

Tiago da Silva

tdsh97@gmail.com | github.com/tiagodasilva | tiagodasilva.github.io | linkedin.com/in/tiagodasilvah

Summary

I hold a PhD in Applied Mathematics and have a strong interest in Machine Learning (ML) and Bayesian statistics. My current research is centered around the question: how to incorporate Bayesian principles into Artificial Intelligence (AI) systems? The Bayesian paradigm provides a unifying framework for understanding, designing, and implementing ML models; however, their core tenets are seldom explicitly integrated into modern AI systems due to computational intractability. Developing computationally tractable approximations to Bayesian models is therefore the central objective of my ongoing and previous works, which have repeatedly appeared at the most prestigious ML conferences (ICML, NeurIPS, ICLR). Please refer to [my research and teaching statements](#) for further information.

Academic Employment

Postdoctoral associate, MBZUAI.	Abu Dhabi, UAE, 2025/10 – ongoing
Postdoctoral associate, Getulio Vargas Foundation.	Rio de Janeiro, Brazil, 2025/01 – 2025/10
Visiting scholar, Aalto University.	Espoo, Finland, 2024/07 – 2024/10

Education

PhD in Applied Mathematics	2024/03 – 2024/12
School of Applied Mathematics, Brazil. Advisor: Prof. Diego Mesquita .	
- Thesis title: Streaming, Distributed, and Asynchronous Amortized Inference. - Fast-tracked graduation due to recognized academic excellence and scientific productivity.	
MSc in Applied Mathematics	2024/03 – 2024/12
School of Applied Mathematics, Brazil. Advisor: Prof. Diego Mesquita .	
Dissertation title: Expert-Aided Discovery of Ancestral Graphs.	
BSc in Data Science	2020/03 – 2023/12
School of Applied Mathematics, Brazil	
GPA: 9.9/10.0. 1st in class and 1st in the entrance exam. Received a fully-funded scholarship.	

Selected Publications

1. Path-dependent Discrete Amortized Inference	ICML 2026 (oral) (CORE A*)
da Silva, T., Malkin, N., Lahlou, S. - TL;DR: We delineate the theoretical limitations of Markovian amortized samplers. To mitigate them, we augment existing methods with a deterministic, learnable latent dynamical system, improving learning convergence.	
2. When do GFlowNets learn the right distribution?	ICLR 2025 (spotlight) (CORE A*)
da Silva, T., Silva, E., Alves, R., Souza, A., Garg, V., Kaski, S., Mesquita, D. - TL;DR: We show that a GFlowNet’s learning objective may have an unattainable global minimum, which is undetectable by standard diagnostics. To address this, we propose the first tractable metric for assessing GFlowNets.	
3. Generalization and Distributed Learning of GFlowNets	ICLR 2025 (CORE A*)
da Silva, T., Souza, A., Rivasplata, O., Garg, V., Kaski, S., Mesquita D. - TL;DR: We devise the first PAC-Bayesian generalization bounds for GFlowNets. Inspired by them, we also introduce the first general-purpose distributed learning algorithm for GFlowNets, drastically accelerating training convergence.	
4. Streaming Bayes GFlowNets	NeurIPS 2024 (CORE A*)
da Silva, T., Souza, D., and Mesquita, D. - TL;DR: We design a method to update GFlowNets trained on a streaming Bayesian posterior. Experiments show a drastic reduction in training time when compared against learning from scratch a model based on the entire dataset.	
5. On Divergence Measures for Training GFlowNets	NeurIPS 2024 (CORE A*)
da Silva, T., Silva, E., and Mesquita, D. - TL;DR: We empirically show that the inefficacy of divergence-based objectives for GFlowNets is due to their large gradient variance. We then develop variance reduction techniques that significantly accelerate training convergence.	
6. Embarrassingly Parallel GFlowNets	ICML 2024 (CORE A*)

- [da Silva, T.](#), Souza, A., Carvalho, L., Kaski, S., and Mesquita, D.
- TL;DR: We propose a divide-and-conquer approach to train a log-pool of GFlowNets in an embarrassingly parallel fashion. Results show a significant speed up in learning when the unnormalized target is expensive to evaluate.

Teaching activity

1. **Co-supervision**, Getulio Vargas Foundation & MBZUAI 2024-ongoing
 - Assisted BSc, MSc, and PhD students in both formulating problems and writing papers (e.g., manuscripts 1-2 below).
 - Coached students in effectively responding to reviewer feedback.
2. **Teaching assistant**, Getulio Vargas Foundation 2021-2025
 - TA'd for machine learning, statistics, time series, data analysis, probability theory, and linear algebra courses.
 - Partnered with faculty to design and grade homeworks and exams.
 - Held office hours to support students, and provided structured solutions for take-home assignments ([here](#) and [here](#)).
3. **Industry Project Co-supervisor**, Rei do Pitaco 2025/08–2025/11
 - Co-supervised three undergraduate students in scoping and solving an open-ended industry problem.
 - Guided the team through problem formulation, solution design, implementation, and presentation.
4. **Guest Lecturer**, graduate-level courses on Machine Learning and Reinforcement Learning. 2024, 2025

Peer-reviewed research

1. **Particle GFlowNets: Rethinking Generative Marginalization Models** UAI 2026 (CORE A)
 - [da Silva, T.](#), Mesquita, D., Lahlou, S.
2. **Path-dependent Discrete Amortized Inference** ICML 2026 (oral) (CORE A*)
 - [da Silva, T.](#), Malkin, N., Lahlou, S.
3. **Avoid What You Know: Divergent Trajectory Balance for GFlowNets** ICML 2026 (CORE A*)
 - Dall'Antonia, P., [da Silva, T.](#), Csillag, D., Lahlou, S., Mesquita, D.
4. **Boosted GFlowNets: Improving Exploration via Sequential Learning** AISTATS 2026 (CORE A)
 - Dall'Antonia, P., [da Silva, T.](#), de Souza, D. A., Mattos, C. L. C., and Mesquita, D.
5. **When do GFlowNets learn the right distribution?** ICLR 2025 (spotlight) (CORE A*)
 - [da Silva, T.](#), Silva, E., Alves, R., Souza, A., Garg, V., Kaski, S., Mesquita, D.
6. **Generalization and Distributed Learning of GFlowNets** ICLR 2025 (CORE A*)
 - [da Silva, T.](#), Souza, A., Rivasplata, O., Garg, V., Kaski, S., Mesquita D.
7. **Streaming Bayes GFlowNets** NeurIPS 2024 (CORE A*)
 - [da Silva, T.](#), Souza, D., and Mesquita, D.
8. **On Divergence Measures for Training GFlowNets** NeurIPS 2024 (CORE A*)
 - [da Silva, T.](#), Silva, E., and Mesquita, D.
9. **Expert-aided Causal Discovery of Ancestral Graphs** LatinX @ NeurIPS 2024
 - [da Silva, T.](#), Silva, E., Góis, A., Heider, D., Kaski, S., Mesquita, D., Ribeiro, A.
10. **Embarrassingly Parallel GFlowNets** ICML 2024 (CORE A*)
 - [da Silva, T.](#), Souza, A., Carvalho, L., Kaski, S., and Mesquita, D.
11. **Exploring scientific literature by textual and image content using DRIFT** Computer & Graphics 2022
 - Pocco, X., [da Silva, T.](#), Poco, J., Nonato, L. G., Gomez-Nieto, E.

Community service

I served as a reviewer for ICML 2025, NeurIPS 2025, ICLR 2025, AISTATS 2026, and ICML 2026.

Languages

Portuguese (Native), English (Fluent)

Honors & Awards

Top-ranked doctoral thesis for Mathematics and Statistics, Brazilian Graduate Education Agency (CAPES) 2025

Top-ranked doctoral thesis for Computer Science, Brazilian Computer Society (SBC)	2025
Award for Academic Excellence , Brazilian Society of Applied and Computational Mathematics (SBMAC).	2023
First place , School of Applied Mathematics entrance exam.	2020
I was awarded <u>19 prizes in scientific competitions</u> during high school, including:	
William Glenn Whitley Prize for achieving the highest score on the State Mathematical Olympiad.	2019
Top score in the country , Brazilian Mathematical Olympiad of Public Schools.	2019
Top score in the country , Brazilian Mathematical Olympiad of Public Schools.	2018
Gold medal , Brazilian Chemistry Olympiad.	2018
Gold medal , Brazilian Mathematical Olympiad of Public Schools.	2017
Gold medals , State Chemistry Olympiad. Highest score in 2019.	2016-2019

Employment

Rei do Pitaco	<i>Santa Catarina, Brazil</i>
• Data Scientist.	2025/01 – 2025/11
▶ Spearheaded the development of the company’s bet builder, allowing customers to create customized bets. I designed, implemented, and maintained both the algorithm and the web API responsible for the real-time odds calculations. ▶ Handling up to 1.2 million requests per day, the service accounted for 50% of the company’s sportbook revenue.	
Proffer (Price Monitoring & Optimization startup)	<i>Rio de Janeiro, Brazil</i>
• Software Engineer.	2024/04 – 2025/09
▶ Extended the company’s API to support larger workloads via asynchronous processing and caching. ▶ Led the adoption of Infrastructure as Code (IaC) using Terraform as a standard for the company’s web scraping services. ▶ Assisted in the migration of the company’s infrastructure between cloud providers (AWS and Azure).	
Rei do Pitaco (largest fantasy sports company in Brazil)	<i>Rio de Janeiro, Brazil</i>
• Data Science intern.	2023/01 – 2023/07
▶ Designed predictive models to define the opening lines of bets on the outcomes of sport events (bookmaking). ▶ Deployed and upheld the created models within applications serving thousands of concurrent users.	

Skills

Computer languages: Proficient with <u>Python</u> and <u>SQL</u> . Competent with R and Stan and Rust. Familiar with C++.
Scientific computing frameworks: PyTorch, PyTorch Geometric, GPyTorch, NumPy, SciPy, JAX, Flax.
Technologies: Git, Linux, Docker, FastAPI, Datadog, Kubernetes, Argo, Tailscale, Kafka.

References

Diego Mesquita	diego.mesquita@fgv.br
Amauri Souza	amauri.souza@aalto.fi