

THOMAS PALMEIRA FERRAZ

🏠 Grenoble, France

🌐 [thomasferraz](#) 🌐 [thomas-ferraz](#) 🌐 [thomas-ferraz.github.io](#)

Education

Université Grenoble-Alpes, France **Sept. 2025 – Sept. 2028 (expected)**

PhD in Computer Science

Topic: Interpretability and Efficiency for Reasoning and Planning

Télécom Paris, ENS Paris-Saclay, France **Sep. 2022 – Aug. 2023**

PhD Track / Master M2 - Applied Mathematics, MVA - Mathematics, Vision, and Learning

GPA: 18.1/20

Télécom Paris, Institut Polytechnique de Paris, France **Sept. 2021 – Aug. 2022**

PhD Track / Master M1 - Computer Science, Data and Artificial Intelligence

GPA: 16.5/20

Universidade de São Paulo, Institute of Math and Statistics, Brazil **Sep. 2020 – Dec. 2021 (interrupted)**

MSc. Computer Science – Thesis topic: Adversarial Robustness on NLP

GPA: 4.0/4

Universidade de São Paulo, Escola Politécnica, Brazil (#1 in Latin America) **Feb. 2015 – Dec. 2020**

Engineering Degree – Major in Electrical Engineering, Electronics and Computer Systems

GPA: 8.6/10 - 1st/40

Research Interests

NLP, Low-Resource NLP, Large Language Models, Few-shot Learning, Adversarial Robustness, Machine Translation, Speech Processing, Efficient ML, Multilingual Models.

Work Experience

NAVER LABS Europe **Since April 2025**

Research Scientist / PhD Student CIFRE

Meylan, France

- Research on improving the interpretability and efficiency of large language models, with a particular emphasis on planning, reasoning, and AI agents.

Meta **Jun. 2024 – Oct. 2024**

Research Scientist Intern

Bellevue, WA, United States

- Working on Sparse Architectures and Mixture-of-Experts for Efficient ML, with applications on Recommendation Systems, on the ML Ads Infra Team. Paper in preparation for **ICML 2025**.
- Our proposed approach keep the same model quality in Foundation Models, with 20% faster training and 34% faster inference, reducing the company spend in energy and hardware resources for production models.

Amazon **Oct. 2023 – Feb. 2024**

Applied Science Intern

Sunnyvale, CA, United States

- Worked at AGI Foundation organization (formerly Alexa AI), on the Post-Training Amazon Nova Language Generation team. Our focus was on Benchmarking and Enhancing LLMs on Multi-constrained Instruction-Following. We also explore real world user requests and LLM-as-a-Judge performance.
- We create an human-centered benchmark that demonstrate even the best proprietary LLM (GPT-4) fails to meet user requirements in 21% of the cases.
- Our proposed LLM Self-Correction approach was able to improve Mistral v0.2 performance by up to 34%, beating GPT-4 on instruction-following benchmarks.
- Papers presented at **EMNLP 2024** and **Workshop on System 2 Reasoning at NeurIPS 2024**.

NAVER LABS Europe **Mar. 2023 – Sep. 2023**

Research Scientist Intern

Grenoble, France

- Knowledge Distillation and bias control of Multilingual Speech Models, with focus on Low-Resource Languages.
- Our method reduce by 35% de gap between small and large Speech models, being particularly effective in Low-Resource Languages.
- 2nd place Naver's Intern of the Year Award. Paper presented at **ICASSP 2024**.

Apple **Mar. 2022 – Aug. 2022**

Research Intern, NLP

Paris, France / Barcelona, Spain

- Worked on the Siri Data Augmentation Team, on the topics of **zero-shot multilingual paraphrasing**, **adversarial robustness** and **graph-to-text generation**.
- Designed graph-based cross-lingual paraphrasing and machine translation methods to generate synthetic training data for unseen domains, enhancing multilingual NLU systems.

- Developed efficient systems delivering high-quality cross-lingual paraphrasing in low-resource scenarios.

Universidade de São Paulo

Volunteer Research Assistant, NLP

Apr. 2021 – Sep. 2023

São Paulo, Brazil (remote)

- Part-time volunteer work collaborating in NLP research projects under Profs. Anna Costa and Roseli Lopes, covering the follow topics: Adversarial Robustness of Language Models; Zero-Shot Learning for Low-Resourced Scenarios; Language Models for American Indigenous Languages; Graph Neural Networks; Topic Modeling.
- Advising undergraduate interns, assisting research development, and participating in collaborations with universities abroad (Columbia, Tallinn, Porto U.).
- 7 papers approved (2 as first author), in peer-reviewed conferences and journals (see “Main Publications”).

Accenture (Former: Nectar Consulting)

Software Engineer Intern (2020) / Machine Learning Engineer (2021-22)

Jan. 2020 – Feb. 2022

São Paulo, Brazil

- Development of Machine Learning models for consumer interaction applications, deploying them into production on the Azure platform. I got Microsoft’s DP-100 Data Scientists Associate certification and other certifications related to Data and AI.
- During the internship, the work included employing best practices for design, test, deployment, and maintenance of large-scale software systems using PHP / JavaScript / CSS / HTML, modeling and handling data with MySQL, and implementing and testing REST APIs.

Escola Politécnica, Universidade de São Paulo

Research Intern / Teaching Assistant

Aug. 2015 – Jul. 2018

São Paulo, Brazil

- I worked as a paid teaching assistant in the first-year class **Algorithms and Data Structures**. I also taught reinforcement classes for other subjects voluntarily.
- I worked on three paid internship research projects, the ones that follow: Applications of Quaternions and Lie Groups to Robotics (Aug. 2015 - Jul. 2016); Definition and Analysis of Correction Automation Process for undergraduate courses (Aug. 2016 - Jul. 2017); Machine Learning Techniques applied to Genomic Analysis (Aug. 2017 - Jul. 2018).

Skills

Major Programming Languages: Python (7-year), (3-year), Hugging Face Transformers (4-year), Scikit-learn, C/C++ (4-year), MATLAB, PHP, SQL, NumPy, Pandas, Seaborn, Spark, Github, others.

General Knowledge / Frameworks and Technologies:

Natural Language Processing (5-year), Algorithms and Data Structures, OpenCV/Computer Vision, TensorFlow, PyTorch

Natural Languages: Portuguese (native), English (fluent), French (C1), Spanish (B2), Catalan (A2).

Publications

Conferences

- [To be submitted] FERRAZ T.P. and others. **Not All Layers Need Every Input: Dynamic Input Routing for Efficient Large-Scale Model Training and Serving**. ICML 2025.
- FERRAZ T.P. and others. **LLM Self-Correction with DeCRIM: Decompose, Critique, and Refine for Enhanced Following of Instructions with Multiple Constraints**. EMNLP 2024. [Paper](#). [Poster](#). [Video](#).
- FERRAZ T.P. and others. **Multilingual DistilWhisper: Efficient Distillation of Multi-task Speech Models via Language-Specific Experts**. ICASSP 2024. [Paper](#).
- CROSO, B.* , FERRAZ T.P.* , LOPES, R.D. **Enriching GNNs with Text Contextual Representations for Detecting Disinformation Campaigns on Social Media**. Learning on Graphs Conference (LoG 2024). [Paper](#). [Poster](#).
- ALCOFORADO, A., FERRAZ T.P. and others. **From Random to Informed Data Selection: A Diversity-Based Approach to Optimize Human Annotation and Few-Shot Learning**. PROPOR 2024. [Paper](#).
- ALCOFORADO, A., FERRAZ T.P. and others. **BRNews: A Decade of Brazilian History Unveiled in a Comprehensive Dataset**. PROPOR 2024.
- [To be submitted] FERRAZ T.P. and others. **Diversity Matters: Self-Training Zero-Shot Text Classification via Clustering Selection**. ACL/EMNLP/NAACL.
- ELOY, A., FERRAZ T.P. and others. **Science and engineering for what? A large-scale analysis of students’ projects in science fairs**. 17th International Conference of the Learning Sciences (ICLS 2023). June, 2023.

- ALCOFORADO, A., FERRAZ T.P. and others. **ZeroBERTo: Leveraging Zero-Shot Text Classification by Topic Modeling**. PROPOR 2022 – 15th International Conference on the Computational Processing of Portuguese. March, 2022. [Paper](#). (**Honorable Mention**)

Journals

- [Accepted] ALCOFORADO, A., FERRAZ T.P. and others. **Estratégias de Seleção Informada de Dados para Aprendizado com Dados Escassos e Desbalanceados** (*Informed Data Selection Strategies for Few-Shot Learning on Imbalanced Data*). Linguamática Journal.
- [Submitted in Aug/24] ELOY, A., FERRAZ T.P. and others. **STEM for What? Analyzing Project Themes in K-12 Science and Engineering Fairs at Scale**. International Journal of Science Education.
- [Submitted in Mar/23] FERRAZ T.P. and others. **Improving Large-Scale Speech Recognition Robustness**. Image Processing On Line - IPOL Journal.
- ALCOFORADO, A., FERRAZ T.P. and others. **Augmented Democracy: Artificial Intelligence as a Tool to Fight Disinformation**. Advanced Studies – Special Issue: AI and Democracy (h-5 index: 30). Accepted in Dec/22.
- FERRAZ T.P. and others. **Explainable AI as a tool for mitigating the lack of transparency and legitimacy in internet content moderation**. Advanced Studies – Special Issue: AI and Democracy (h-5 index: 30). Accepted in Dec/22.

Thesis

- FERRAZ T.P. **Efficient Compression of Multitask Multilingual Speech Models**. Master Dissertation, Télécom Paris (2023). [Thesis](#). [Master Defense Video](#).

Workshops / Symposia

- FERRAZ T.P. and others. **LLM Self-Correction for Multi-Constrained Instruction Following**. The First Workshop on System-2 Reasoning at Scale at NeurIPS (Sys2Reasoning at NeurIPS 2024). [Paper](#).
- FAMA, I., FERRAZ T.P. and others. **No Argument Left Behind: Overlapping Chunks for Faster Processing of Arbitrarily Long Legal Texts**. 15th Symposium in Information and Human Language Technology (STIL 2024). [Paper](#).
- FERRAZ T.P. and others. **DEBACER: a method for slicing moderated debates**. ENIAC 2021 – The 18th National Meeting on Artificial and Computational Intelligence. November, 2021.
- FERRAZ T.P. and COLON, D. **Description of Direct Kinematics of Robots using Quaternions**. 24th USP International Symposium of Undergraduate Research (SIICUSP 2016). (**Spotlight Paper**)
- FERRAZ T.P. and others. **Smartlines: Optimizing Ground Public Transport in São Paulo using Machine Learning and Computer Vision**. International Meeting for Leadership in Engineering (EILE 2016).

Major Projects

NLP for indigenous languages | Since Oct. 2022

The project aims to develop tools and models for native languages of Brazil and South America. We investigate how we can leverage lexical resources, dictionaries and other structured data to pre-train Language Models with a much smaller amount of data. We start working with discriminative tasks, but the aim is Machine Translation.

NLP-BR | Since Oct. 2021

Project with about 15 collaborators, in three universities, and aims to develop models, tools and algorithms for processing the Portuguese language. The project has generated contributions in the areas of **zero-shot learning**, **topic modeling**, **word embeddings**, **domain-specific language modeling**.

LIBRA | Since Jan. 2020

My Capstone Engineering Thesis and follow-up

I am the coordinator of this project that focuses on the development of models for moderating content on the internet in Portuguese, with a focus on fake news, hate speech and privacy invasion. Some examples of models under development:

Use of propagation characteristics to analyze publications on social networks via **Graph Neural Networks (GNNs)**; Attacks and defenses to certify **Adversarial Robustness** of language models; **Language modeling** for user-generated data on social media (BERTweet-BR). Project has received some awards like Tallinn University Grant, AWC Institute TIM, USP-Santander Innovation. We were also invited by the Brazilian Supreme Court (STF) to collaborate in the debate on social media moderation regulation in Brazil.

AI Augmented Democracy | Feb. 2021 - Dec. 2021

Project of the University of São Paulo, in partnership with Columbia University (USA) and U. Católica Lisboa (PT), which seeks to process the minutes of the Assembly of the Republic of Portugal (Portuguese Parliament) and identify the participations, opinions and feelings of the deputies in relation to proposed themes and bills, especially those involving the agenda of corruption. The project tasks include **Text Classification and Zero-Shot Learning**, **Text Summarization**, **Topic Modeling**.

Main Awards and Fellowships

- Honorable Mention as the best student in the Einstein's 2014 class;
- Bronze Medal and Merit Award at OBMEP - Brazilian Olympiad of Math (2014 and 2011);
- Best project in Introduction to Electrical Engineering (2015);
- Fellowship Instituto Semear (2015);
- 3rd place Big Data Brasil Datathon (2020);
- AWC TIM Institute (2020);
- USP Santander Innovation Award (2020);
- Top 1 of the Class 2020 of Electrical Engineering;
- ACL Student Reseach Mentorship (2021-22);
- 2nd place for Best Paper PROPOR 2022 (Honorable mention).

Leadership / Extracurricular

Paper Reviewing / Program Committee

- **ICLR:** 2025
- **ICASSP:** 2025
- **COLING:** 2025
- **AAAI:** 2025
- **ACL ARR:** 2025
- **BRACIS (Springer):** 2022, 2023, 2024
- **Workshops at NeurIPS:** 2024

Grêmio Politécnico of USP | Jan. 2019 – Jan. 2020

Political Coordinator / Legal Director

- Led Brazil's largest student association with a \$5 million annual budget, representing 7,000 members, and achieved a historic opposition election victory.
- Organized events featuring prominent figures from academia, industry, and politics, including unicorn CEOs, members of Congress, and presidential candidates.
- Contributed to the financial turnaround of Grêmio, the first administration to publish financial statements online, amidst its worst financial year in two decades.

CEE POLI-USP | Jul. 2016 – Dec. 2018

President

- Served as a student representative on the Electronic Systems Engineering Department Board and five relevant committees, leading curriculum reviews for the Engineering course
- Successfully contributed to projects with private financing, including space renovations for students, organization of academic events, and hosting the major parties at Escola Politécnica in 2017 and 2018.

Other Activities

- **ACL Student Mentorship (Since Sep. 2021)**
- **BeeData - USP's Data Science League (Sep. 2020 - Jul. 2021)**
- **TEDx USP (Jun. 2017 - Dec. 2017)**
- **Sports:** Running, Swimming, Skiing