Ben Peters

Curriculum Vitae

Education

2018–2022 **PhD Electrical and Computer Engineering**, *Instituto Superior Técnico, University of Lisbon*, Lisbon
With Honour and Distinction
Thesis title: Sparse and Linguistically Informed Sequence-to-Sequence Modeling
Advisor: André F. T. Martins

2015–2018 **MSc. Language Science and Technology**, *Saarland University*, Saarbrücken
Excellent
Thesis title: Massively Multilingual Neural Grapheme-to-Phoneme Conversion
Advisors: Jon Dehdari and Josef van Genabith

2010–2014 **BA Linguistics**, *University of Washington*, Seattle
Cum Laude, with honors in Linguistics (specialization: computational syntax)
Thesis title: Building a Precision Grammar of Meithei
Advisor: Emily M. Bender

Work Experience

2025–present **Researcher**, *INESC-ID, Lisbon*
I research speech-text integration for AMALIA, a European Portuguese LLM. This includes maintaining codebases for training and evaluating LLMs, creating and curating datasets for speech-centric and tasks, organizing external collaborations, applying for grants, and supervising junior students.

2022–2025 **Post-doctoral Researcher**, *Instituto de Telecomunicações, Lisbon*
I researched speech translation, robust machine translation, and speech-text LMs within UTTER, a multi-institution Horizon Europe project. In the course of the project, I contributed to Tower, a widely-used LLM for translation, and led the development of Spire, an LLM specialized for ASR and speech translation.

2018–2022 **PhD Student Researcher**, *Instituto de Telecomunicações, Lisbon*
I researched machine translation, character-level NLP, decoding strategies, and sparsity in neural networks within DeepSPIN, an ERC-funded project about deep structured prediction in NLP. This work led to publications in top-tier venues and multiple shared task wins.

2017–2018 **Scientific Research Assistant**, *German Research Center for Artificial Intelligence (DFKI), Saarbrücken*
I worked on the Quality Translation 21 and DeepLee projects, focusing on open-source development of a neural machine translation library. I was supervised by Josef van Genabith.

2015–2017 **Scientific Research Assistant**, *Saarland University, Saarbrücken*
I provided technical and linguistic help in a computational psycholinguistics lab at Saarland University. My work included writing data processing software for eyetracking experiments and tuning parsing systems. I was supervised by Vera Demberg.

benzurdopeters@gmail.com • *bpopeters.github.io*
ben-peters-b29128a0 • *bpopeters* • *0000-0003-3317-5686*
WFcYhToAAAAJ

2014–2015 **Contract Annotator**, *Pactera Technology*
I provided intent and entity annotations of transcribed conversations between Microsoft Cortana and users in order to supply data to the development team.

Teaching Experience

2021–present **Assistente Convidado (Invited Assistant Professor)**, *Department of Informatics, Instituto Superior Técnico*
I deliver lectures and design assignments for masters-level courses in Deep Learning.

2019–present **Senior Monitor**, *Lisbon Machine Learning School (LxMLS)*
I organize, teach, and design curriculum for one of Europe's largest annual machine learning summer schools.

MSc Students Advised

2026–present César Camus-Emschwiller. Continuous representations for Portuguese speech-text integration (with Alberto Abad and Bruno Martins).

2025–present Artur Monteiro. Discrete representations for Portuguese speech-text integration (with Alberto Abad and Bruno Martins).

2024–present Anil Keshwani. Discrete representations for speech LMs (with André F.T. Martins).

Conference Publications

2026 Francisco Teixeira, Carlos Carvalho, Mariana Julião, Catarina Botelho, Rubén Solera-Ureña, Sérgio Paulo, Thomas Rolland, **Ben Peters**, Isabel Trancoso, and Alberto Abad. "FalAR: A Large-scale Speaker-Annotated European Portuguese Speech Corpus of Parliamentary Sessions". In proc. LREC.

2025 **Ben Peters***, Kshitij Ambilduke*, Sonal Sannigrahi*, Anil Keshwani, Tsz Kin Lam, Bruno Martins, André F. T. Martins, and Marcely Zanon Boito. 2025. "From Tower to Spire: Adding the Speech Modality to a Translation-Specialist LLM". In Findings of EMNLP.

2025 **Ben Peters** and André F. T. Martins. 2025. "Did Translation Models Get More Robust Without Anyone Even Noticing?". In proc. ACL.

2024 "Tower: An Open Multilingual Large Language Model for Translation-Related Tasks". Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, **Ben Peters**, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G.C. de Souza, and André F.T. Martins. In proc. COLM.

2021 **Ben Peters** and André F. T. Martins. 2021. "Smoothing and Shrinking the Sparse Seq2Seq Search Space". In proc. NAACL.

2019 **Ben Peters**, Vlad Niculae, and André F. T. Martins. 2019. "Sparse Sequence-to-Sequence Models". In proc. ACL.

benzurdopeters@gmail.com • *bpopeters.github.io*
ben-peters-b29128a0 • *bpopeters* • *0000-0003-3317-5686*
WFcYhToAAAAJ

Workshop Publications

2025 Giuseppe Attanasio, Sonal Sannigrahi, **Ben Peters**, and André F. T. Martins. 2025. "Instituto de Telecomunicações at IWSLT 2025: Aligning Small-Scale Speech and Language Models for Speech-to-Text Learning". In proc. IWSLT.

2022 **Ben Peters** and André F. T. Martins. 2022. "Beyond Characters: Subword-level Morpheme Segmentation". In proc. SIGMORPHON. **Shared Task winner**.

2020 **Ben Peters** and André F. T. Martins. 2020. "One-Size-Fits-All Multilingual Models". In proc. SIGMORPHON. **Shared Task winner**.

2019 **Ben Peters** and André F. T. Martins. 2019. "Sparse Two-Headed Models for Inflection". In proc. SIGMORPHON. **Shared Task Interpretability Prize winner**.

2017 **Ben Peters**, Jon Dehdari, and Josef van Genabith. 2017. "Massively Multilingual Neural Grapheme-to-Phoneme Conversion". In proc. BLGNLP.

Publications Under Review

2026 Emmanouil Zaranis*, António Farinhas*, Saul Santos*, Beatriz Canaverde*, Miguel Moura Ramos*, Aditya K Surikuchi, André Viveiros, Baohao Liao, Elena Bueno-Benito, Nithin Sivakumaran, Pavlo Vasylenko, Shoubin Yu, Sonal Sannigrahi, Wafaa Mohammed, **Ben Peters**, Danae Sánchez Villegas, Elias Stengel-Eskin, Giuseppe Attanasio, Jaehong Yoon, Stella Frank, Alessandro Suglia, Chrysoula Zerva, Desmond Elliott, Mariella Dimiccoli, Mohit Bansal, Oswald Lanz, Raffaella Bernardi, Raquel Fernández, Sandro Pezzelle, Vlad Niculae, and André F. T. Martins. "Movie Facts and Fibs (MF^2): A Benchmark for Long Movie Understanding".

Academic Service

2024–2025 **FSTP Project Reviewer and Sponsor**, *UTTER Project*
I reviewed applications for Financial Support for Third Parties (FSTP) grants. As sponsor of one funded project and co-sponsor of another, I advised grant recipients and ensured that they delivered their project goals on time.

2023–2024 **PhD Application Evaluator**, *ELLIS PhD Program*
I screened applications for the ELLIS PhD Program, providing recommendations for PIs of which applicants to interview.

2019–present **Program Committee**, *Association for Computational Linguistics*
I frequently serve as a reviewer for ACL Rolling Review (ARR). Before the introduction of ARR, I reviewed for several editions of ACL, EMNLP, NAACL, and EACL.

Open Source Projects

SPITE, *https://huggingface.co/collections/bpop/spite*
Speech Pseudolabeling for Instruction-based Translation Ensembling (SPITE) is a growing speech translation corpus containing high-quality synthetic translations from several models, as well precomputed quality metrics to enable rapid filtering.

SpireLM, *https://github.com/utter-project/SpireLM*
I am the primary developer of SpireLM, a library that provides fast implementations for learning and applying discrete speech tokenizers. Although the project began as a repository for the "Spire" paper, development is ongoing and the package now includes additional techniques.

benzurdopeters@gmail.com • *bpopeters.github.io*
ben-peters-b29128a0 • *bpopeters* • *0000-0003-3317-5686*
WFcYhToAAAAJ

entmax, *https://github.com/deep-spin/entmax*

I maintain the entmax package, which contains efficient PyTorch implementations of entmax, a family of softmax alternatives, and their corresponding loss functions.

OpenNMT-py, *https://github.com/OpenNMT/OpenNMT-py*

I maintained and contributed to the Python version of the OpenNMT system for neural machine translation. My contributions include implementing factored word embeddings and adding new features for attention mechanisms.

Technical Skills

Python Programming, *10+ years experience*

PyTorch, NumPy, transformers, Axolotl, vLLM, Matplotlib, scikit-learn, pandas, NLTK

High-Performance Computing, *2 years experience*

I frequently run large jobs on HPC clusters such as MareNostrum 5 and Helios. In order to do so, I have become an expert in **containerization** tools such as Singularity and Docker, and in the challenges of **multi-node** GPU training and inference.

GPU Kernel Programming, *2 years experience*

I use Triton to write optimized kernels for deep learning functions and layers.

Workflow Management, *2 years experience*

I am a power user of ducttape, a tool which allows complex, multi-configuration pipelines (preprocessing, training, inference, evaluation, etc.) to be codified as a dependency graph. This allows massive arrays of jobs to be run with a single command, reducing required effort and the risk of human error.

Miscellaneous

I have years of experience with shell scripting and programming in Java.

Languages

English	Native	Spanish	Proficient
Portuguese	Proficient	Turkish	Basic skills
German	Proficient	Arabic	Basic skills

benzurdopeters@gmail.com • *bpopeters.github.io*
ben-peters-b29128a0 • *bpopeters* • *0000-0003-3317-5686*
WFcYhToAAAAJ