Tiago Pimentel Martins da Silva

Postdoc at ETH Zürich
Brauerstrasse, 27
Zürich, Switzerland

tpimentelms@gmail.com
Phone: +41 76 593 33 45
https://tpimentelms.github.io/
Google Scholar

Education

University of Cambridge *Oct. 2019 – Oct. 2023*
PhD in Computer Science
Advisors: Simone Teufel and Ryan Cotterell
Thesis: *On the Optimality of the Lexicon*
★ **Gates** Cambridge Scholar
★ **Facebook Fellow**

Universidade Federal de Minas Gerais (UFMG) *GPA*: 4.0/4.0 *March 2016 – Feb. 2018*
Master in Computer Science Advisors: Adriano Veloso and Nivio Ziviani

Universidade de Brasília (UnB) *GPA*: 4.669/5.0 *March 2009 – Dec. 2014*
B.E., Mechatronics Engineering Advisor: Carla Koike ★ **Valedictorian**, highest GPA in class

Research Interests

Interpretability and Analysis of Language Models

- I ask questions such as: (i) what does it mean for a model to understand language? or (ii) how does this model process language? I also use information theory, statistics and causality to analyse specific properties of language models, such as memorisation or tokenisation bias.

Computational and Psycho-linguistics

- I use methods from NLP, coupled with ideas from information theory, to analyse how humans read sentences and how communicative pressures shape natural languages.

Language Modelling and Generation

- I propose and analyse methods to train and decode text from language models. I also investigate how the way we encode language models' inputs (i.e., how we tokenise text) affects their performance.

Selected & Awarded Publications

The Non-Linear Representation Dilemma: Is Causal Abstraction Enough for Mechanistic Interpretability?
NeurIPS 2025 (spotlight). *D. Sutter, J. Minder, T. Hofmann, T. Pimentel*

Tokenisation is NP-Complete
ACL 2025 (★ **senior area chair highlights** award). *P. Whittington, G. Bachmann, T. Pimentel*

Using Information Theory to Characterize Prosodic Typology: The Case of Tone, Pitch-Accent and Stress-Accent
ACL 2025 (★ **senior area chair highlights** award). *E. Wilcox, C. Ding, G. Acampa, T. Pimentel, A. Warstadt, T. Regev*

Causal Estimation of Memorisation Profiles
ACL 2024 (★ **best paper** award). *P. Lesci, C. Meister, T. Hofmann, A. Vlachos, T. Pimentel*

Large-Scale Evidence for Logarithmic Effects of Word Predictability in Naturalistic Reading
PNAS 2024. *C. Shain, C. Meister, T. Pimentel, R. Cotterell, R. Levy*

Revisiting the Optimality of Word Lengths
EMNLP 2023 (★ **outstanding paper** award). *T. Pimentel, E. Wilcox, C. Meister, R. Cotterell, K. Mahowald*

Language Model Quality Correlates with Psychometric Predictive Power in Multiple Languages
EMNLP 2023 (★ **outstanding paper** award). *E. Wilcox, C. Meister, R. Cotterell, T. Pimentel*

Information-Theoretic Probing for Linguistic Structure
ACL 2020. *T. Pimentel, J. Valvoda, R. Maudslay, R. Zmigrod, A. Williams, R. Cotterell*

Meaning to Form: Measuring Systematicity as Information
ACL 2019 (★ nominated for best paper award). *T. Pimentel, A. McCarthy, D. Blasi, B. Roark, R. Cotterell*

Work Experience

Postdoc – ETH Zürich *Oct. 2023 – Present*

- Research on natural language processing and on machine learning interpretability.
- Advised 8 students (5 Master's, 3 Bachelor's theses); 5 resulted in publications at **NeurIPS 2025** (spotlight), **ACL 2024** (Findings), **TACL 2024**, **EMNLP 2024** (Findings), **EMNLP 2025**.

Research Intern – Google *Aug. 2022 – Oct. 2022*

- Research project on the structure with which metaphors are encoded in language models

Research Intern – Google *June 2021 – Sep. 2021*

- Research project on applying upside-down reinforcement learning in recommendation systems
- Contributed to the open source Jax Recommenders python library

Data Scientist – Kunumi *March 2018 – Sep. 2019*

- Designed an active anomaly detection tool for an insurance company client
- Started an inner development program to capacitate employees in deep learning and NLP

Teaching Experience

Tutorial Organiser

Generating Text from Language Models
ACL 2023. *A. Amini, R. Cotterell, J. Hewitt, C. Meister, L. Malagutti, T. Pimentel*

Information Theory in Linguistics: Methods and Applications
COLING 2022 and **ESSLLI 2021**. *R. Cotterell, R. Futrell, K. Mahowald, C. Meister, T. Pimentel, A. Williams, and A. Arora*

Adjunct Lecturer – IESB *Aug. 2018 – April 2019*

Lectured *Unsupervised Learning* and *Reinforcement Learning* in the AI postgraduate specialisation.

Voluntary Teacher – EduBot *March 2014 – June 2015*

Taught computer science and robotics to high school students from under-resourced public schools.

Teaching Assistant

University of Cambridge (2020–2022), UFMG (2017), UnB (2010–2011)

Selected Invited Talks

Promises and Limitations of Causality for ML Interpretability *2026*
At: EPFL, University of Tübingen, University of Oxford, University of Edinburgh, UT Austin

From Surprisal to Processing Difficulty *2024–2025*
At: Universitat Pompeu Fabra, Ev Fedorenko's Lab @ MIT, UT Austin, University of Edinburgh

An Information-theoretic Analysis of Language Models' Representations *2023*
At: University of Cambridge, IST Austria, Roger Levy's Group @ MIT, INSAIT

Revisiting the Optimality of Word Lengths *2024*
At: ETH Zürich, University of Cambridge, IT University of Copenhagen, LIRMM/CNRS

Academic Service

Workshop Organization. The First Workshop on LLM Memorization (L2M2) @ ACL 2025
Shared Task Organization. SIGMORPHON's Shared Task 0 in 2023, 2021, and 2020
Senior Area Chair. ACL 2025 & 2026, EMNLP 2026 (Interpretability and Analysis of Models for NLP)
Area Chair. ARR 2024–2025 (6 cycles), EMNLP 2023 (★ best area chair award), CoNLL 2026, CogSci 2026
Paper Reviews. ACL, EMNLP, NAACL, EACL, CoNLL, NeurIPS, ICLR, ICML, CoLM, Cognition, and more

References

Roger Levy
Professor
MIT
rplevy@mit.edu

Kyle Mahowald
Assistant Professor
UT Austin
mahowald@utexas.edu

Simone Teufel
Professor
University of Cambridge
sht25@cl.cam.ac.uk

Thomas Hofmann
Professor
ETH Zürich
thomas.hofmann@inf.ethz.ch