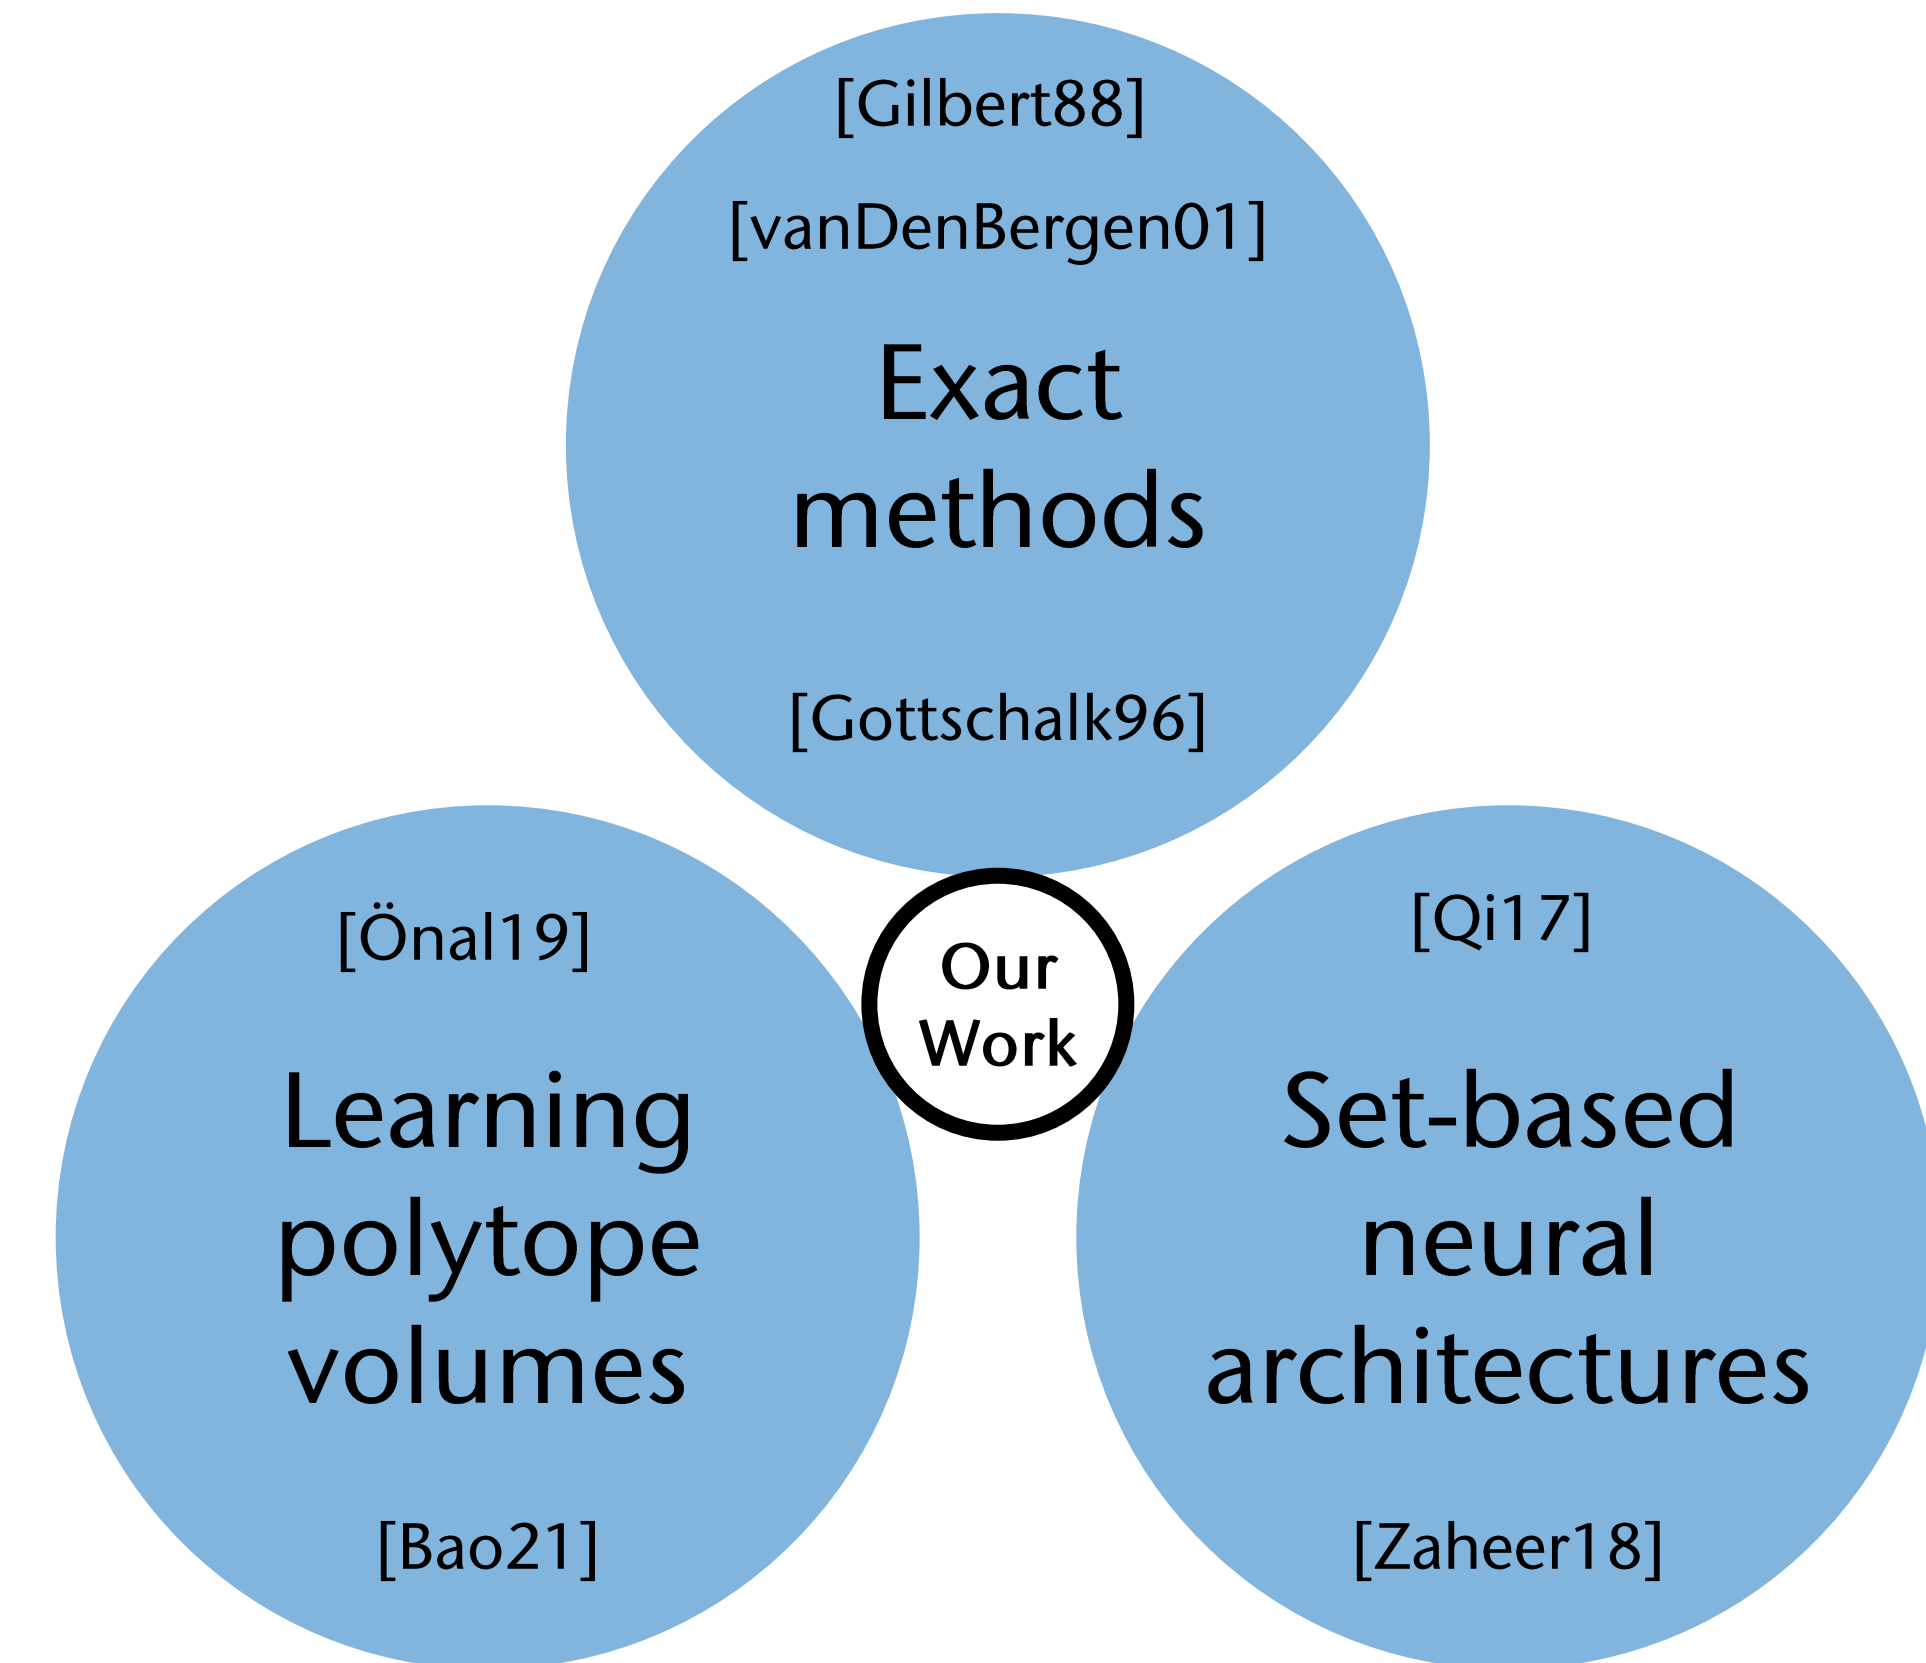
# Problem and Gap

- Classical exact methods rely on sequential branching logic
  - This causes thread divergence and poor GPU utilization
- Volume computation is even more expensive than intersection testing
  - Requires constructing the overlap geometry
  - Becomes difficult to use at large scale
- Question: Can batched neural inference predict tet–tet intersection and overlap volume?

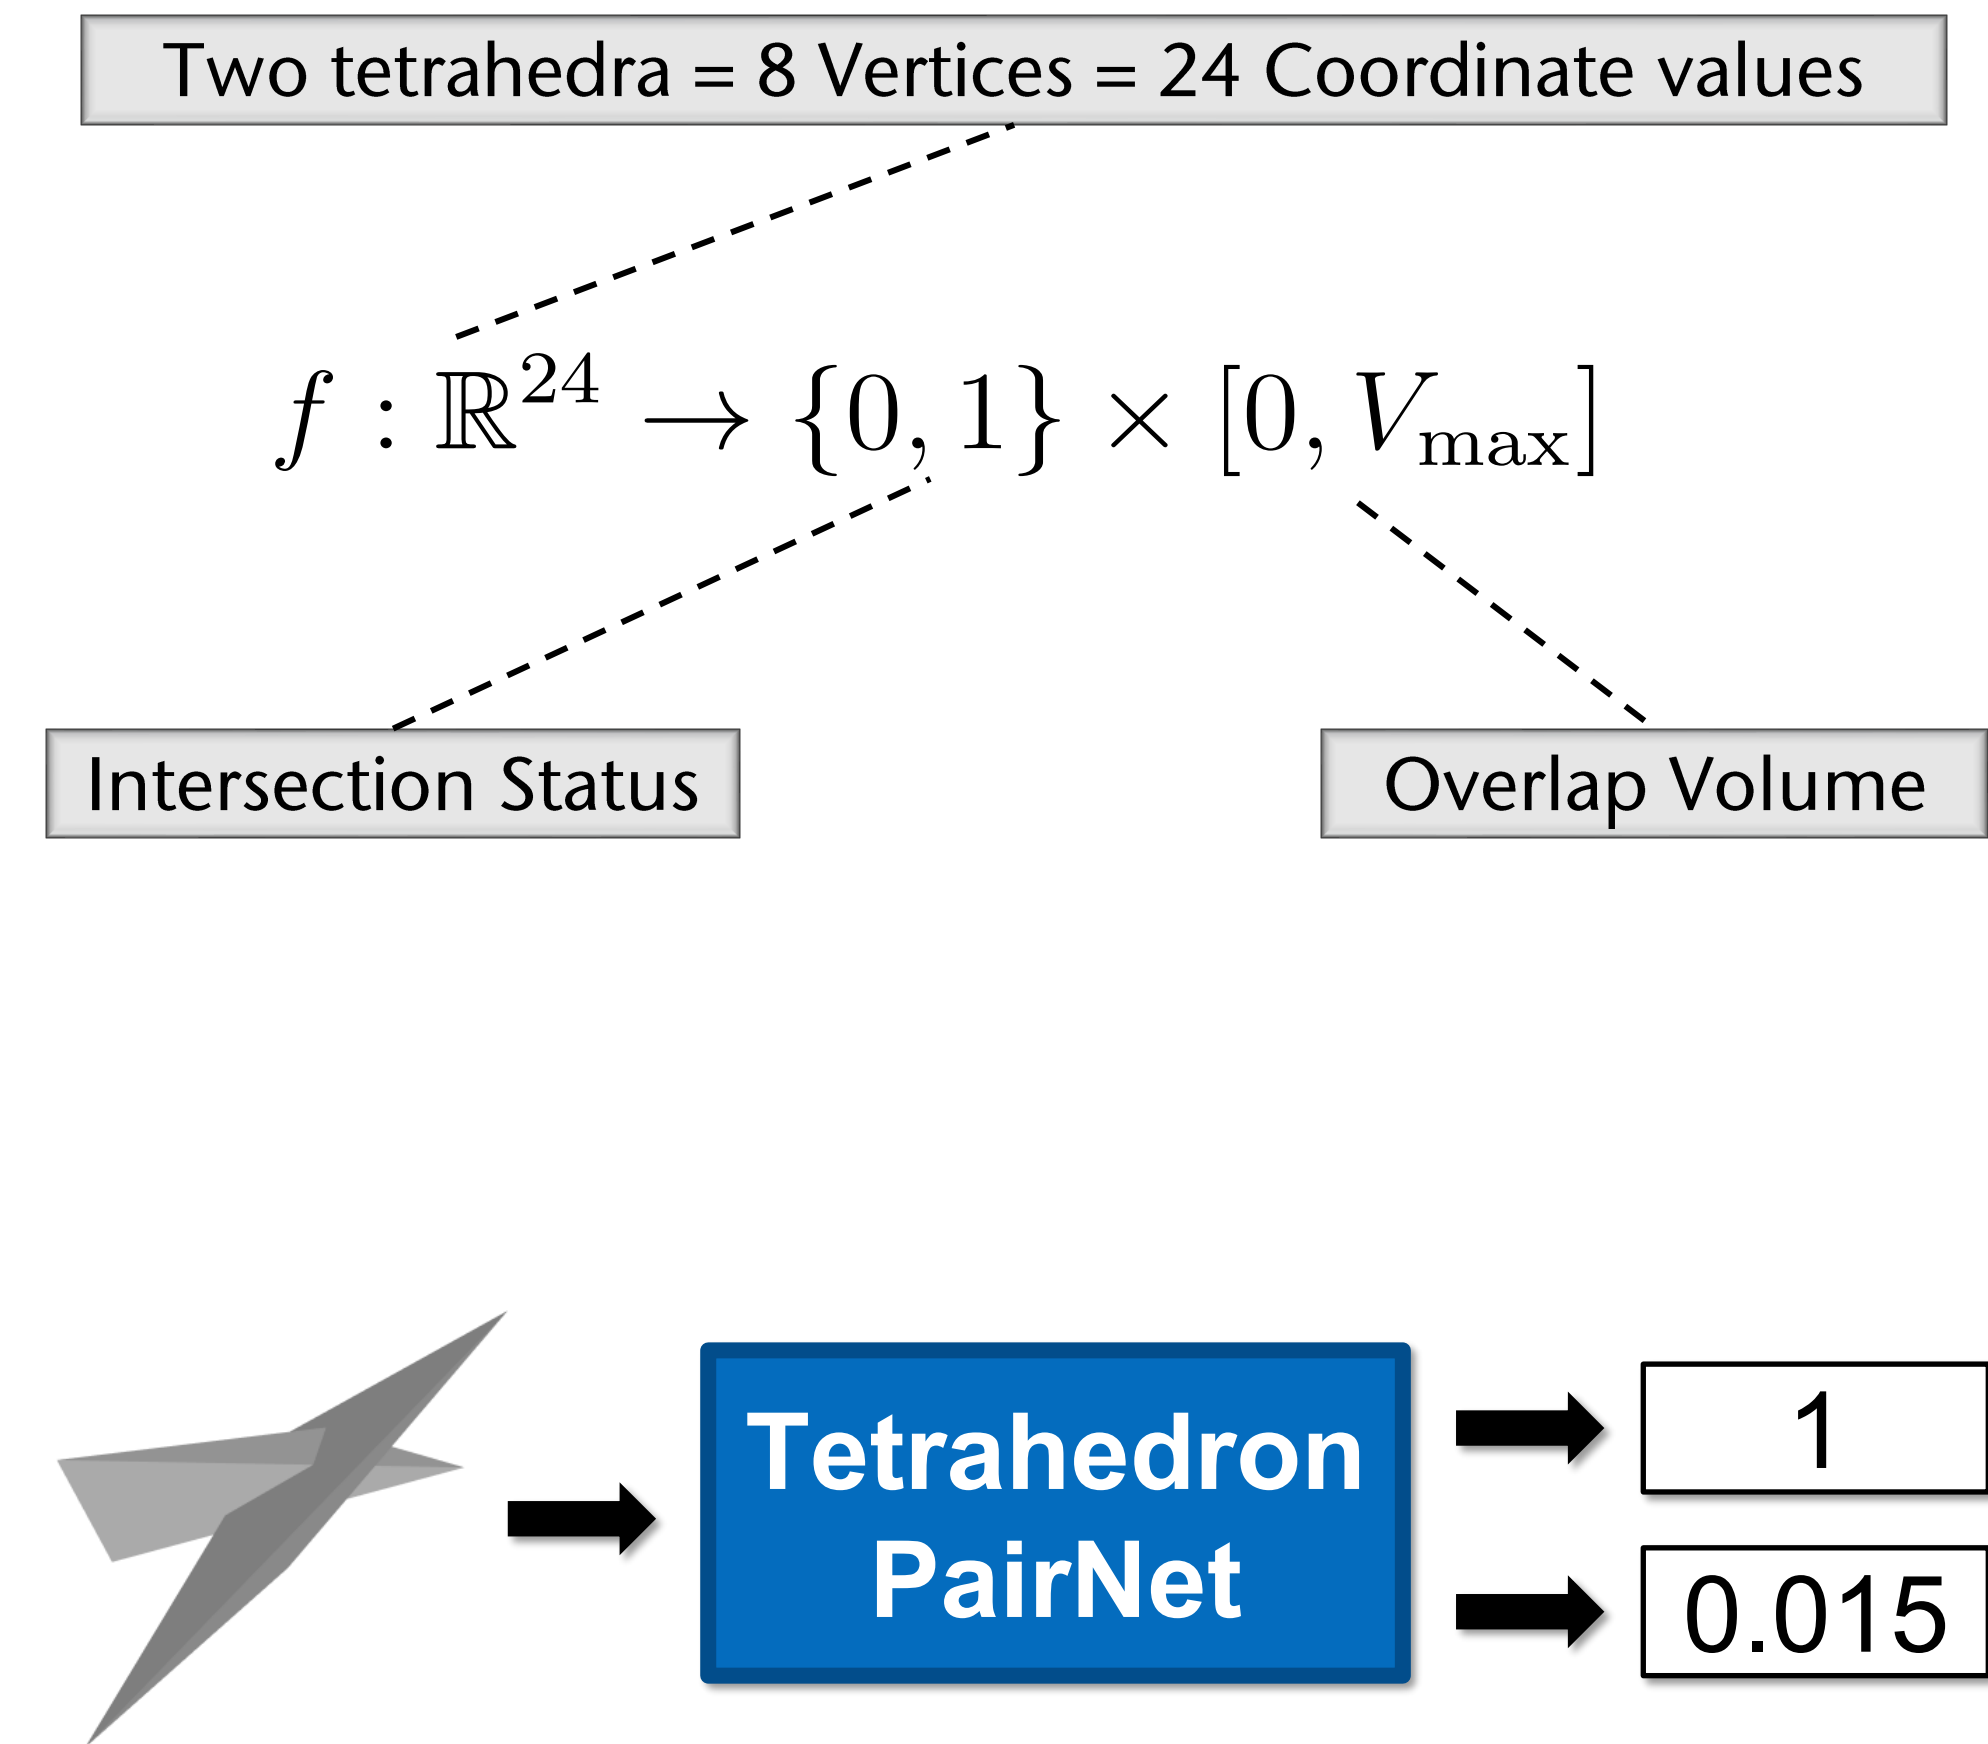
```
pair<bool, simplex> gjk(support_func):
...
if dot(cross(normal, C-A), -A) > 0:
    if dot(C-A, -A) > 0:
        simplex = [C,A]
        D = cross(cross(C-A, -A), C-A)
        return
    else:
        simplex = [A]
        D = -A
        return
else
    ...
...
```



- Exact collision methods: GJK, EPA, SAT overlap tests
- Learning for geometry exists, but not directly for this task
- PointNet and DeepSets provide permutation-invariant set processing
- Our work: Joint tetrahedron intersection & overlap volume prediction

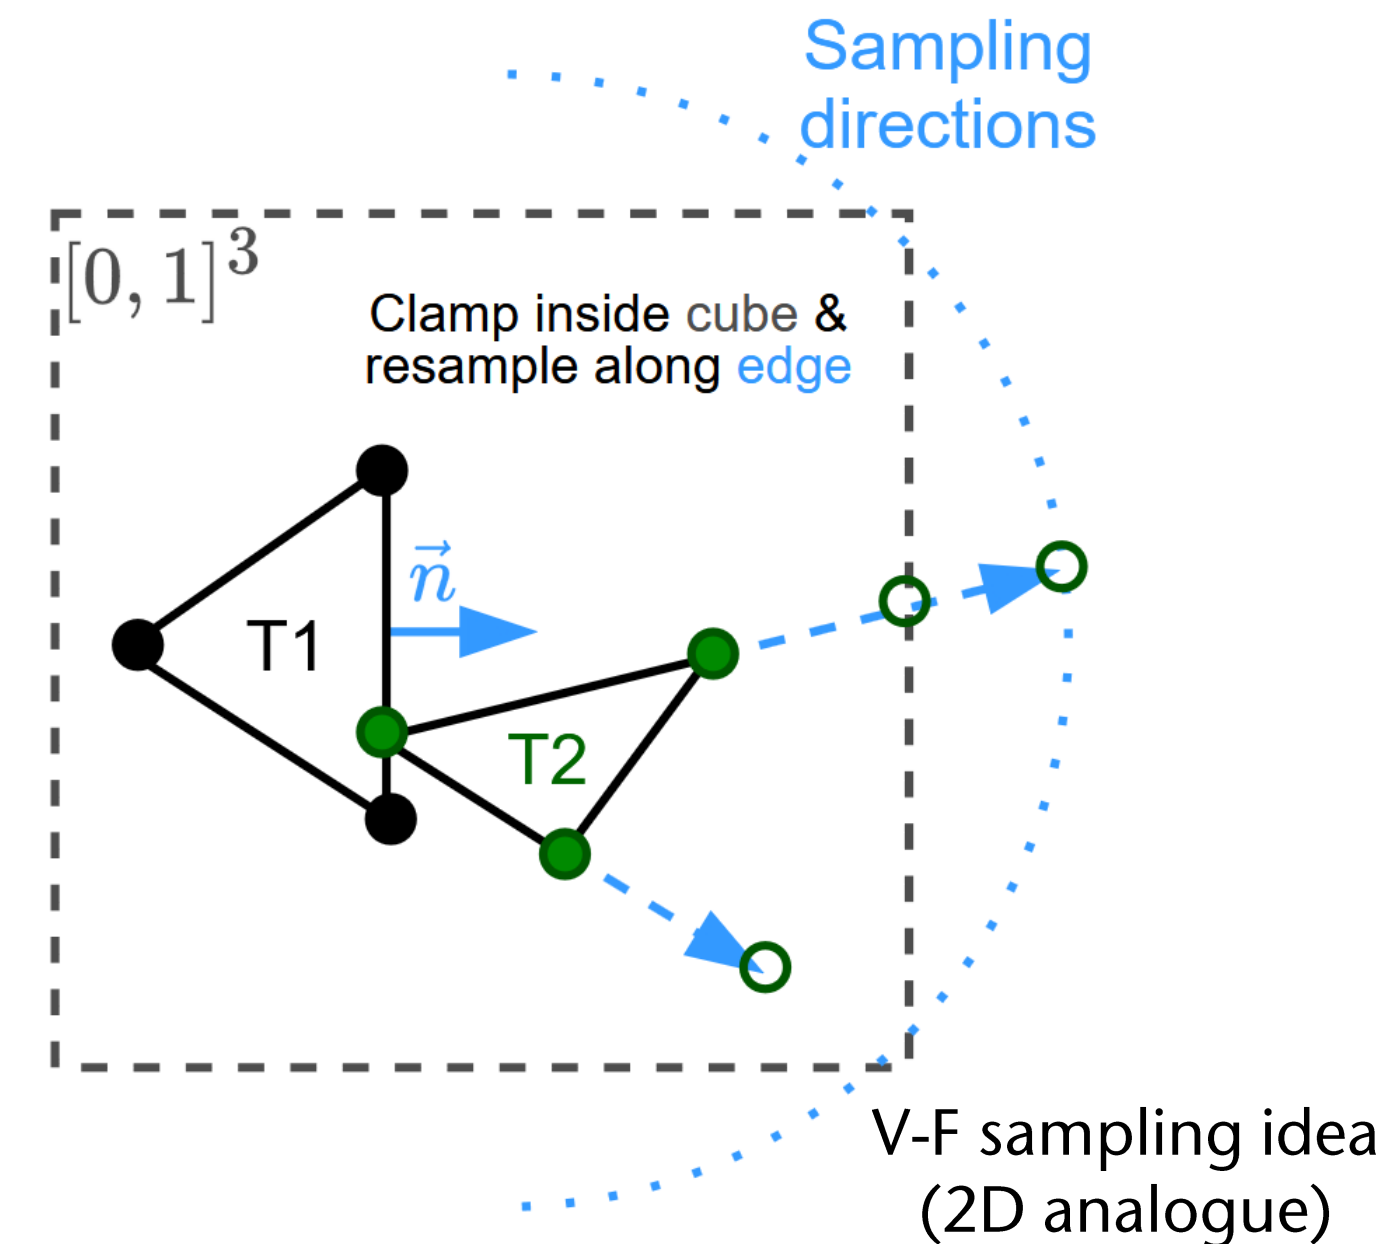
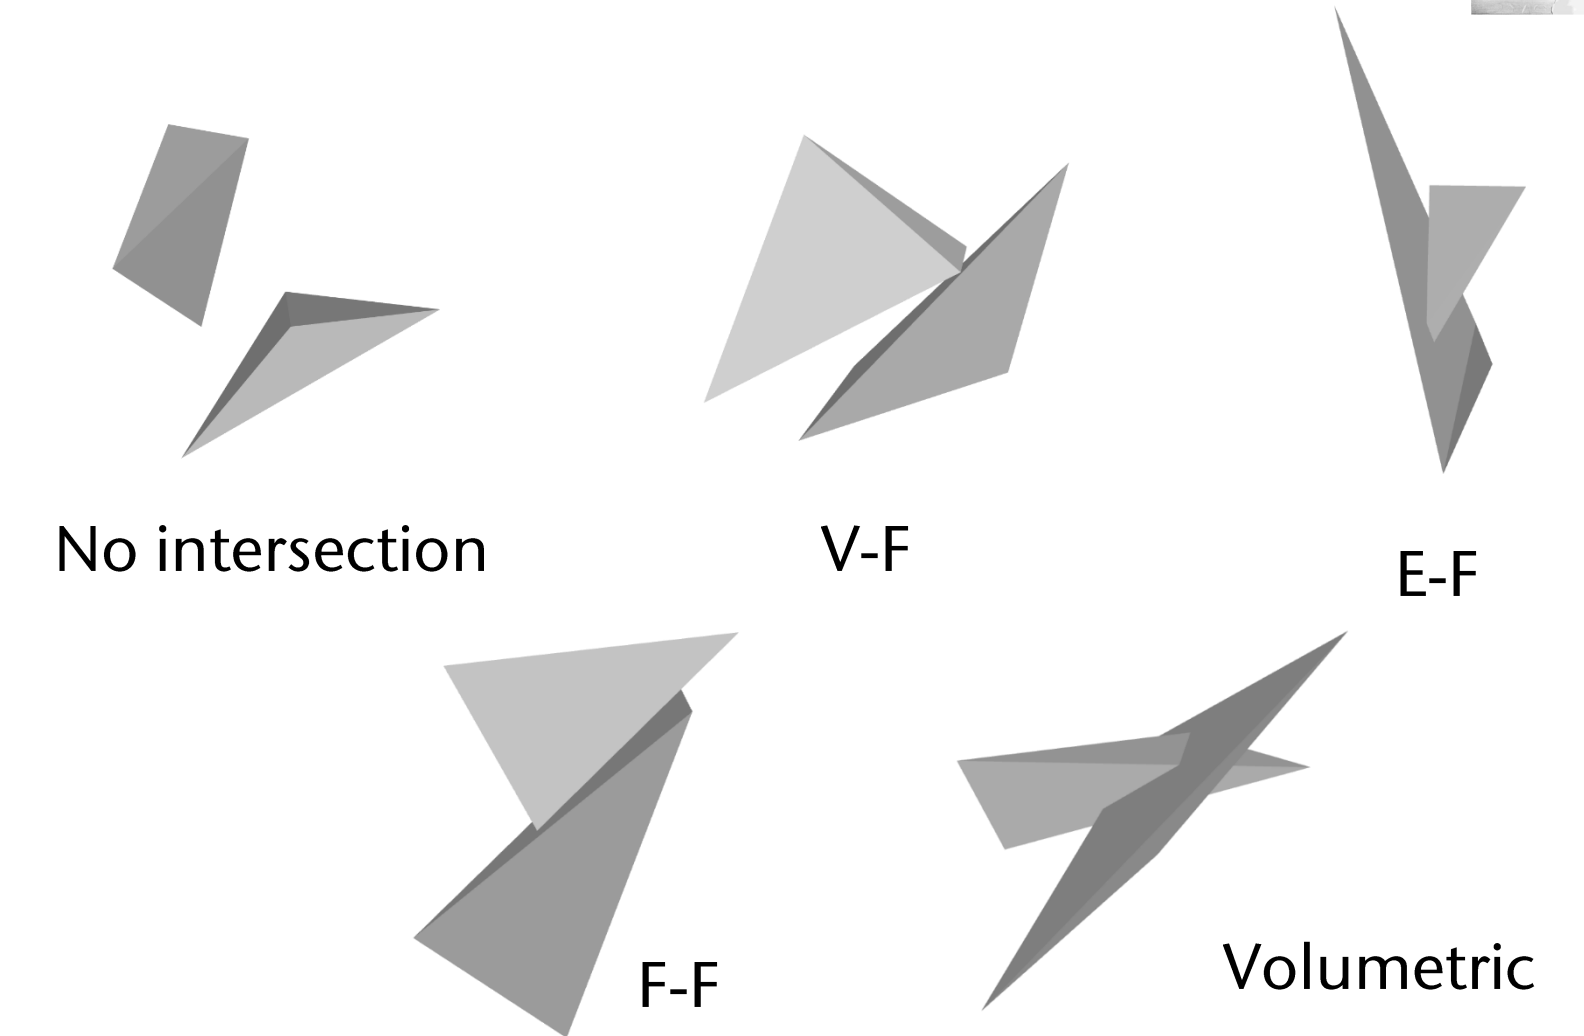


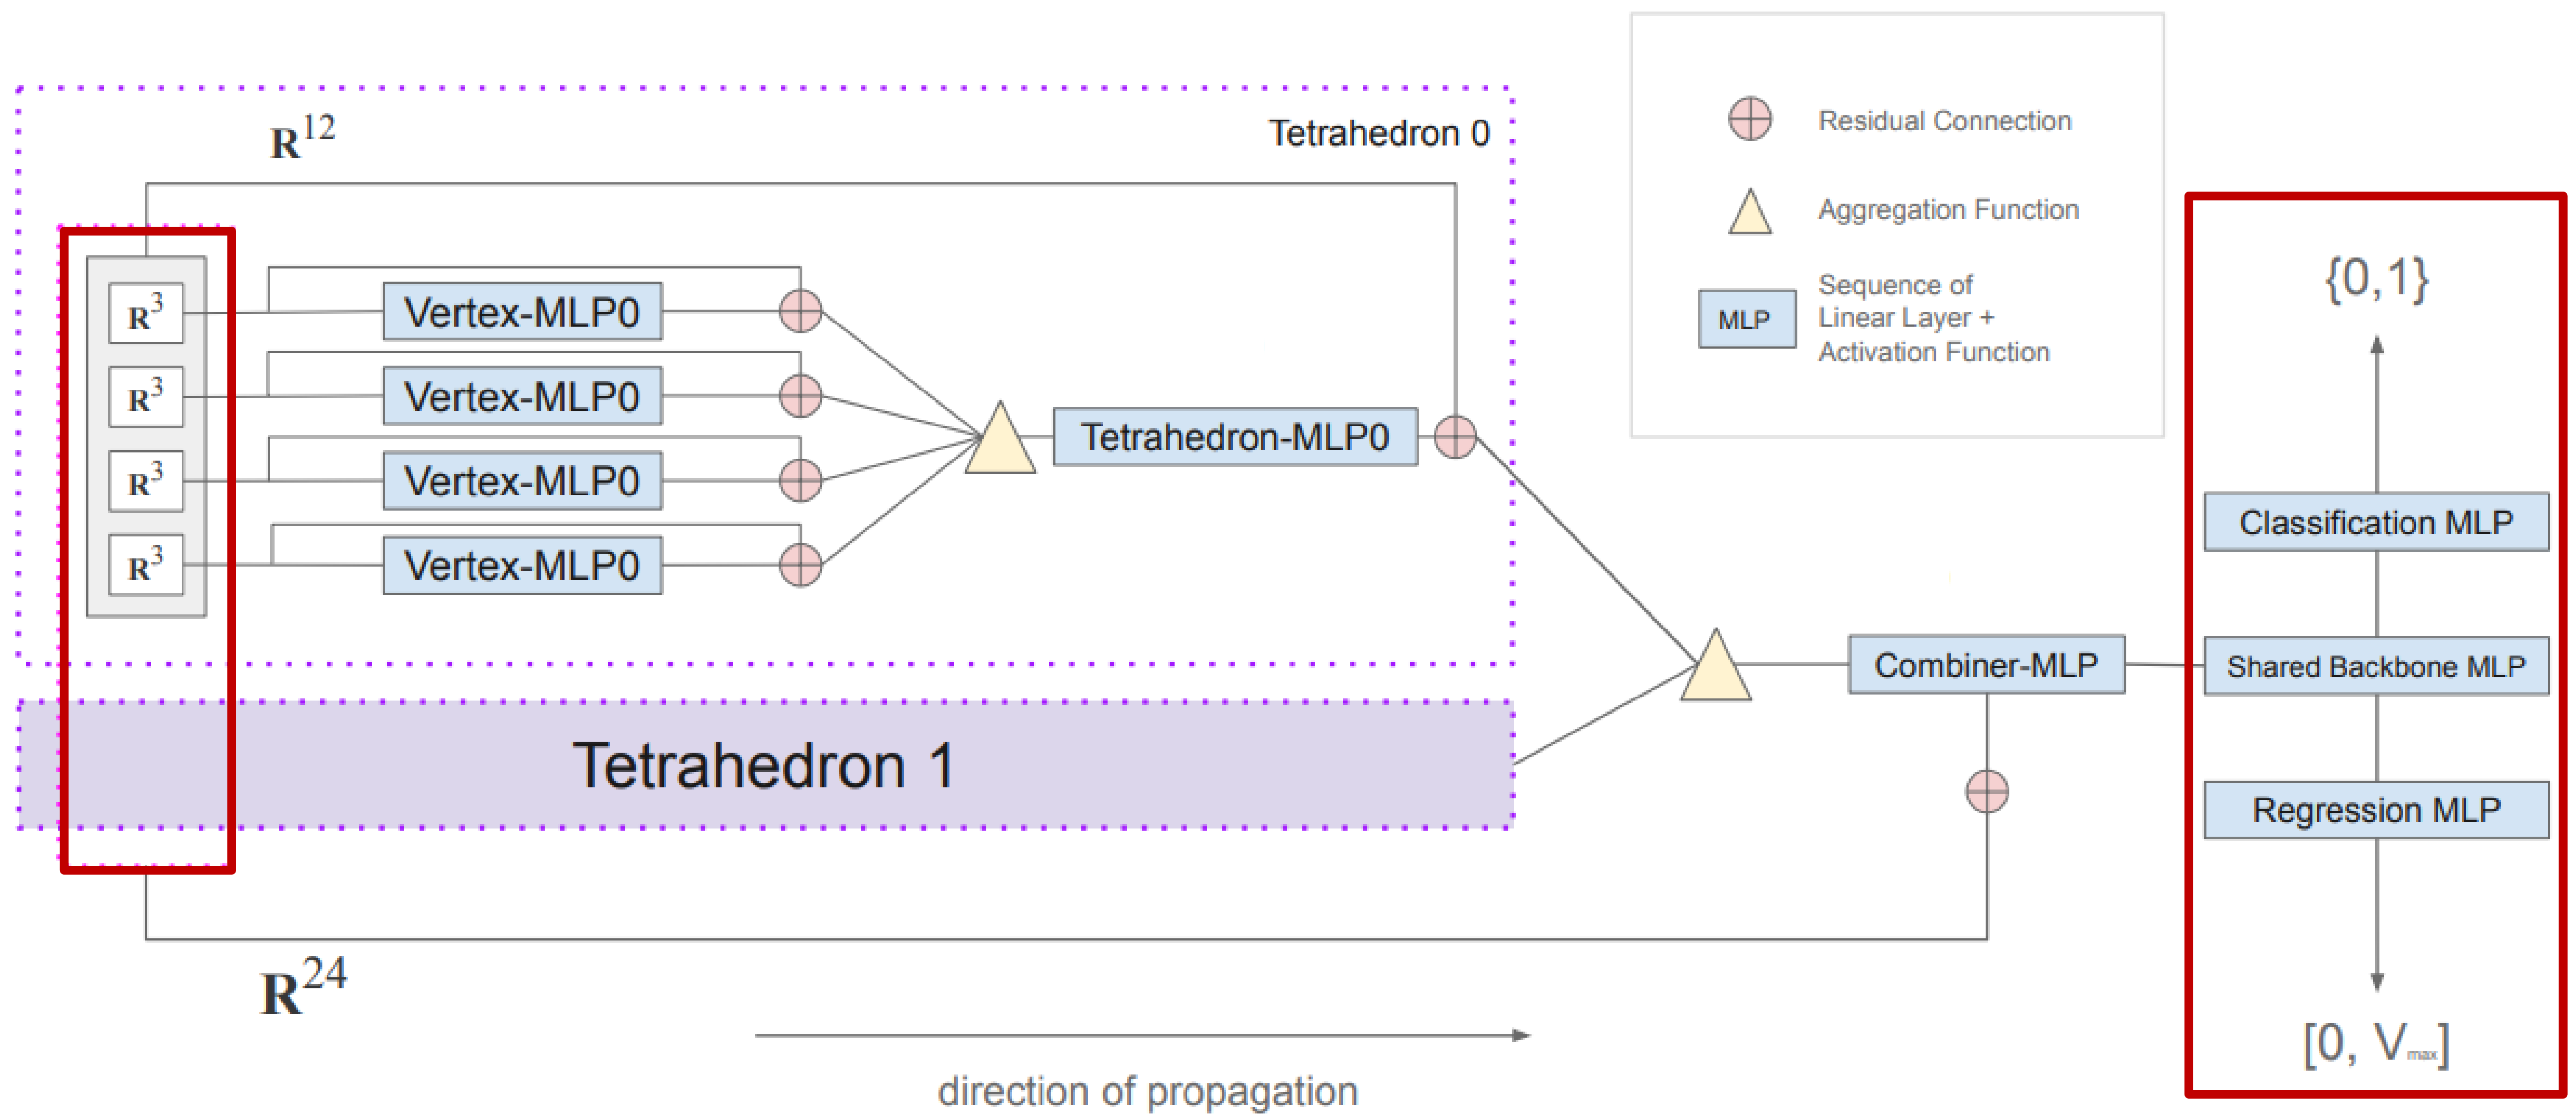
- Reframe tetrahedron intersection as a learned inference problem
- Permutation-invariant tetrahedron-pair network
  - Jointly predict: Binary intersection and volumetric overlap
- Synthetic data generation for several contact topologies
- 99.3% accuracy and 93x / 16,821x speedups



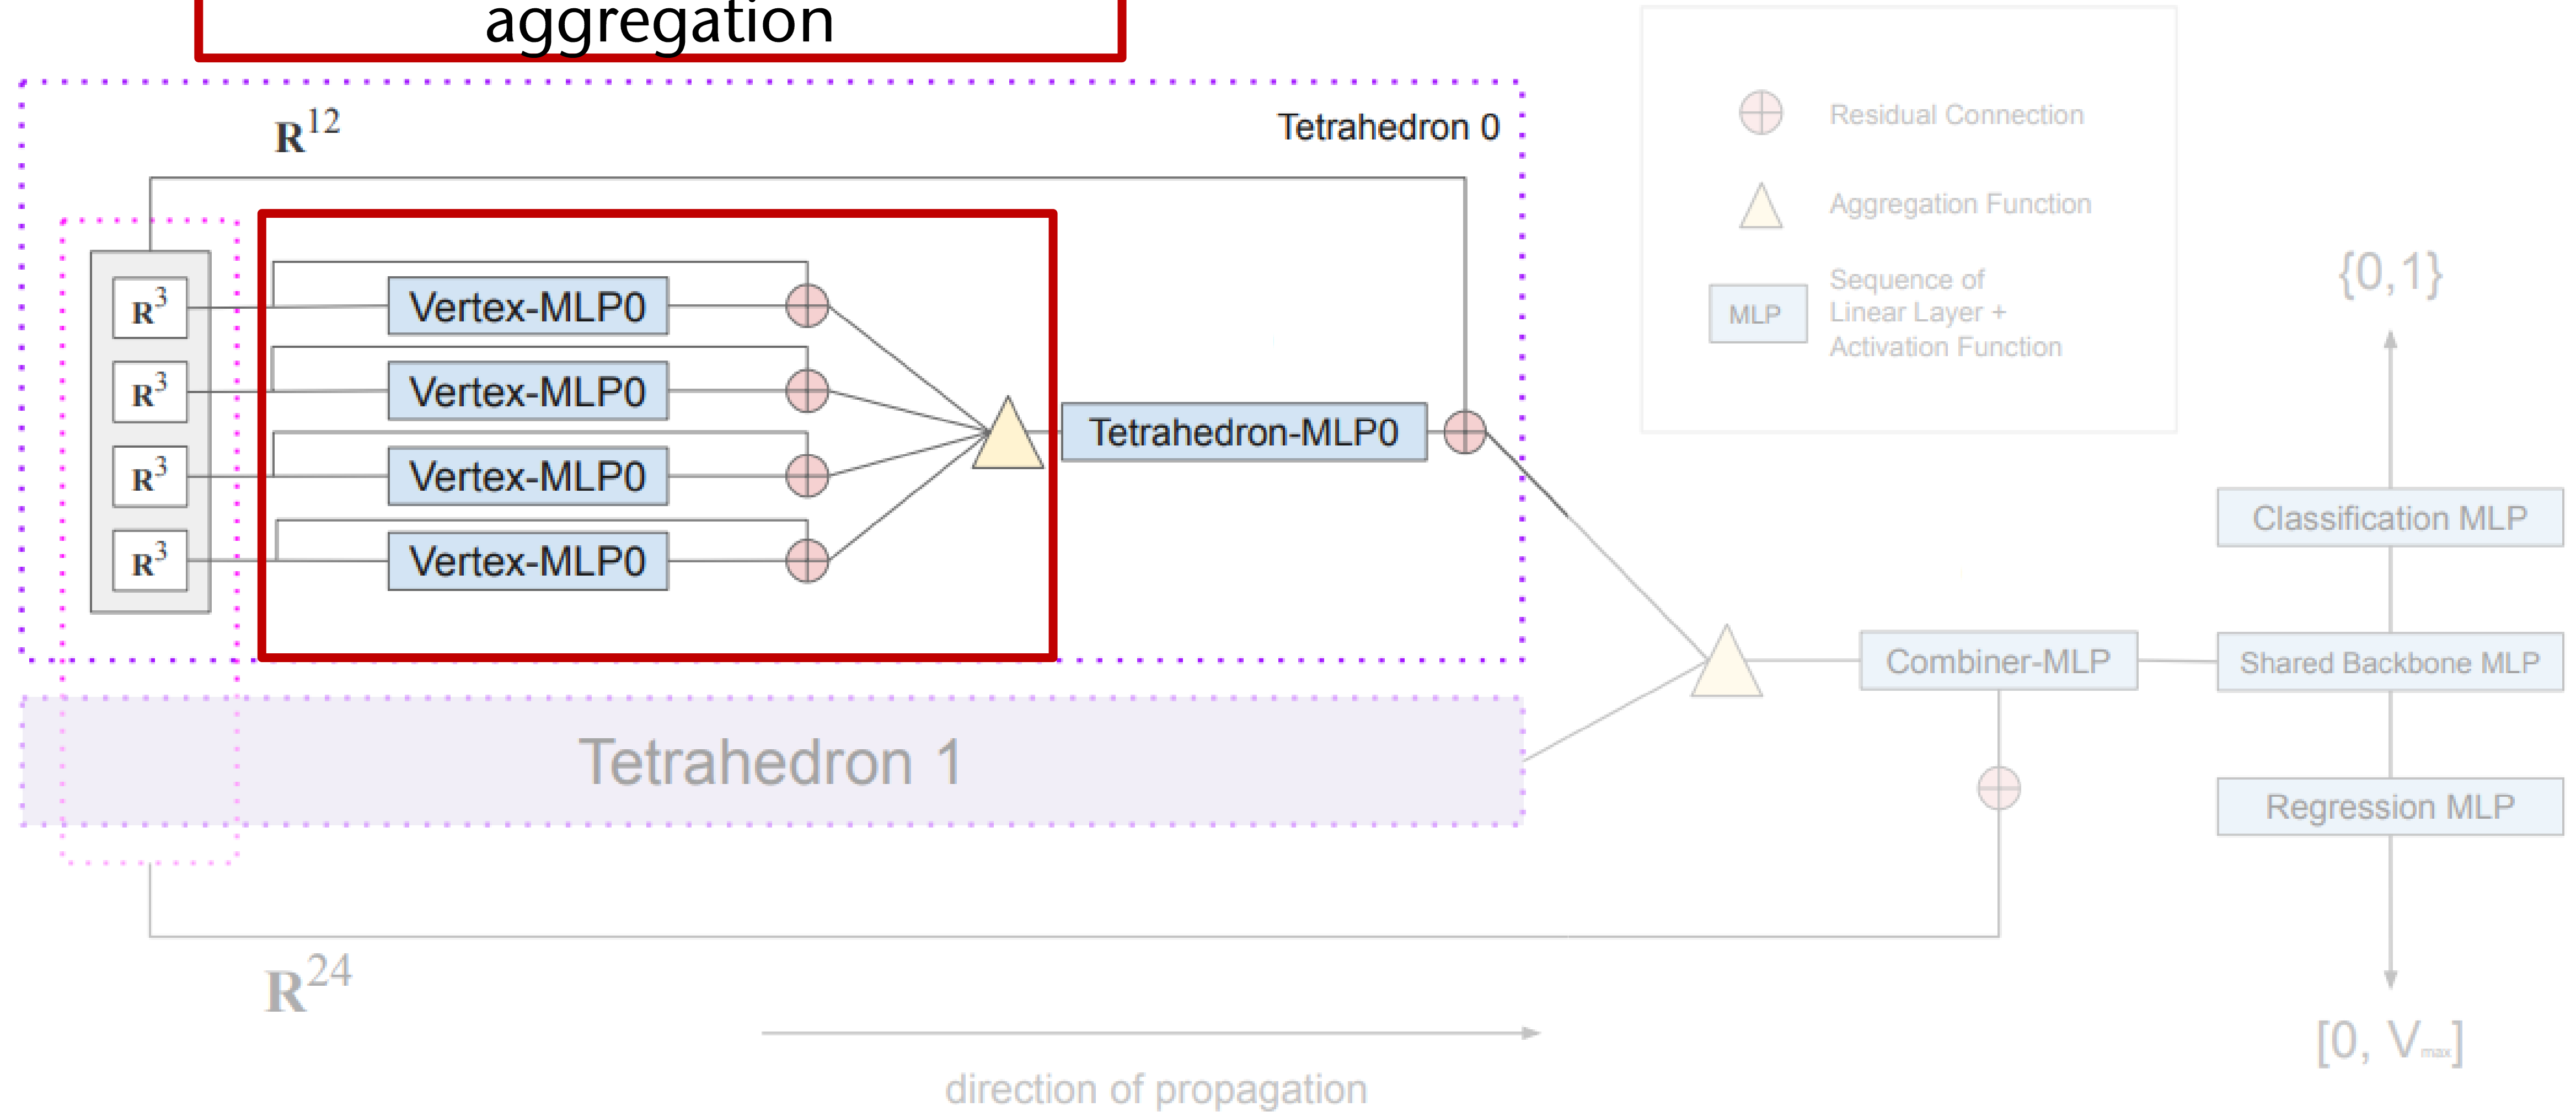
# Training Dataset Generation

- We use CGAL as exact ground truth
- Dataset generation algorithms based on
  - Rejection sampling for no-intersection and volumetric cases
  - Geometric constraints for other contact cases
- Stratified dataset avoids trivial zero-volume bias
- 12M training samples

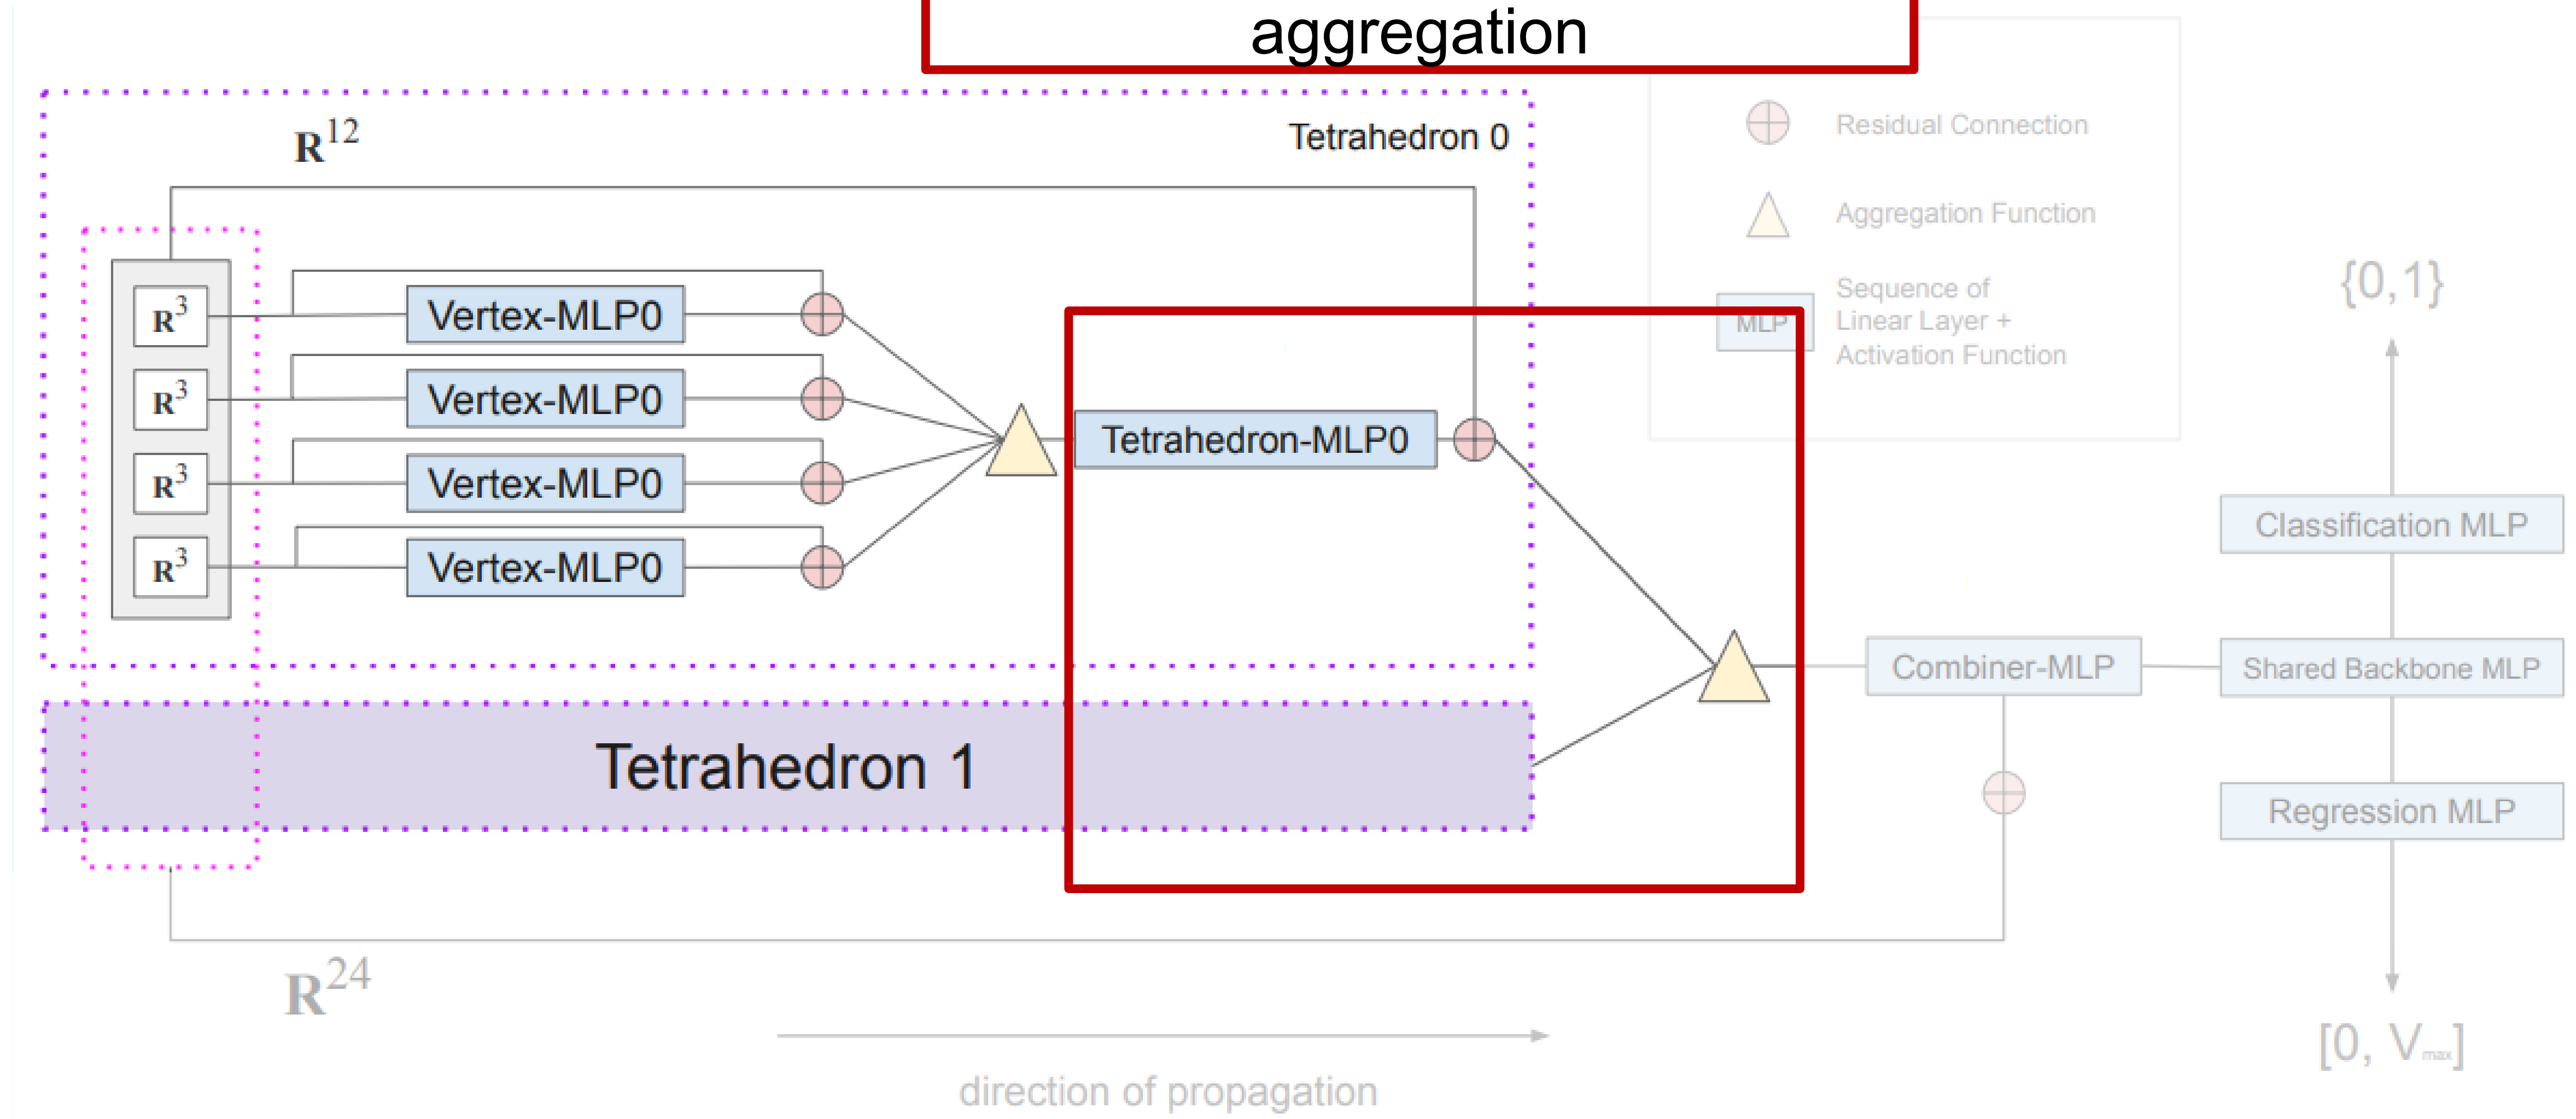




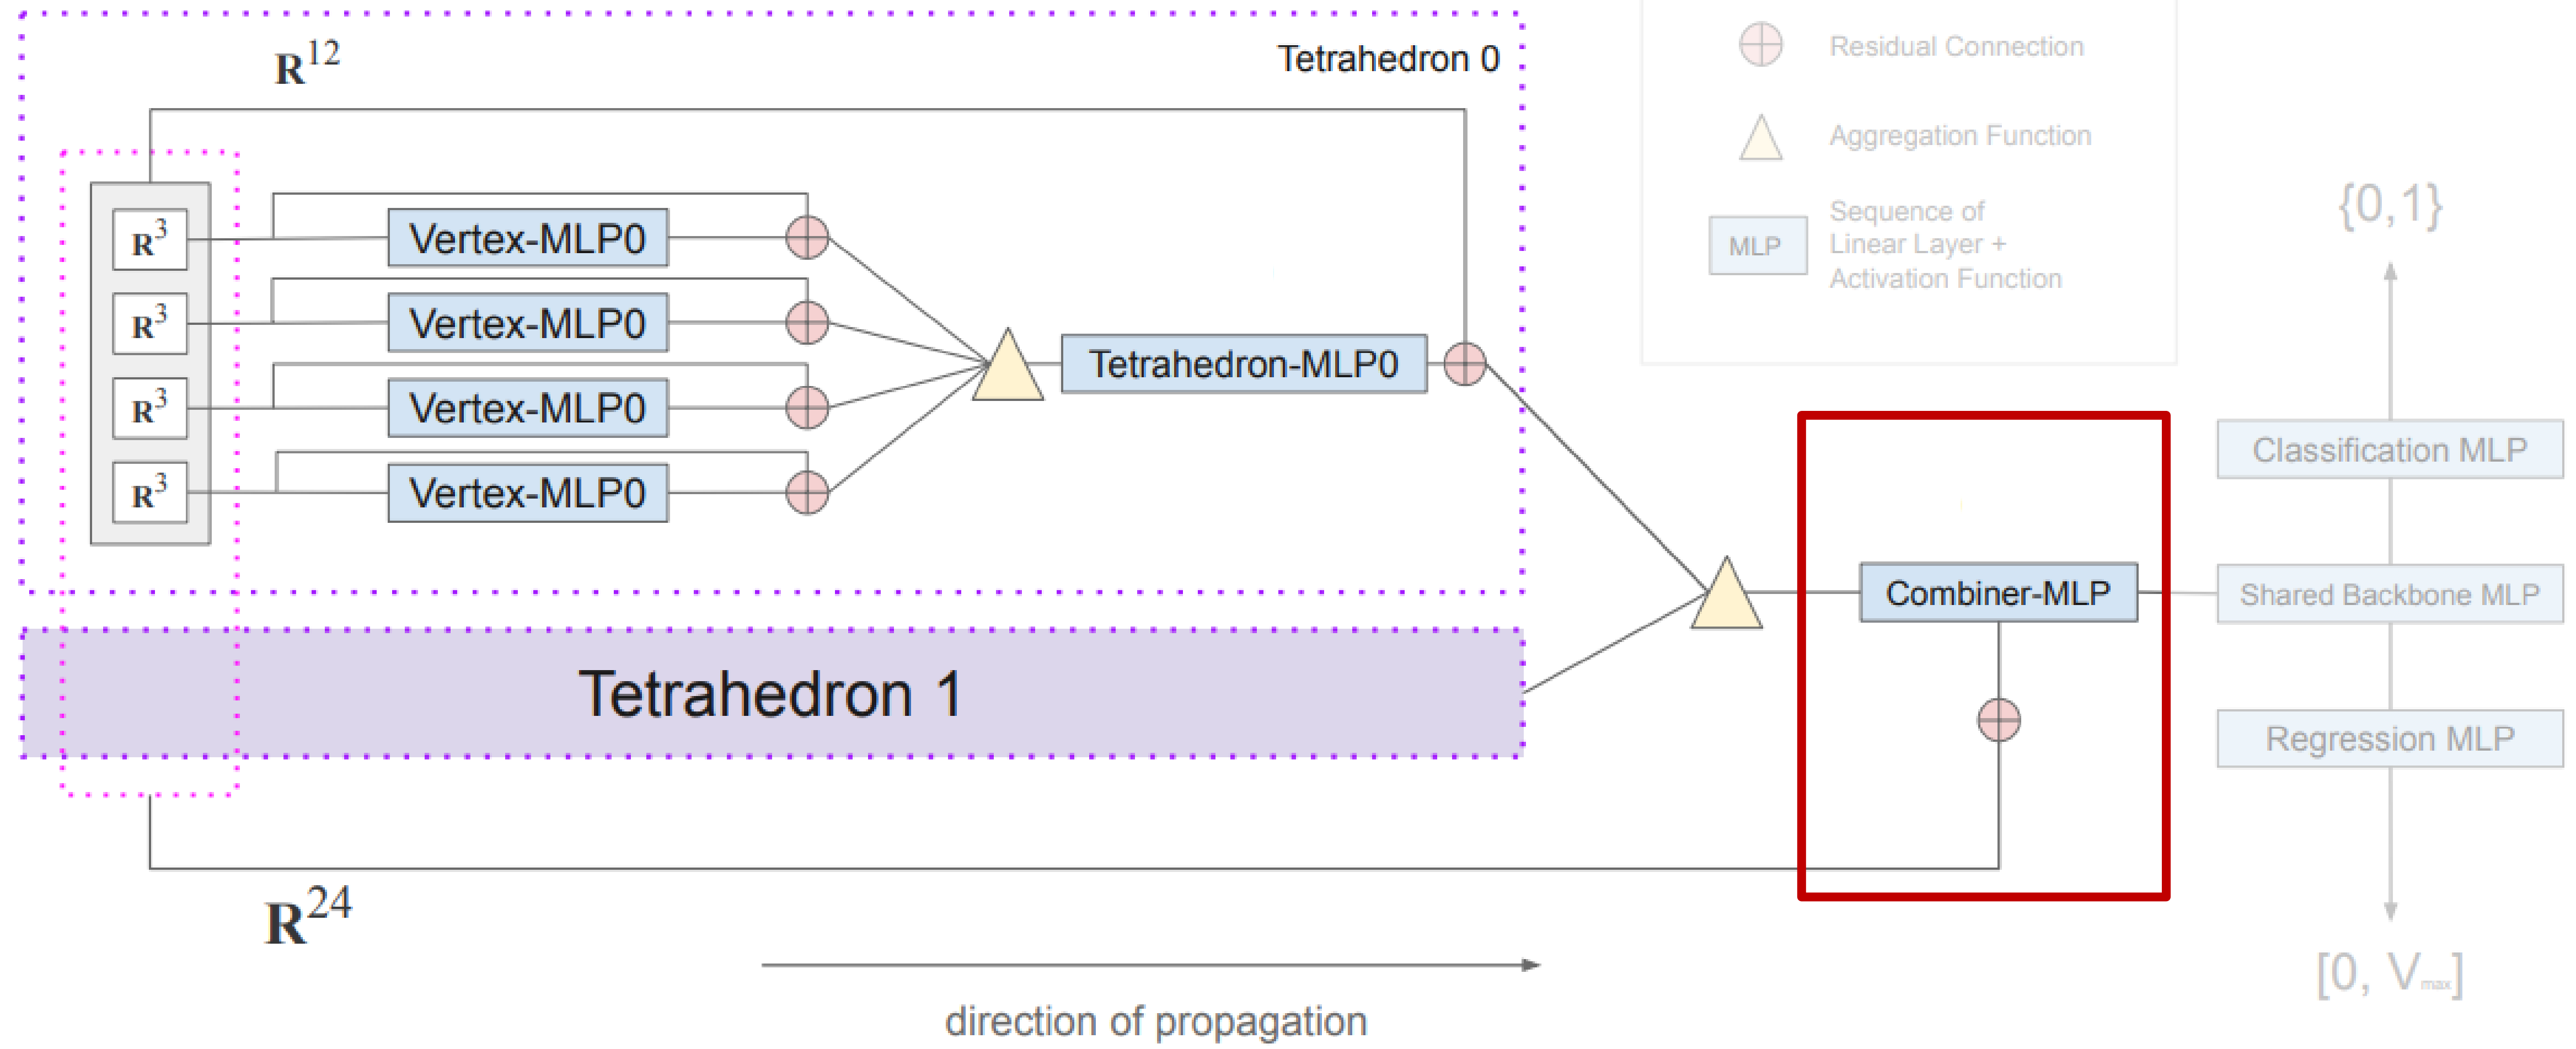
Vertex-level shared MLP + aggregation

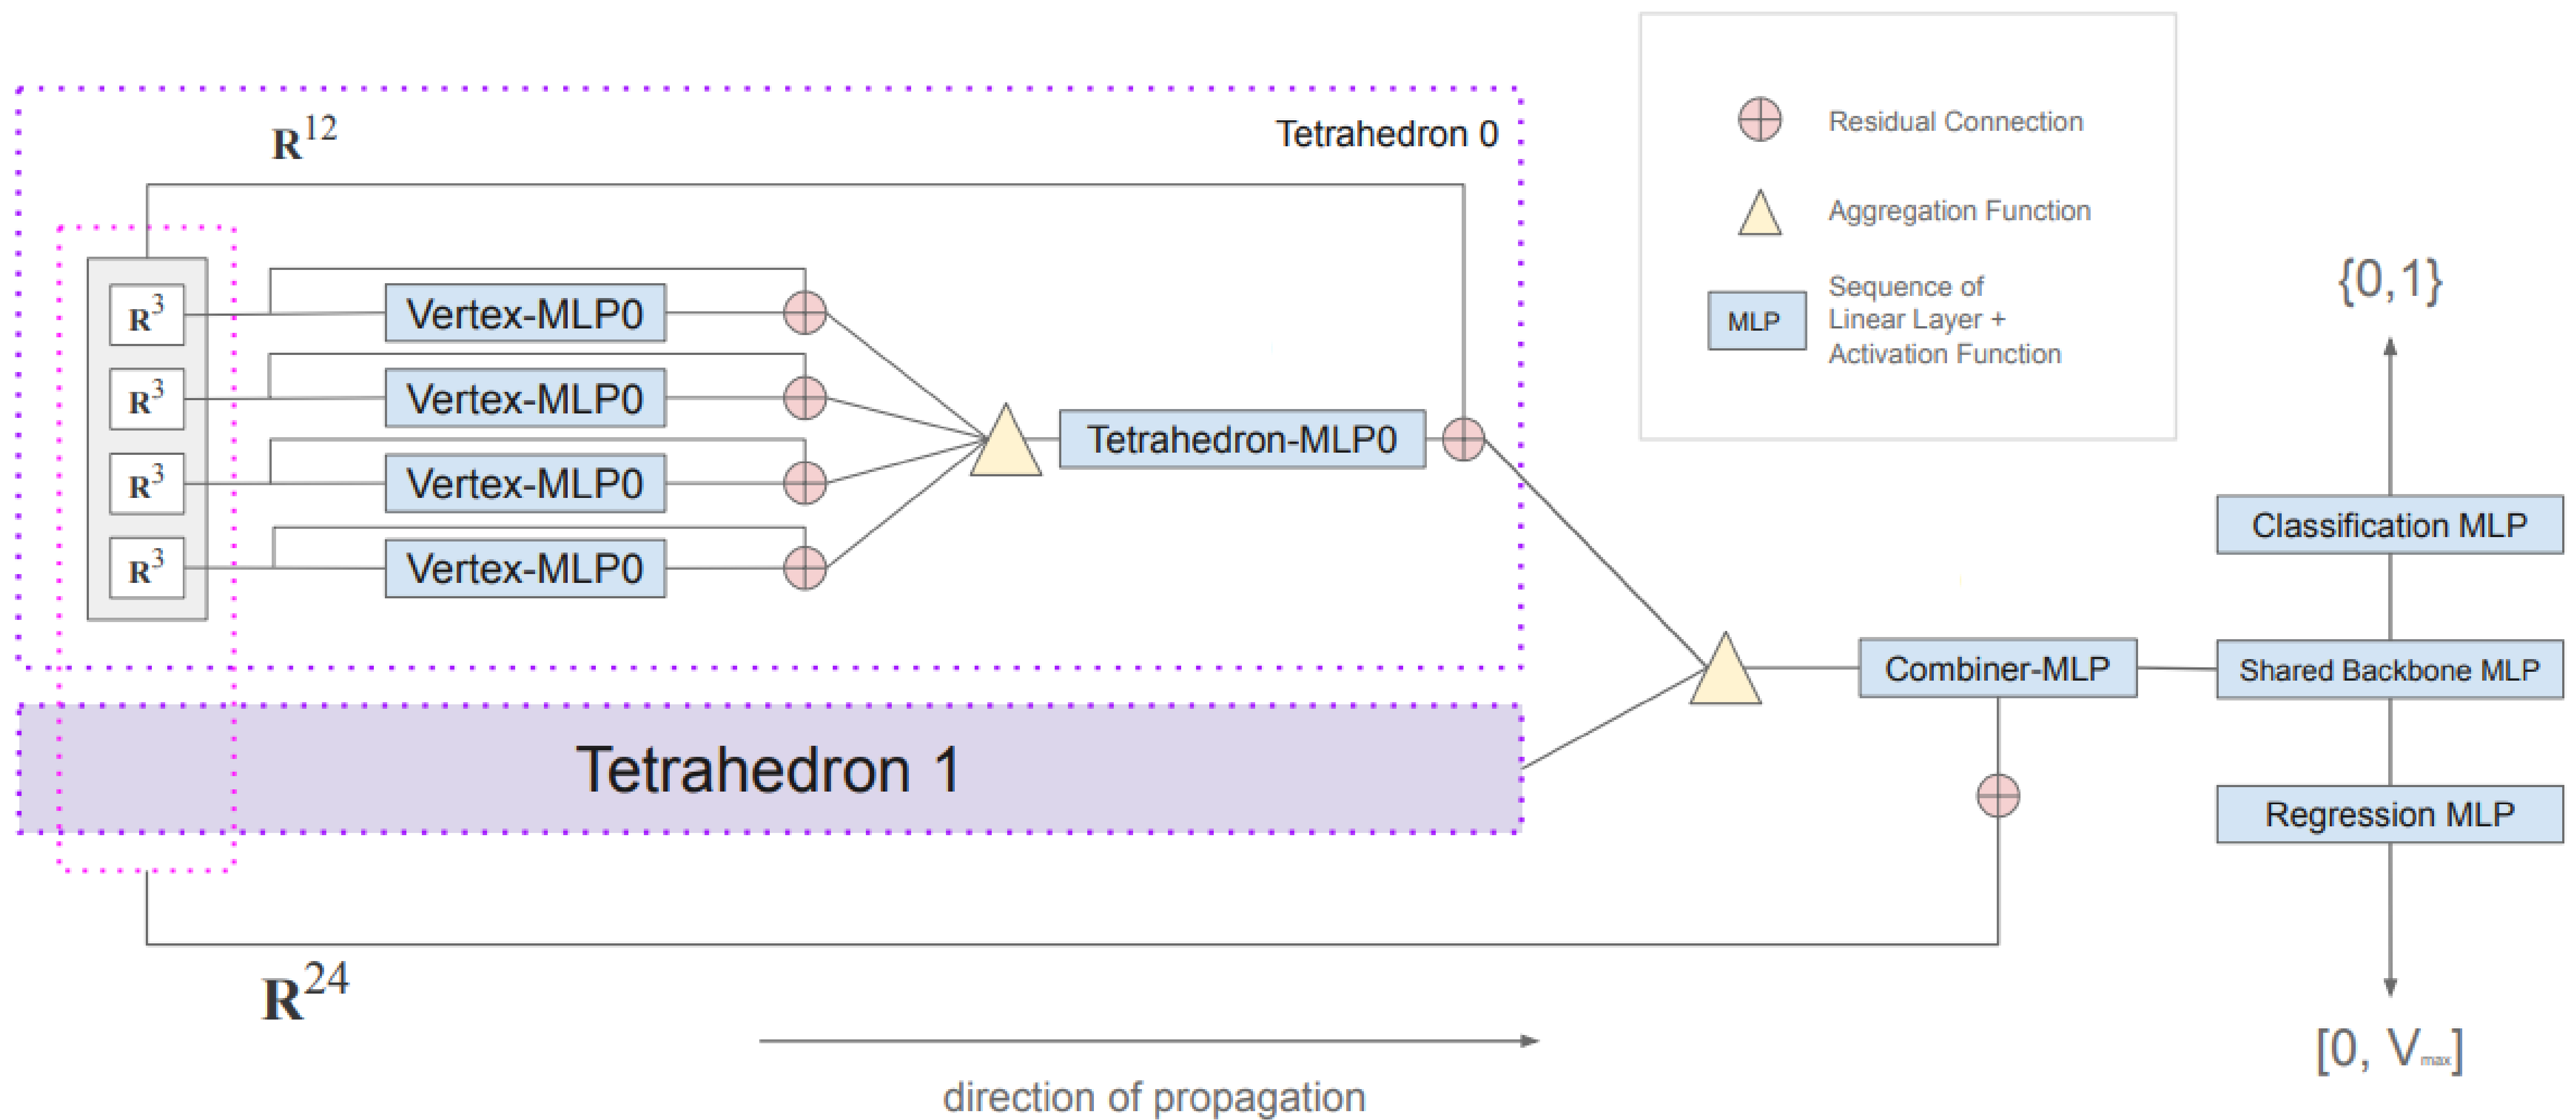


Tetrahedron-level MLP + aggregation




Tetrahedron embeddings combined with 24D input



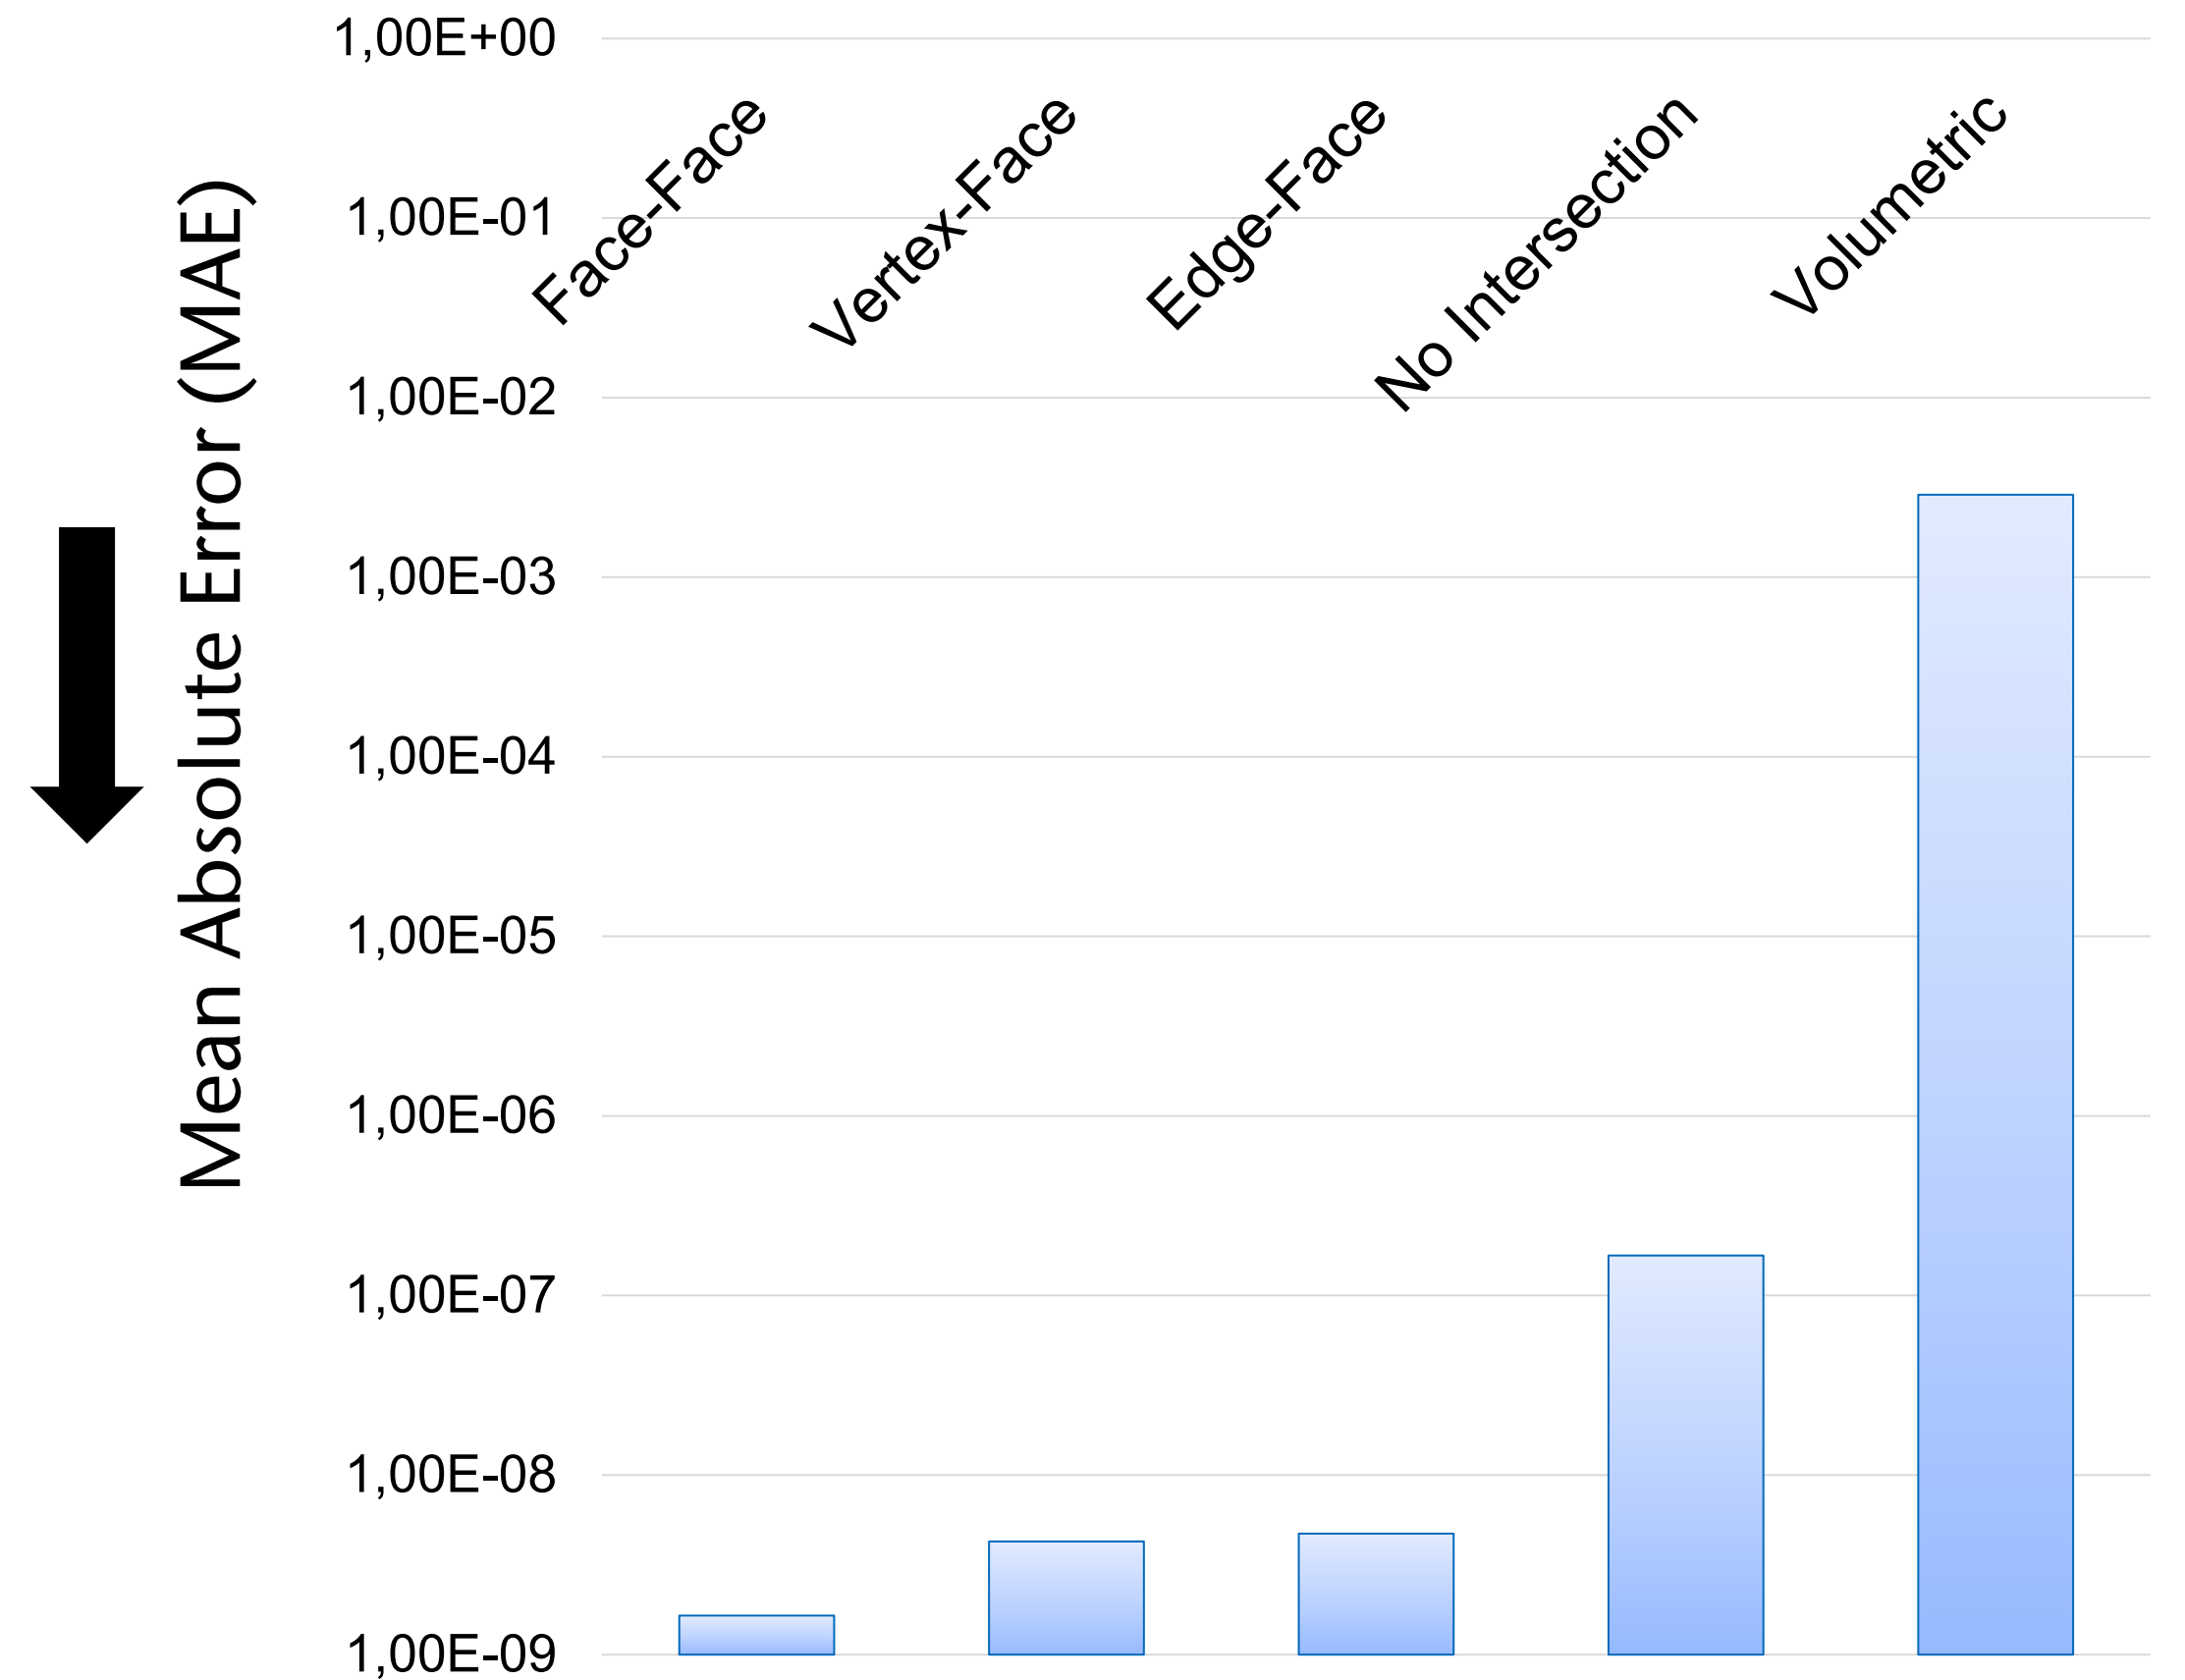


### Collision Accuracy by Case



Mean collision accuracy: 99.37%

### Volume MAE by Case ( $0 < V < \frac{1}{3}$ )



Mean tetrahedron volume in unit cube:  $1.38 \times 10^{-2}$

Volume MAE:  $2.88 \times 10^{-3}$



# Performance Comparison

Method	Throughput	Latency (Pre. + Inf.)	Speedup Detection	Speedup Volume
CGAL (Detection, CPU)	6,352 samples/s	157 $\mu$ s	Baseline	CGAL::do_intersect
CGAL (+ Volume, CPU)	35 samples /s	28.31 ms	Baseline	CGAL::Nef_polyhedron_3
TetrahedronPairNet (CPU, BS=2048)	271,313 samples/s	3.69 $\mu$ s (1.80 + 1.89)	<b>43×</b>	<b>7,751×</b>
TetrahedronPairNet (GPU, BS=2048)	588,744 samples/s	1.70 $\mu$ s (1.11 + 0.59)	<b>93×</b>	<b>16,821×</b>


Large speedups on CPU and GPU

Note: CGAL is exact; TetrahedronPairNet is approximate



- Current limitations:
  - Missing training cases
  - Volume regression degrades for deeper penetration
- Best deployment model:
  - Use as GPU-resident pre-filter
  - Fall back to exact predicates for hard / low confidence cases
- Take-home message:
  - Tet–tet **intersection** and **approximate volume** can be predicted jointly by one massively parallel learned operator



# Thank You for Your Attention!



GitHub



erendropedro@gmail.com  
{meissenhelter, zach}@cs.uni-bremen.de

<https://cgvr.cs.uni-bremen.de>

# Training and Architecture Parameters

Training Configuration		Model Architecture	
<b>Training Samples</b>	12,000,000	<b>Precision</b>	float32
<b>Optimizer</b>	AdamW	<b>Total Parameters</b>	73,370
<b>Learning Rate</b>	$10^{-3}$	<b>Model Size</b>	633 KB
<b>Batch Size</b>	2048	<b>Activation Function</b>	ReLU
<b>Epochs</b>	50	<b>Aggregation Function</b>	Max-Pooling
<b>Loss Function</b>	0.5 BCE + 0.5 RMSLE	<b>Vertex MLP</b>	[24, 24]
<b>Class Balance</b>	0.5 positive/negative	<b>Tetrahedron MLP</b>	[48, 48]
<b>Input Normalization</b>	PCA + Unit	<b>Combiner MLP</b>	[96, 96]
<b>Data Distribution</b>	[50%, 5%, 7%, 20%, 18%]	<b>Shared Backbone</b>	[1024]
<b>Volume Scaling</b>	$10^3$	<b>Classification Head</b>	[128, 64, 32, 1]
		<b>Regression Head</b>	[128, 128, 1]

# Volume Bin Accuracy

Test Subset / Metric	Detect Acc.	Vol. MAE	Notes
<i>Degenerate &amp; Non-Intersecting Cases</i>			
No Intersection	98.02%	$1.67 \times 10^{-7}$	Baseline negative class
Vertex–Face	99.07%	$4.27 \times 10^{-9}$	Correctly predicts $\approx 0$ volume
Edge–Face	99.97%	$4.72 \times 10^{-9}$	Correctly predicts $\approx 0$ volume
Face–Face	99.92%	$1.65 \times 10^{-9}$	Correctly predicts $\approx 0$ volume
<i>Volumetric Intersection Analysis</i>			
<b>Volumetric (Overall)</b>	<b>99.07%</b>	<b><math>2.88 \times 10^{-3}</math></b>	<b>Volumetric collision case</b>
<i>Small</i> ( $v < 0.0067$ )	—	$1.47 \times 10^{-3}$	88% Bin Acc. (34% of samples)
<i>Medium</i> ( $0.0067 \leq v < 0.0133$ )	—	$2.80 \times 10^{-3}$	60% Bin Acc. (33% of samples)
<i>Large</i> ( $0.0133 \leq v < 0.0200$ )	—	$4.44 \times 10^{-3}$	41% Bin Acc. (33% of samples)

Strategy	Mechanism & Properties
<b>Principal Axis Alignment (PCA)</b>	<p><b>Mechanism:</b> Applies a rigid rotation derived from the PCA of <math>\mathcal{T}_1</math> to both tetrahedra, aligning <math>\mathcal{T}_1</math>'s principal axes with the global basis.</p> <p><b>Properties:</b> Standardizes orientation while strictly preserving distances, scale, and volume.</p>
<b>Unitary Tetrahedron Alignment (Unit)</b>	<p><b>Mechanism:</b> Applies an affine transformation that maps <math>\mathcal{T}_1</math> to the unit simplex <math>\{(0, 0, 0), (1, 0, 0), (0, 1, 0), (0, 0, 1)\}</math> and expresses <math>\mathcal{T}_2</math> in this new frame.</p> <p><b>Properties:</b> Preserves boolean intersection status; intersection volume is scaled linearly by the determinant of the affine map.</p>

### TetrahedronPairNet: Accuracy Heatmap

