

Jacopo Minniti

jacopominniti.com | jacopo.minniti004@gmail.com | [Github](#) | [Linkedin](#)

EDUCATION

Minerva University

Sep 2023 – May 2027

Bachelor of Science in Computational Sciences

- Concentrations: AI & Mathematics; Minor: Neuroscience. **GPA 4.00 (major); 3.94 (overall).**
- Relevant Courses: ML, Bayesian Statistics, Abstract Algebra, Real Analysis, ODEs, Optimization Theory.

PUBLICATIONS

Learned Relay Representations for Forward-Thinking Discrete Diffusion Models

Submitted to NeurIPS 2026

- Introduced learned relay representations trained via Backpropagation Through Time to pass continuous latent states across denoising steps, enabling forward-thinking in discrete diffusion models.
- Demonstrated significant reasoning and coding gains, improving LLaDA-8B accuracy on MATH500 by 7.6% and raising HumanEval performance from 23.78% to 32.32%.

Challenges in Inference Time Scaling with Uncertainty Aware Tree Search

ICLR 2026, Workshop on Scaling Post-training for LLMs

- Designed inference algorithm to reallocate CoT budget to ambiguous states using a process uncertainty model.
- Showed that entropy is a reliable uncertainty metric in-distribution but fails under search-induced distribution shift due to overoptimization, with correlation to rollout variance collapsing from $r = 0.62$ to 0.03 .

Effects of Semantic Relatedness and Vocabulary Knowledge on New Word Meanings

Submitted to Journal of Memory and Language

- Conducted experiments with 300+ participants to evaluate the acquisition of 32 novel meanings for familiar words.
- Demonstrated that semantic relatedness increased recall accuracy ($p < .001$) and that higher baseline vocabulary scores predicted improved form-meaning association ($R = 0.2$), mitigating age-related learning declines.

PROFESSIONAL EXPERIENCE

Member of Technical Staff, [Sakana AI](#), Tokyo

May 2026 – Sep 2026

- Developed a theory and experimental framework for decomposing transformer learning into discrete, reusable computations and ordered causal dependencies.
- Used this framework to guide training of a transformer, improving exact-sequence accuracy by 75.3% on NumberNaming dataset.

Research Assistant, [University of Toronto](#), Remote

Sep 2025 – Present

- Reduced compute waste in diffusion language models by introducing a continuous, stateful information flow across denoising steps.
- Applied Backpropagation Through Time to enhance long-horizon planning in token demasking.

AI Safety Researcher, [University of Chicago](#), [XLab](#), Chicago

Jun 2025 – Aug 2025

- Analyzed how diverse uncertainty signals are represented within LLM latent states during reasoning, assessing the extent to which models internally track their own uncertainty.
- Reached 89% AUROC by training a Qwen3 LLM to predict reasoning entropy at the CoT step level.

SKILLS & ACCOMPLISHMENTS

- *Programming Languages:* Python, Java, JavaScript, Dart, Matlab, SQL
- *Technical Skills:* AWS, Linux, scikit-learn, PyTorch, React, Docker, Git, RAG, EEGLAB
- *Certifications:* SAT (1580/1600), Stanford ML, Dart & Flutter (Udemy), Next.js (Udemy)
- [Talk at the University of Chicago](#) on prospects of using machine learning to categorize inductive biases.
- [Winner of NASA Space Apps 2023](#): customized A^* algorithm for safe evacuation routing in wildfire scenarios.
- Secured [Double Prizes](#) at CalHacks 2023, Berkeley University: built AutoML querying system for SQL database.