

Thiago Tarraf Varella

thiago.t.varella@gmail.com · +1 (561) 870-3264 · [GitHub](#) · [LinkedIn](#) · thiagotvarella.github.io

SUMMARY

Machine Learning Scientist combining 10 years of computational behavioral modeling with production ML engineering at Meta. My PhD research focused on mathematically defining how biological agents change behavior, using RL, Bayesian inference, and thermodynamic models to balance energetic cost against social gains. At Meta, I translate complex systems into production-grade ML pipelines, optimizing GPU infrastructure at scale. Seeking to blend behavioral science and scalable ML into agentic systems.

WORK EXPERIENCE

Meta, ML Research Scientist

Jun 2025 – Present

- Built the first ML system deployed on a GPU job-scheduling team that had run exclusively on deterministic algorithms, independently defining the ML roadmap and designing the full system architecture: data pipelines, feature engineering, model training, scheduler API integration, and observability
- Used a transformer-based encoder to generate embeddings from heterogeneous job metadata, feeding them into a downstream NN to predict runtime, data transfer, and resource availability, to collectively target an identified \$500K+/day in OpEx waste
- Proactively traced an orphaned-job bug through a multi-system kill-signal chain, escalated to system owners to identify the root cause, and fixed it via UI and backend changes, preventing an expected \$4M+/year in waste.
- Improved flexible GPU-host selection accuracy from 50% to 60% for a preemption system that dynamically prioritizes higher-value training jobs, reducing daily job interruptions from higher-priority jobs by 22% and improving training efficiency
- Built a SQL telemetry pipeline over TBs of GPU data-transfer patterns, enabling ML features and anomaly detection that contributed to a \$90M GPU starvation reduction. All work shipped via trunk-based DevOps with CI/CD, code reviews, and tests

Princeton University, Graduate Research Scientist & Research Assistant

Jan 2017 – May 2025

- **Architected a custom RL environment for primate vocal motor learning:** the actor maps biomechanical states (air pressure, vocal fold tension) to continuous actions, the critic approximates a value function, and reward is negative MSE against target acoustic output. Main findings:
 - Exploration: high-exploration agents significantly outperformed imitation learning; premature end to the plasticity window prevents learning entirely
 - Social contingency: poorly timed external feedback corrupted the intrinsic reward signal, suggesting that incentive-based interventions can backfire in behavior change and destabilize agentic learning loops
- **Engineered a multi-agent Bayesian decision loop** where agents updated belief states to maximize Expected Information Density via Bayesian filter updates, modeling information-seeking as the governing agentic strategy over social accommodation
- **Developed a thermodynamic landscape model** framing behavioral transitions not as rational state machines, but as a balance between energetic cost and information gain across hidden motivational states. Achieved R^2 of 0.86–0.96 (vs. 0.54–0.80 linear/recurrent baselines) in humans, bats, and monkeys, predicting developmental timing under novel experimental perturbations
- Designed scalable signal processing pipelines for 100k+ longitudinal recordings in audio, EMG, MCG, EEG, and ECoG, modalities whose acquisition and processing methods transfer directly to human health monitoring
- Sole architect of full ML stack across all projects, from synthetic data generation through model design, training, evaluation, and publication. Published 3 peer-reviewed papers, gave invited international talks, led a team of 4, and mentored undergraduates

EDUCATION

Princeton University, *Ph.D., Psychology & Computational Neuroscience*, GPA 3.9/4.0

May 2025

University of São Paulo, *B.Sc., Mathematics Applied to Neuroscience*, GPA 8.9/10

May 2018

SKILLS

ML/AI: PyTorch · Actor-Critic RL · Bayesian Inference · Generative Models · Gaussian Mixture Models · Information Theory

Core: Python · MATLAB · R · C++ · SQL · CI/CD · Multi-GPU · Signal Processing (EMG, EEG, ECoG, Audio, Behavioral)

ADDITIONAL TRAINING & INTERESTS

AI: Agentic AI (Andrew Ng) · Neural Networks: Zero to Hero (Karpathy) · Deep Learning (Neuromatch) · Advanced Modeling of Behavior (BAMB) · Brain, Minds & Machines (CBMM, MIT) · Multi-GPU Training (Princeton PICSciE)

Interests: Piano (classical/jazz) and singing (solo/a cappella) · partner dancing · education · hiking

SELECTED PUBLICATIONS & CONFERENCES

- | | |
|--|------|
| T. T. Varella et al., IEEE International Conference on Development and Learning (ICDL) [poster] | 2024 |
| T. T. Varella et al., <i>Nature Communications Biology</i> , 7, 1098. doi:10.1038/s42003-024-06764-8 | 2024 |
| T. T. Varella et al., <i>PLoS Computational Biology</i> , 18(6). doi:10.1371/journal.pcbi.1010173 | 2022 |
| T. T. Varella et al., Computational and Systems Neuroscience (COSYNE) | 2020 |