

Andrea Emir Sevincel

Milan, IT | +39 3348967889 | aemirsevincel@gmail.com | andrasevincel.github.io

Final-year AI student working on 3D Deep Learning, Geometry Processing and Geometric Deep Learning

Education

B.Sc. in Artificial Intelligence (Joint Degree Program)

09/2023 - 03/2027 (expected)

University of Pavia, University of Milan, University of Milan-Bicocca

Relevant Coursework: Machine Learning, Deep Learning, Computer Vision, Data Structures and Algorithms

Preprints

A. E. Sevincel. *What Symmetry Buys a Learned Motion Planner*. [arXiv:2609.10033](https://arxiv.org/abs/2609.10033), 2026. Submitted to ICRA 2027.

Experience

Visiting Student Intern

10/2026 - 01/2027 (expected)

Hong Kong University of Science and Technology (HKUST)

- Incoming visiting research intern in Prof. Yuan Liu's group, working on 3D generative modelling.
- B.Sc. thesis research in Intelligent Graphics Lab, funded by both an Erasmus+ Traineeship grant and HKUST.

Research Collaborator

05/2026 - Present

MeshAnyOrder - Autoregressive Mesh Generation (7-author collaboration; equal-contribution first author)

- Designed and implemented the rotation-equivariant and canonicalization components of an autoregressive mesh transformer, extending its translation equivariance to full rigid motions.
- Identified OBJ text parsing as the dominant cost against a 150 ms training-step budget on four A100 GPUs, and specified an HDF5 format with per-field data types that cuts index-array memory 4x.

Research Intern

07/2026 - 08/2026

Shanghai Jiao Tong University (SJTU)

- Showed that the frame defined by a query's start and goal removes five of the six rigid degrees of freedom at no model cost, raising the held-out collision-free rate from 14.6%, no better than a straight line, to 51.1%.
- Proved the remaining rotation admits no continuous canonical frame, then showed that augmentation, test-time symmetrization and an exactly equivariant backbone each add under one point, with the backbone reaching that accuracy in a quarter of the epochs.
- Found that splitting by trajectory leaks near-duplicate paths into validation (nearest-neighbour error 0.030 against 0.162), and built PointMass3D, the benchmark behind the study, with splits that hold out whole start-goal pairs.

Research Contributor, Embedding Stability for Single-Cell Data

06/2026 - 07/2026

MIT CSAIL - Mantis AI - Kellis Lab

- Showed that repeated UMAP runs on identical data differ partly by an arbitrary rotation or reflection, so layout comparisons in Mantis conflated this gauge freedom with genuine structural change.
- Removed that freedom by generalized orthogonal Procrustes alignment and derived per-point instability and trustworthiness scores that separate reproducible structure from layout artifacts on PBMC3k.

Projects

Generative Search for Non-Rupert Convex Polyhedra

02/2026 - Present

- Self-directed project: found that the passage search labeling 3,000 generated polyhedra returns only the trivial value of 1, even for the octahedron (known value 1.061), so the conditional VAE trained on those labels could not target non-Rupert shapes.

Generative Modeling of Pareto-Optimal 3D Shapes (Project Lead)

03/2026 - 04/2026

- Led a 3-person team in designing a pipeline that optimizes neural signed distance functions under competing geometric objectives and trains a conditional VAE on the results to generate shapes from target properties, writing the full technical specification.

Technical Skills

- **Programming Languages:** Python, C, C++
- **Machine Learning & Tools:** PyTorch, PyTorch Geometric, JAX, NumPy, SciPy, Open3D, OpenCV, Weights & Biases, Git, Docker, Linux