

Beatrice Bevilacqua

Machine Learning Research Scientist

beatricebevilacqua4@gmail.com

📈 1.4k cit - h-index 15 |   beabevi

Research scientist developing and evaluating large-scale generative models for biological data, spanning pretraining, post-training, and rigorous evaluation across biological contexts. My broader research focuses on generalization, equivariance, and causal learning, with publications at *Cell*, ICLR, NeurIPS, and ICML.

Experience

Arc Institute | *Machine Learning Research Scientist*

May 2025 - Present | Palo Alto, CA

- Led the architecture and end-to-end implementation of an internal decoder-only autoregressive model for cellular biology, from pretraining on observational single-cell data through post-training on perturbational data.
- Delivered the strongest internal results to date for perturbation-specific gene recovery, improving over baselines by 43% on cancer cell-line data.
- Implemented multi-node training with PyTorch FSDP, supporting both pretraining and post-training.
- For the earlier STATE project, formalized the model description and co-developed the theoretical analysis proving that its solution family can recover continuous optimal-transport maps; published in *Cell*.

Google DeepMind | *Research Scientist Intern*

Sep 2022 - Jan 2023 | London, UK

- Developed a synthetic-data generation method based on causal interventions to learn invariant representations for neural algorithmic reasoning, enabling execution on inputs $4\times$ larger than those seen during training and improving baseline performance by $3\times$.
- Work with Petar Veličković resulted in an ICML publication and a U.S. patent application.

Meta FAIR | *Research Scientist Intern*

May 2022 - Aug 2022 | Menlo Park, CA

- Designed a self-supervised objective, formalized through a causal graph, to learn image representations invariant to day-to-night domain shifts.

Sapienza University | *Research Collaborator*

Jan 2020 - Jul 2020 | Rome, Italy

- Developed a post-hoc subgraph attribution method for interpreting knowledge-graph predictions while preserving the original model's performance.

Selected Publications (full list on Google Scholar)

1. A. K. Adduri, D. Gautam, **B. Bevilacqua**, et al. "Predicting cellular responses to perturbation across diverse contexts with State." **Cell**, 2026.
2. M. Eliasof, K. S. I. Mantri, **B. Bevilacqua**, B. Ribeiro, and C.-B. Schönlieb. "Bridging Input Feature Spaces Towards Graph Foundation Models." **ICLR**, 2026.
3. **B. Bevilacqua**, J. Robinson, J. Leskovec, and B. Ribeiro. "Holographic Node Representations: Pre-training Task-Agnostic Node Embeddings." **ICLR**, 2025.
4. Y. Shen, J. Zhou, **B. Bevilacqua**, J. Robinson, C. Kanatsoulis, J. Leskovec, and B. Ribeiro. "Zero-Shot Generalization of GNNs over Distinct Attribute Domains." **ICML**, 2025.
5. M. Eliasof*, **B. Bevilacqua***, C.-B. Schönlieb, and H. Maron. "GRANOLA: Adaptive Normalization for Graph Neural Networks." **NeurIPS**, 2024. *Equal contribution.

6. Y. Zhang, **B. Bevilacqua**, M. Galkin, and B. Ribeiro. "TRIX: A More Expressive Model for Zero-shot Domain Transfer in Knowledge Graphs." **LoG**, 2024.
7. **B. Bevilacqua**, M. Eliasof, E. Meirom, B. Ribeiro, and H. Maron. "Efficient Subgraph GNNs by Learning Effective Selection Policies." **ICLR**, 2024.
8. **B. Bevilacqua**, K. Nikiforou, B. Ibarz, I. Bica, M. Paganini, C. Blundell, J. Mitrovic, and P. Veličković. "Neural Algorithmic Reasoning with Causal Regularisation." **ICML**, 2023.
9. F. Frasca*, **B. Bevilacqua***, M. Bronstein, and H. Maron. "Understanding and Extending Subgraph GNNs by Rethinking Their Symmetries." **NeurIPS Oral**, 2022. *Equal contribution.
10. **B. Bevilacqua***, F. Frasca*, D. Lim*, B. Srinivasan, C. Cai, G. Balamurugan, M. M. Bronstein, and H. Maron. "Equivariant Subgraph Aggregation Networks." **ICLR Spotlight**, 2022. *Equal contribution.
11. **B. Bevilacqua**, Y. Zhou, and B. Ribeiro. "Size-Invariant Graph Representations for Graph Classification Extrapolations." **ICML Long Talk**, 2021.

Education

Purdue University | *Ph.D. in Computer Science, GPA: 4.0/4.0* 2020 - 2025 | West Lafayette, IN

- Research: foundation models, generalization, graph representation learning, and causal machine learning.
- Advisors: Prof. Bruno Ribeiro (Purdue) and Prof. Haggai Maron (Technion, NVIDIA).

Sapienza University | *M.Sc. in Computer Science and Engineering, with honors* 2017 - 2019 | Rome, Italy

Sapienza University | *B.Sc. in Computer Science and Engineering, with honors* 2014 - 2017 | Rome, Italy

Awards and Recognition

Rising Stars in EECS, MIT 2024

Top Reviewer, NeurIPS 2022, 2023, 2024, 2025

Employee Recognition Award, Purdue University 2024

Andrews Ph.D. Fellowship, Purdue University 2020 - 2022

Additional Recognition

Reviewer for ICML, NeurIPS, ICLR, TPAMI, and TMLR. Co-inventor on U.S. Patent Application US20240256879A1 on self-supervised neural algorithmic reasoning.