SIEVE

Overcoming Topological Obstruction in Equivariant Self-Supervised Learning

Jongann Lee*, Sun Woo Park*, Yun Young Choi

SolverX · University of Illinois Urbana-Champaign · Max Planck Institute for Mathematics

*Equal contribution

NeurIPS 2026

In 3D, the coordinates are the content

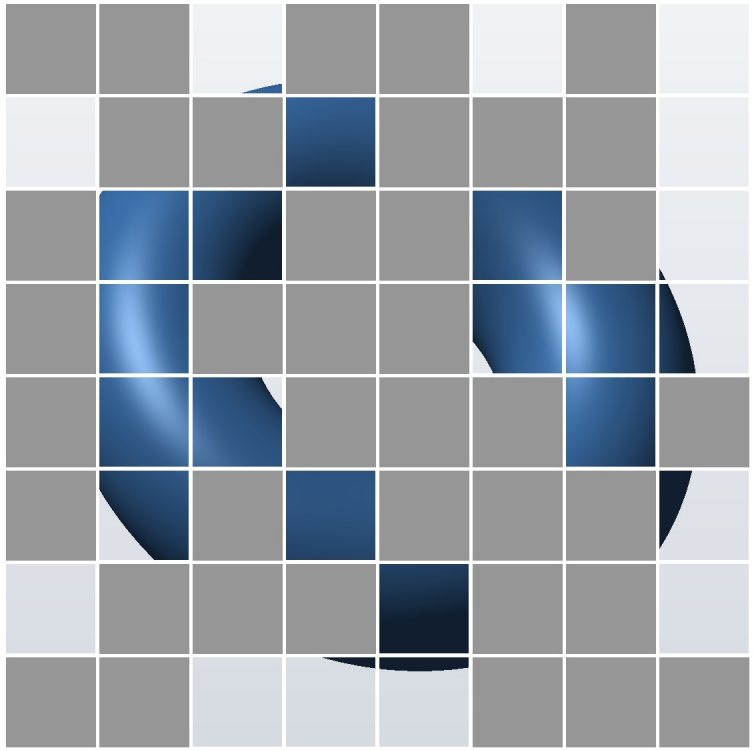


Image: the mask keeps the position

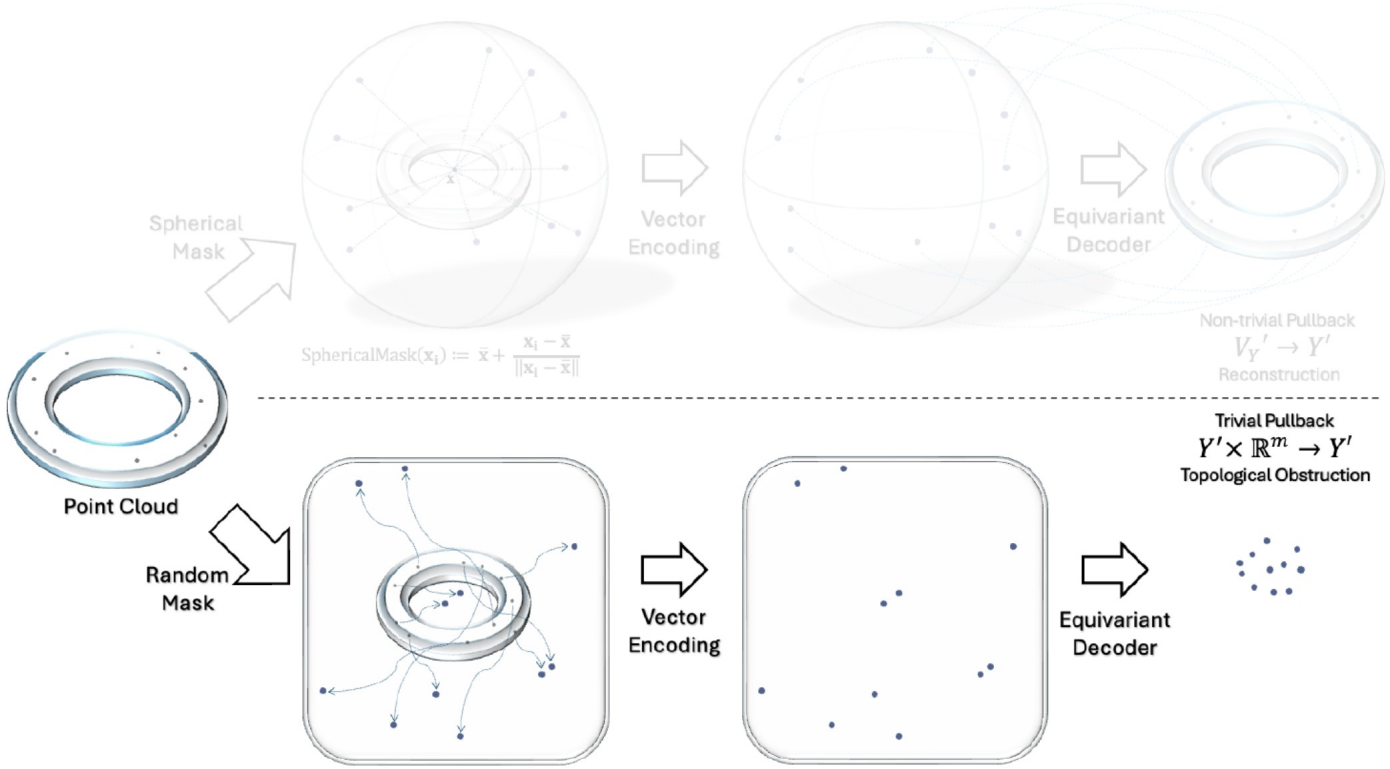


Point cloud: the position is gone too

Theorem 2.6: a topological obstruction

Equivariant networks with constant or random masks cannot approximate non-trivial vector bundles.

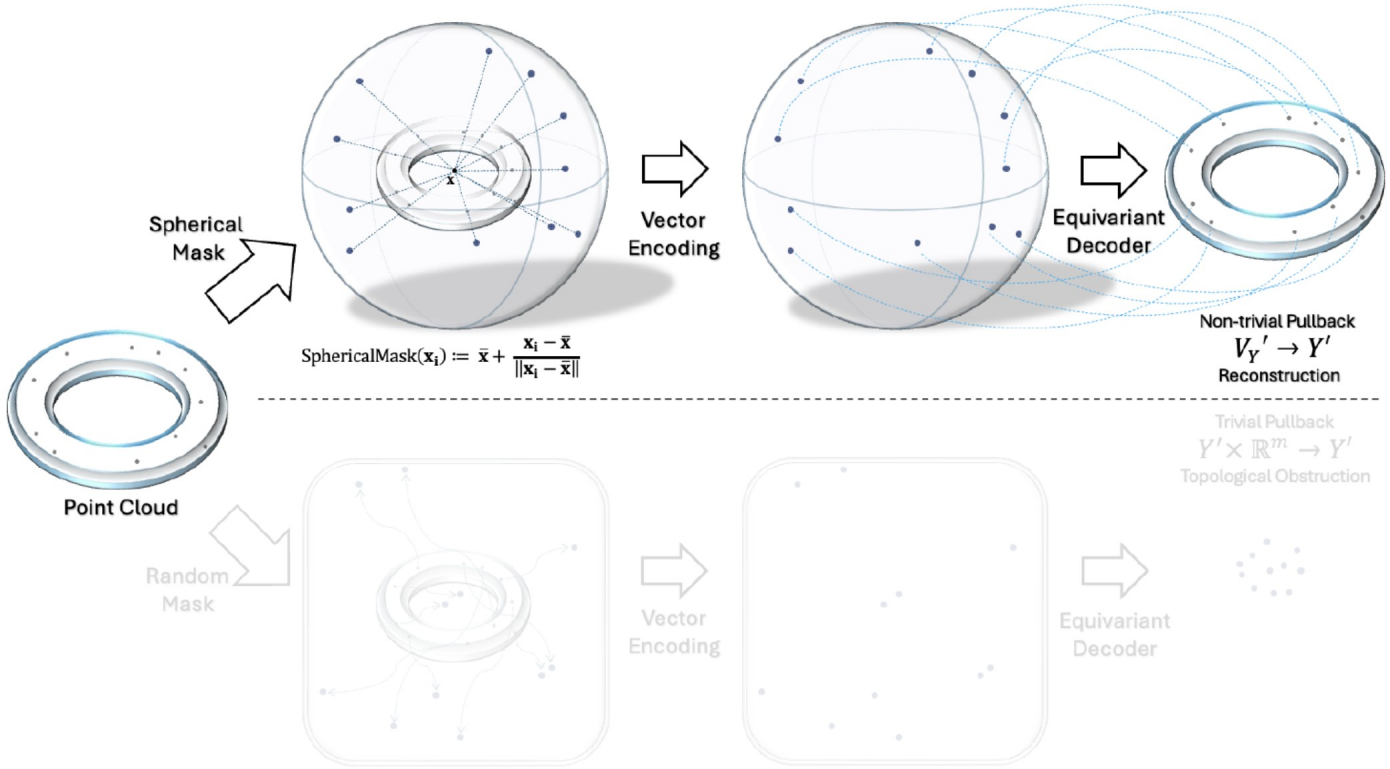
Consequence: such networks cannot reconstruct point clouds of different topologies.



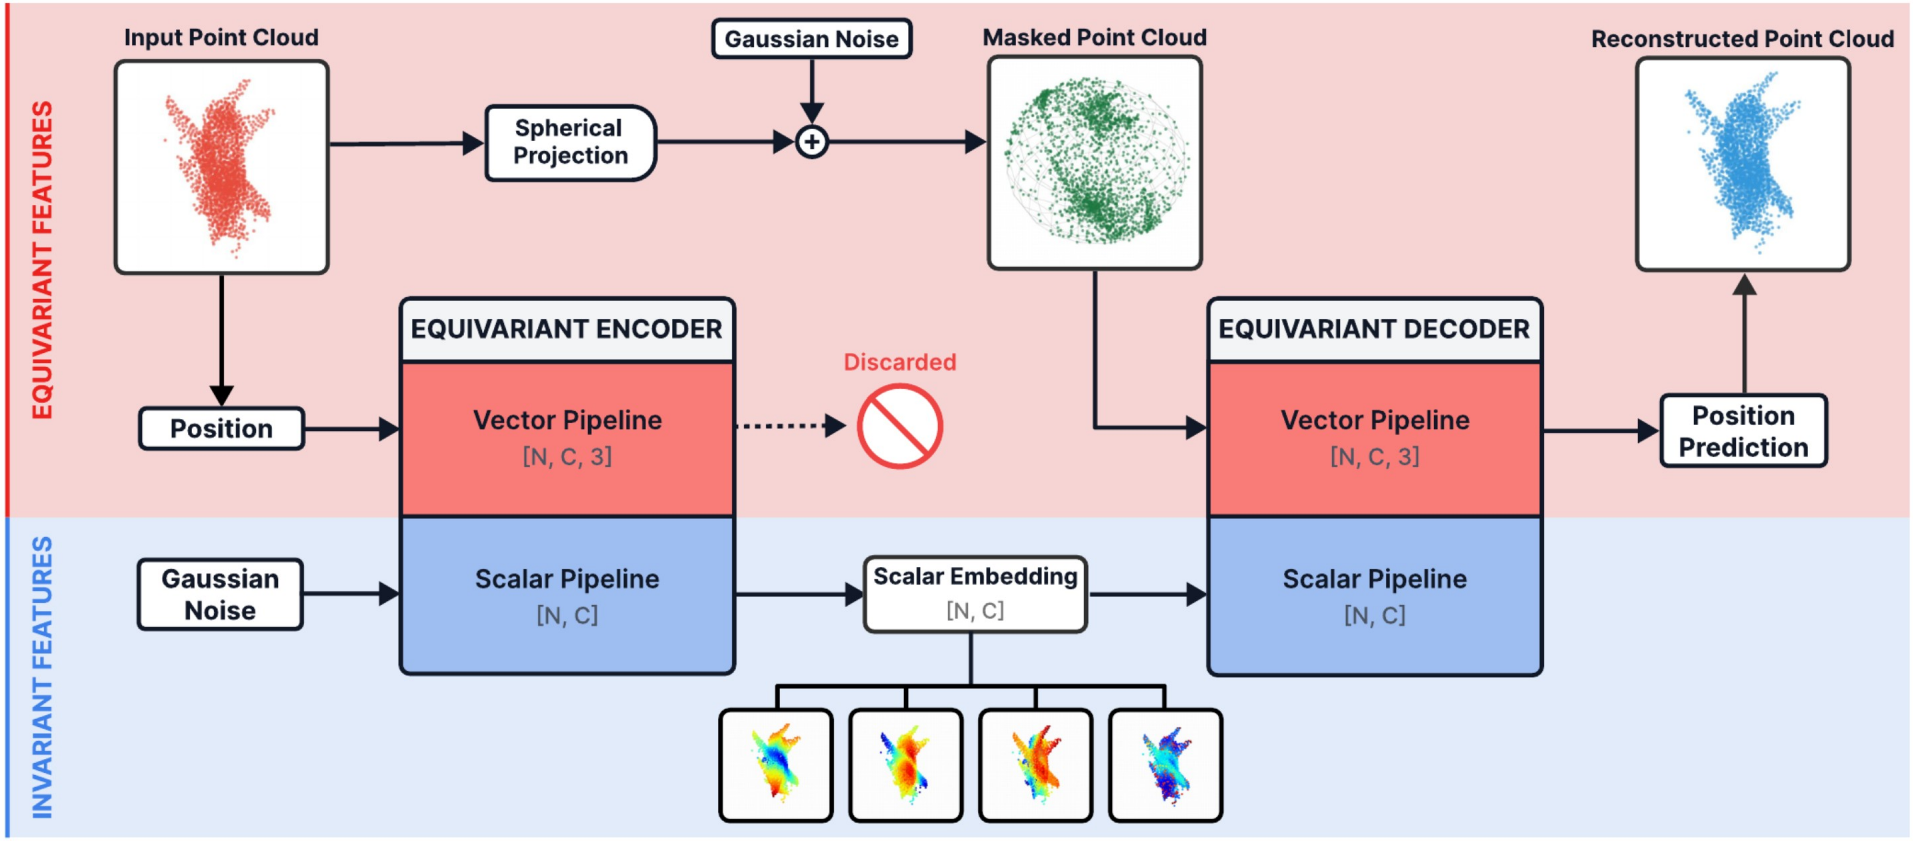
Theorem 3.1: the spherical projection mask

With the spherical projection mask, SIEVE can approximate pullbacks of non-trivial bundles over S^2.

Consequence: SIEVE can reconstruct point clouds of different topologies.

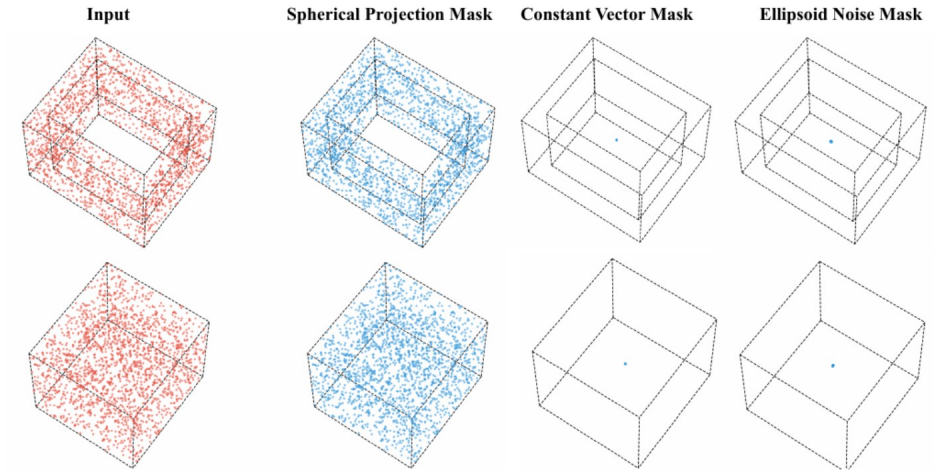


SIEVE: only invariant scalars get through

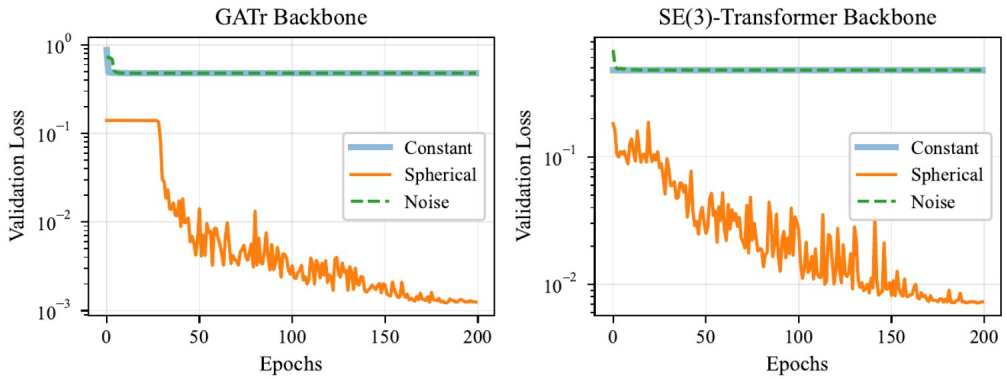


Vector output discarded; only invariant scalars reach the decoder. Pre-trained on 876k OpenShape models.

The predicted collapse is real



Reconstructions: square torus (top) and cuboid (bottom)



Validation L1 over 200 epochs, on two equivariant backbones

Only the spherical mask learns (L1 below 0.01). Constant and noise masks stall near 0.5 and output a single point.

Drop-in features that generalize

Frozen SIEVE encoder. Its features are concatenated with xyz at the input of standard models.

Task	Metric	Baseline	+ SIEVE
RNA segmentation (PointNet++)	Accuracy ↑	71.6%	86.3%
Shape completion, unseen categories (PCN)	Chamfer ↓	11.98	10.49
Registration, unseen categories (DCP)	Rotation RMSE ↓	2.64	0.94

Completion: 3D-EPN, Chamfer distance $\times 10^4$ averaged over the 8 unseen categories. Registration: ModelNet40, 20 held-out categories.

Gains are largest on unseen categories: completion error −12.5% (vs −1.6% on seen), rotation error 2.8× lower.

Takeaways

1. Constant and random masks provably collapse equivariant MAEs.
2. A spherical projection mask avoids the obstruction by construction.
3. SIEVE is an equivariant self-supervised learning method that uses a spherical projection mask and an equivariant bottleneck.
4. Frozen SIEVE features can easily be applied to any point cloud task and especially improve performance on unseen shapes.

Paper: SIEVE, NeurIPS 2026 · Lee, Park, Choi