

Gonçalo Faria

PH.D. STUDENT

✉ gfaria@cs.washington.edu | 🏠 www.goncalofaria.com | 📺 [goncalorafaria](https://www.youtube.com/channel/UCgongcalorafaria) | 📄 [goncalorafaria](https://www.linkedin.com/in/goncalorafaria) | 🐦 [@goncalorafaria](https://twitter.com/goncalorafaria)

Education

Paul G. Allen School of Computer Science & Engineering, University of Washington

Seattle, WA, USA

PH.D. IN COMPUTER SCIENCE AND ENGINEERING

2024 - now

- Advised by Noah Smith.

Instituto Superior Técnico, University of Lisbon

Lisbon, Portugal

MASTER'S DEGREE IN DATA SCIENCE AND ENGINEERING

2019 - 2021

- Was part of DeepSPIN research group (<https://deep-spin.github.io/>) led by André Martins.
- Collaborated closely with Professors **Mário Figueiredo and André Martins** for my master's dissertation, focusing on the interplay between causality and machine learning.
- Researched on latent interventions to decode cause-effect relationships, blending causal discovery, deep latent variable models, structured prediction, and approximate posterior inference techniques. Developed a method that leverages neural networks and variational inference to uncover causal relationships in environments with entirely latent interventions. The method was published in Faria, Martins, & Figueiredo (2022).
- Dissertation received a (20/20) grade (<https://arxiv.org/abs/2203.02336>).
- 18/20 average, class A in ECTS grading scale.

University of Minho

Braga, Portugal

BACHELOR DEGREE IN COMPUTER SCIENCE ENGINEERING

2016 - 2019

- 17/20 average, class A in ECTS grading scale

Professional Experience

2024 - now **Researcher Assistant**, University of Washington

- Working on test-time alignment
- ARK Lab
- Advised by Noah Smith

2023 - 2024 **NLP Researcher**, Instituto de Telecomunicações

- I am part of the alignment teams working on the TowerLLM project.
- This was worked done with **André Martins**, as part of the Sardine group. It resulted in two publications.
- NEURIPS2024 <https://arxiv.org/abs/2406.00049>
- EMNLP2024 <https://arxiv.org/abs/2410.07779>

- 2021 - 2023 **Lead ML Engineer**, Dotmoovs
- **Initial Role & MVP Development:**
 - System for counting soccer juggles using object detection and pose estimation.
 - Enabled users to upload videos and receive a count of their soccer ball juggles with different body parts.
 - **Post-funding Expansion:**
 - Spearheaded the assembly and led an ML engineering team.
 - Bolstered the model serving stack and rectified MVP's public launch edge cases.
 - Incorporated multi-object tracking, and finetuned SOTA vision models using annotated user data.
 - Led the establishment of data engineering pipelines to support business intelligence applications and develop recommender systems.
 - **Enhancing Activity Recognition Technology:**
 - Improved kinematics models to minimize juggle-counting errors.
 - Introduced expanded capabilities to recognize more sports and activities.
 - Utilized annotated user data for training juggle and freestyle football trick detection.
 - **Fitness, and Support:**
 - Expanded the activity recognition capabilities to over 100 fitness exercises.
 - Introduction of NLU pipelines for the personal trainers' support channels.
 - Using an LLM provider, created an AI coach to help the user with in-app navigation and training assistance.
- 2021 - 2022 **Invited Teaching Assistant**, Instituto Superior Técnico
- I was part of the team that taught the first MSc class exclusively about deep learning at IST. It was a valuable experience.
- 2021 - 2021 **Student Researcher**, Instituto de Telecomunicações
- Was part of DeepSPIN research group (<https://deep-spin.github.io/>) led by André Martins.
 - Collaborated closely with Professors **Mário Figueiredo and André Martins** for my master's dissertation, focusing on the interplay between causality and machine learning.
 - Researched on latent interventions to decode cause-effect relationships, blending causal discovery, deep latent variable models, structured prediction, and approximate posterior inference techniques. Developed a method that leverages neural networks and variational inference to uncover causal relationships in environments with entirely latent interventions. The method was published in Faria, Martins, & Figueiredo (2022).
- 2020 - 2021 **Student Researcher**, Unbabel
- I was involved in addressing the challenge of selecting the correct stylistic register when performing machine translation. This was a significant issue because while some languages (e.g., English) do not differentiate between stylistic registers, this distinction is crucial in other languages (e.g., Portuguese, where "you" can be translated as "tu" or "você"). The work I did was train transformers to perform register selection and conversion for many languages. We evaluate our system both intrinsically using the labeled dataset and for style transfer using Unbabel's own robust human evaluation process and, extrinsically, the system's quality as part of Unbabel's translation pipeline. Overall, the models worked well and improved the MT results on all the language pairs we worked on.
 - The project moved into production and is now being used by many users.
 - This was worked done with **Christine Maroti**, and **André Martins**.

- 2018 - 2019 **Research Scholar**, Fundação Calouste Gulbenkian - Programa Novos Talentos Inteligência Artificial
- *In my undergraduate final year, I was fortunate to be selected for the “New Talents in Artificial Intelligence” program by the Gulbenkian Foundation, which provided a unique opportunity to delve into Deep Learning research whilst still an undergraduate. The program brought together like-minded individuals from universities across the country and facilitated collaborations with leading Portuguese researchers in the field. Through this platform, I had the privilege to meet some of the pioneers in the field, such as Stuart Russell, and to start my own research journey where I worked on my first research project, which tried to address a few limitations of what was then a new and exciting idea called Capsule Networks.*
 - **Project:** *Equivariant models for discriminative learning with unstructured data*(<https://arxiv.org/abs/1909.12737>)
 - **Presented at:** INTERNATIONAL CONFERENCE ON OPTIMIZATION AND LEARNING 2020
 - **Tutor:** Professor Doutor Cesar Analide
- 2016 - 2016 **Supermarket Cashier**, Modelo Continente Hipermercados SA, Guimarães, Portugal
- *The lesson I took away was a deep understanding that every customer interaction is an opportunity to make a lasting impression. This foundational experience became a cornerstone for my entrepreneurial ambitions, emphasizing that an essential component of success lies in valuing and understanding every individual we serve.*

Publications

PUBLISHED

- Sweta Agrawal, José G. C. de Souza, Ricardo Rei, António Farinhas, **Gonçalo R. A. Faria**, Patrick Fernandes, Nuno M Guerreiro, Andre Martins; Modeling User Preferences with Automatic Metrics: Creating a High-Quality Preference Dataset for Machine Translation. **EMNLP 2024.**
- Gonçalo R. A. Faria**, Sweta Agrawal, António Farinhas, Ricardo Rei, José G. C. de Souza, André F. T. Martins; QUEST: Quality-Aware Metropolis-Hastings Sampling for Machine Translation. **NEURIPS 2024.**
- Gonçalo R. A. Faria**, André F. T. Martins, Mário A. T. Figueiredo; Differentiable Causal Discovery Under Latent Interventions. **CLEAR 2022.**

Awards

- 2024 **Fully Funded Ph.D. Scholarship**, Paul G. Allen School of Computer Science & Engineering
Tuition and stipend to pursue a Ph. D. in Computer Science at Paul G. Allen School of Computer Science & Engineering, University of Washington.
- 2022 **APRP Prize - Best Master Thesis**, Portuguese Pattern Recognition Association
Associação Portuguesa de Reconhecimento de Padrões(APRP) establishes the APRP prize for the best master’s thesis, with the purpose of distinguishing work of high merit in the area of pattern recognition in Portugal.
- 2022 **Prémio Prof. Luís Vidigal**, IST (Instituto Superior Técnico) - University of Lisbon
It is intended to award Masters of any master’s degree at IST annually, authors of the best Master’s Thesis whose theme falls within the scientific areas of Electrical and Computer Engineering and Informatics.
- 2020 **Prémio Almedina**, University of Minho
Awarded annually to the best scholarship student of each Organic Teaching and Research Unit.

Invited Talks

- Causal Data Science Meeting 2022 : Causal Discovery under Latent Interventions
- Deep Learning Sessions Portugal : Causal Discovery and Deep Learning
- Causality Discussion Group : Causal Discovery under Latent Interventions

Teaching Experience

- 2023 **13th Lisbon Machine Learning Summer School(LxMLS)**, Lab Monitor *IST*
<http://lxmls.it.pt/2023/>
- 2022 **Deep Learning**, Invited Teaching Assistant *IST*
- 2021 **11th Lisbon Machine Learning Summer School(LxMLS)**, Lab Monitor *IST*
<http://lxmls.it.pt/2021/>
- 2017 **CeSIUM - University of Minho Computer Engineering Students Center**, Colaborator in the Pedagogical Department *Braga*

Complementary Education

- 2018 **Structuring Machine Learning Projects**, Andrew Ng via Coursera (deeplearning.ai)
An online non-credit course authorized by deeplearning.ai and offered through Coursera. Credential ID: 6NNX58JYWP2L
- 2018 **Improving Deep Neural Networks: Hyperparameter tuning, Regularization and Optimization**, Andrew Ng via Coursera (deeplearning.ai)
An online non-credit course authorized by deeplearning.ai and offered through Coursera. Credential ID: EAMS2XZZ8MBD
- 2018 **Neural Networks and Deep Learning**, Andrew Ng via Coursera (deeplearning.ai)
An online non-credit course authorized by deeplearning.ai and offered through Coursera. Credential ID: JNQJW39U2H6X
- 2017 **Machine Learning**, Andrew Ng via Coursera (Stanford University)
An online non-credit course authorized by Stanford University and offered through Coursera. Credential ID: 87CHL2E4476N