Christian Kittim Belardi

ckb73@cornell.edu | Personal Website | Google Scholar

PhD candidate in Computer Science at Cornell University researching diffusion and language models, with industry experience leading DARPA and NASA research projects.

Education

Doctor of Philosophy, Computer Science — Expected 2027
Cornell University, College of Computing and Information Science — Ithaca, NY
- Co-advised by Professor Carla Gomes and Professor Kilian Weinberger.

Master of Science, Computer Science — December 2024
Cornell University, College of Computing and Information Science — Ithaca, NY

Bachelor of Science, Computer Science, *cum laude* — May 2018
Cornell University, College of Engineering — Ithaca, NY

Selected Publications

Belardi, C., Lovelace, J., Weinberger, K. Q., Gomes, C. P. (Under Review). Inverse Problem Solving by Walking the Sphere. The International Conference on Learning Representations (ICLR).

Belardi, C., Lovelace, J., Weinberger, K. Q., Gomes, C. P. (2026). Adaptive Moments are Surprisingly Effective for Plug-and-Play Diffusion Sampling. The International Conference on Learning Representations (ICLR).

Lovelace, J., **Belardi, C.,** Kundurthy, S. R., Sudhakar, S., Weinberger, K. Q. (2026). Language Denoising Objectives Extend the Value of Limited Data. The Conference on Neural Information Processing Systems (NeurIPS).

Lovelace, J., **Belardi, C.,** Kundurthy, S. R., Sudhakar, S., Weinberger, K. Q. (2026). Prescriptive Scaling Laws for Data Constrained Training. The Conference on Language Modeling (COLM).

Lovelace, J., **Belardi, C.**, Zalouk, S., Polavaram, A., Kundurthy, S. R., Weinberger, K. Q. (2025). Stop-Think-AutoRegress: Language Modeling with Latent Diffusion Planning. The Conference on Language Modeling (COLM).

Belardi, C. K.*, Lee, C. H.*, Wang, Y., Lovelace, J., Weinberger, K. Q., Muller, D. A., Gomes, C. P. (2025). Improving Multislice Electron Ptychography with a Generative Prior. The Computer Vision for Materials Science (CV4MS) Workshop at the International Conference on Computer Vision (ICCV).

Zhou, J. P.*, **Belardi, C. K.***, Wu, R.*, Zhang, T., Gomes, C. P., Sun, W., Weinberger, K. Q. (2025). On Speeding Up Language Model Evaluation. The International Conference on Learning Representations (ICLR).

* equal contribution.

Academic Experience

Graduate Research Assistant — Jan 2023 - Present
Cornell University, College of Computing and Information Science — Ithaca, NY

- Inverse problem solving with generative models
  - Demonstrated that projected gradient descent in the latent space of a sphere encoder is often superior to diffusion-based solvers for inverse problems, achieving the best reconstructions on a variety of benchmarks.
  - Stabilized noisy likelihood scores in plug-and-play diffusion sampling via adaptive moments, improving sample quality and fidelity to the condition across restoration and class-conditional tasks, including a 19+ point gain in ImageNet class-conditional accuracy.
  - Built a diffusion model for multislice electron ptychography that improves reconstruction SSIM by 90% over the state-of-the-art solver and recovers depth structure where existing solvers collapse.
- Language model training and evaluation
  - Developed scaling laws for data-constrained pretraining that capture the damage caused by repeating data, adding an overfitting penalty to the Chinchilla law. The law prescribes better model sizes and epoch counts than Chinchilla and prior data-constrained laws for a given compute budget and amount of unique data.
  - Demonstrated that augmenting the causal language modeling loss with span-corruption denoising objectives cuts overfitting from data repetition by approximately half, improving downstream math performance when training on limited data.
  - Built STAR-LDM, a language model that plans in a continuous latent space via diffusion before generating text. It outperforms similar-sized models on language understanding benchmarks and, in LLM-judged story completion, beats GPT-2 XL and matches models 7–12× its size on most criteria.
  - Designed bandit algorithms that exploit the low-rank structure of evaluation matrices to efficiently identify the best model, prompt, or method, reducing evaluation cost by 85–95%.

Graduate Teaching Assistant Aug 2021 - Dec 2022

Cornell University, College of Computing and Information Science Ithaca, NY

- CS 6784 Advanced Topics in Machine Learning - Deep Learning (Fall 2022)
- CS 4780 Introduction to Machine Learning (Spring 2022)
- CS 4410 Operating Systems (Fall 2021)

Siegel Public Interest Technology Impact Fellow Jun 2022 - Aug 2022

New York Botanical Gardens New York, NY

- Constructed a remote sensing dataset composed of satellite imagery from the European Space Agency's (ESA) Copernicus program and the corresponding tree canopy classifications derived from Global Biodiversity Information Facility (GBIF) and iDigBio.
- Developed deep learning models for biodiversity estimation from remote sensing data.

Industry Experience

Research Scientist Intern May 2025 - Sep 2025

Meta New York, NY

- Collaborated with researchers from FAIR and Reality Labs on fine-tuning language models with memory layers for specialized domains.
- Developed an approach for adding specialized memory, matching the performance of established parameter-efficient fine-tuning methods.

Artificial Intelligence Developer Aug 2018 - Jul 2021

Stottler Henke Associates, Inc. San Mateo, CA

- Wrote and won several grant proposals: **Pathogen Classification Tool** (DARPA Phase II STTR), **Pathogen Classification Tool** (DARPA Phase I STTR), **Question Answer Generation** (ARI BAA) [Co-authored] , **Delay/Disruption Tolerant Network Scheduling** (NASA Phase I SBIR) [Co-authored], **On-The-Fly Parameter Selection** (NGA Phase I SBIR).
- Led and contributed to several NASA and DoD applied research efforts, described below.
  - **Pathogen Classification Tool.** Developed deep learning models to characterize bacteria in image data from the DARPA Friend or Foe program. Led the project through Phase I and into Phase II, managing two full-time employees and an intern. Regularly prepared and presented progress reports to DARPA. Periodically briefed the DARPA Friend or Foe program.
  - **Space Domain Awareness through Heuristics and Fuzzy Logic.** Implemented symbolic reasoning and machine learning algorithms to assess the likely explanations of anomalous behavior in space for Air Force Research Laboratory (AFRL). Participated multiple times in the Air Force's Sprint Advanced Concepts Training (SACT) exercises demonstrating technology on the unclassified operations floor in a live military training exercise.
  - **Anomaly Detection via Topological Feature Maps.** Developed anomaly detection software for long duration space flight using a combination of self-organizing maps and case-based reasoning. Demonstrated work on multiple occasions to NASA program managers and presented this research at the 2020 IEEE Aerospace Conference.
  - **Naval Inverse Synthetic Aperture Radar Classification.** Implemented deep learning models for hierarchical classification of maritime targets in Inverse Synthetic Aperture Radar imagery.

Software Engineering Intern May 2017 - Aug 2017

iRobot Bedford, MA

- Fixed state-machine bugs in C/C++ and supported Alpha 2 testing for the Terra robotic lawn mower.

Awards

Teaching Assistant Award May 2022

Cornell University, College of Computing and Information Science Ithaca, NY

Volunteer Experience

Workshop Leader, Expanding Your Horizons at Cornell (2022-2024) • Mentor, Cornell HS Programming Contest (2022-2023) • Mentor, Bowers CIS Undergraduate Research Experience (Summer 2024)

Skills

Programming Languages: Python, C/C++, Java
Deep Learning: PyTorch, Hugging Face (Transformers, Diffusers, Accelerate, Datasets, etc.), Einops, Kornia, xFormers, FlashAttention
LLM Training & Inference: DDP, FSDP, torch.compile, multi-node GPU training, vLLM, bitsandbytes
Evaluation: lm-evaluation-harness, OLMES, LLM-as-judge pipelines
Scientific Computing: NumPy, SciPy, Pandas, scikit-learn, Matplotlib, Seaborn, Plotly
Infrastructure & Tools: Slurm, Docker, Singularity, AWS, Azure, Weights & Biases, TensorBoard, Optuna, Git, LaTeX