

Thiago Tarraf Varella

+1 (561) 870-3264 | thiago.t.varella@gmail.com | LinkedIn | Google Scholar | GitHub | San Francisco, CA

SUMMARY

ML Research Scientist with a PhD in multi-agent decision-making with reinforcement learning (Nature Communications Biology, PLoS Computational Biology). At Meta, I ship production models for GPU scheduling and infrastructure optimization. Independently building an autonomous ML research system that handles the full loop without intervention: literature discovery, experiment design, cloud dispatch, and structured learning.

SKILLS

ML/AI: PyTorch, reinforcement learning (Actor-Critic), transformers, post-training, multi-agent systems, Bayesian inference, GBDTs
Agentic: LLM orchestration, tool use, eval loops, multi-agent coordination, safety protocols, Anthropic API, Claude Code
Engineering: Python (10+ yrs), C++, SQL, Git, CI/CD, trunk-based DevOps, TB-scale data pipelines, multi-GPU training

EXPERIENCE

Meta · ML Research Scientist · Menlo Park · AI Infrastructure

Jun 2025 – Present

ML Systems for GPU Scheduling

- Built the first ML system on a team running deterministic algorithms. Owned full roadmap: data pipelines, feature engineering, transformer encoder, training, inference, production integration, observability. Targeted \$500K+/day OpEx waste.
- Built ML models for scheduling optimization: job runtime prediction (>95% accuracy, targeting 30% fewer interruptions) and GPU host availability prediction (improved placement accuracy from 50% to 60%, reducing interruptions by ~22%).
- Extended job priority system for RecSys to enable gang-aware preemption, helping define new SLOs and contributing to Muse Spark training.

Agentic AI Infrastructure

- Built agentic evals infrastructure: daily automated evals for knowledge recall, adversarial robustness, task completion. Integrated pre/post A/B testing into the CI/CD pipeline.
- Designed AI triage agent ingesting alerts and emails, producing structured triage cards with category, severity, and routing recommendations.

Platform & Reliability

- Identified reliability bug in multi-system kill-signal chain causing orphaned jobs; traced across multiple systems, escalated to XFN partners, deployed fix eliminating the failure mode (~\$2M annual savings).
- Reduced failure rate by ~25% in certain regions via eligibility filters excluding clusters without secure fiber for multi-host jobs.
- Built SQL telemetry pipeline over TBs of GPU data; contributed to ~\$90M monetization improvement.

Princeton University · Graduate Research Scientist · Princeton Neuroscience Institute

Jan 2017 – May 2025

RL, Multi-Agent Systems & Behavioral Modeling

- Architected custom RL environment for vocal motor learning. High-exploration agents significantly outperformed imitation learning; poorly timed external feedback corrupted the intrinsic reward signal, suggesting incentive-based interventions can destabilize agentic learning loops. Part of PhD dissertation.
- Engineered multi-agent Bayesian decision loop where agents update belief states to maximize Expected Information Density via Bayesian filter updates, modeling information-seeking as the governing agentic strategy. Published in *Nature Communications Biology*.
- Created thermodynamic landscape model for behavioral transitions (R^2 0.86–0.96 vs 0.54–0.80 baselines across humans, bats, monkeys). Published in *PLoS Computational Biology*.
- TA for Computational Models of Cognition and Computational Neuroscience (Neuromatch). 4 invited international talks.

Signal Processing & Infrastructure

- Built pipelines for 100K+ longitudinal recordings (audio, EMG, EEG, ECoG) across five species. Led project involving cross-functional team of 4 (philosophy, medicine, physics, neuroscience).
- Shipped stable-matching scheduler to department staff, replacing hours of manual work with seconds of compute. Used for 4+ years; won Princeton Graduate Student Service Award.

MIT CBMM · Summer Researcher · Advisor: Andrei Barbu

Summer 2019

- Independently analyzed high-frequency human intracranial ECoG data from children/teens watching a movie. Successfully predicted auditory vowel processing using PyTorch.

INDEPENDENT RESEARCH

Autonomous ML Research System Autonomous ML

Claude Code framework applying iterative refinement across the research cycle: literature discovery with convergence detection, prototype scaffolding, and experiment automation with metric-based hypothesis selection.

- **Grid Graph Optimization:** Reduced a combinatorial puzzle to connected dominating sets; derived self-contained lower-bound proof using capacity identities from recent graph theory literature. Identified an open conjecture on corner-constrained cases.
- **Small-Scale Introspective Diffusion LMs:** Converting autoregressive models (GPT-2, Qwen) to diffusion generators via introspective-consistency training. First investigation below 8B parameters; enables self-speculative decoding for faster single-request inference.
- **Photo Restoration Benchmark:** Evaluating how well frontier models (GPT Image 2, InstructPix2Pix) restore damaged photos. Built synthetic degradation pipeline with perceptual similarity scoring and VLM-based quality validation.

Python, PyTorch, Claude Code, Thunder Compute (A100)

Agentic AI Portfolio Agent Infrastructure

(1) AI coaching agents: coordinator routing to sub-agents via tool invocation, persistent memory, safety-protocol layer. (2) FazAI: AI business advisor with orchestrator dispatching worker agents via tool use. (3) [Caramelo](#): adaptive exam-prep with LLM explanation agent, FSRS, 180+ items.

Anthropic API, FastAPI, Flutter, Claude Code

SELECTED PUBLICATIONS

- T. T. Varella et al. "Active sampling in primate vocal interactions." **Nature Comms Bio**, 2024. [\[link\]](#)
- T. T. Varella et al. "Punctuating equilibria during mammalian vocal development." **PLoS Comp Bio**, 2022. [\[link\]](#)
- T. T. Varella, A. A. Ghazanfar. "Cooperative care and the evolution of prelinguistic vocal learning." **Dev Psychobio**, 2021. [\[link\]](#)
- T. T. Varella et al. "Explore-exploit trade-off in vocal learning." **IEEE ICDL**, 2024. [\[poster\]](#)
- T. T. Varella et al. "Large language models in 12 points." **Nexo Politicas Publicas**, 2024. 3x avg readership. [\[article\]](#)
- T. T. Varella et al. "Phase transitions in vocal development are driven by energy-information balance." **COSYNE**, 2020. [\[poster\]](#)

EDUCATION

- Ph.D., Psychology & Computational Neuroscience** · [Princeton University](#) 2025
GPA 3.9/4.0 · Advisor: Asif A. Ghazanfar · NSF Graduate Research Fellowship
- B.Sc., Molecular Sciences (Mathematics Applied to Neuroscience)** · [University of São Paulo](#) 2018
GPA 8.9/10 · Merit-based exchange at Princeton (2017)

AWARDS

- NSF GRFP (~\$100K, 2020–2025) **FAPESP Fellowship** (largest undergrad research award in Brazil, 2017) **Princeton President's Fellowship** (2019)
- Meta Hackathon "Greatest Pitch" (top 7%, 2x: 2025, 2026)

PROFESSIONAL DEVELOPMENT

- [DeepLearning.AI](#) — **Agentic AI (Andrew Ng)**, self-paced: reflection, tool use, planning, multi-agent orchestration 2026
- [Andrej Karpathy](#) — **Neural Networks: Zero to Hero**, self-paced: backprop to GPT, transformers, tokenization 2025
- [Hugging Face](#) — **Deep Reinforcement Learning**, self-paced: policy gradients, PPO, actor-critic, RLHF 2024
- [Princeton PICSciE](#) — **Multi-GPU Training**, 1-day workshop: data and model parallelism in PyTorch 2023
- [BAMB](#) — **Barcelona Summer School**, 2 weeks: advanced modeling of behavior 2023
- [MIT CBMM](#) — **Brains, Minds & Machines (Woods Hole)**, 3 weeks: neuroscience, cognitive science, computational methods 2019

Personal interests: Piano (classical/jazz), singing (solo/a cappella), partner dancing, education, camping