Bernardo Torres

bernardo.torres@telecom-paris.fr
https://bernardo-torres.github.io | github.com/bernardo-torres | linkedin | Google Scholar

PhD candidate in machine learning for audio at Telecom Paris (Institut Polytechnique de Paris), with 6 peer-reviewed publications in TASLP, TISMIR, ICASSP, and ISMIR. Research expertise in self-supervised and unsupervised learning, music source separation, differentiable DSP, and generative audio models.

Education

Telecom Paris, Institut Polytechnique de Paris
Paris, France

Ph.D. in Machine Learning and Signal Processing for Audio and Music — 2023 – Jul. 2026 (expected)

- Thesis: *Audio Analysis with Inductive Biases and Geometric Priors*
- Research in unsupervised and self-supervised learning for audio analysis, music source separation, and music information retrieval.
- Advisors: Prof. Geoffroy Peeters, Prof. Gaël Richard

École Normale Supérieure Paris-Saclay
Gif-sur-Yvette, France

M.S. in Computer Science and Applied Mathematics (MVA). **Highest honors**. — 2021–2022

- MVA (Mathematics, Vision and Learning) — top AI research master in France.
- Coursework: Deep Learning, Time Series, Audio & Speech Processing, Computer Vision.

Telecom Paris, Institut Polytechnique de Paris
Paris, France

Engineering Diploma & B.S. in Computer Science. **GPA: 4.0/4.0**. — 2020–2021

- Majors: Machine Learning, Signal Processing, and Embedded Systems.

Federal University of Minas Gerais (UFMG)
Belo Horizonte, Brazil

B.S. in Electrical Engineering — 2016–2022

- Electrical, Electronics, and Computer Engineering.

Work Experience

Deezer
Paris, France

Research Scientist Intern — Apr. 2025 – Jul. 2025

- Investigated consistency and diffusion model objectives for audio autoencoder training
- Developed an augmentation-based technique to induce desired geometric properties in the latent space of audio autoencoders; published at ICASSP 2026.
- Advisor: Gabriel Meseguer-Brocal

Sony Computer Science Laboratories
Paris, France

Research Scientist Intern — Apr. 2022 – Sep. 2022

- Developed and benchmarked self-supervised foundation models for singing voice representation learning, exploring contrastive, VICReg, and uniformity-alignment objectives; published at ISMIR 2023.
- Advisor: Stefan Lattner

Radix Engineering and Software
Belo Horizonte, Brazil

Data Science Intern — Jul. 2019 – May 2020

- Developed and deployed LSTM networks for multivariate time series analysis and anomaly detection in industrial sensor data.

Publications

Journal Articles

- **Torres, B.**, Peeters, G. and Richard, G., The Inverse Drum Machine: Source Separation Through Joint Transcription and Analysis-by-Synthesis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, 2025.
- Riou, A., **Torres, B. (co-first author)**, Hayes, B., Lattner, S., Hadjeres, G., Richard, G., and Peeters, G., PESTO: Real-Time Pitch Estimation with Self-supervised Transposition-equivariant Objective. *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 2025.

Conference Papers

- **Torres, B.**, Moussallam, M., and Meseguer-Brocal, G., Learning Linearity in Audio Consistency Autoencoders via Implicit Regularization. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2026.

- **Torres, B.**, Peeters, G. and Richard, G., Unsupervised Harmonic Parameter Estimation Using Differentiable DSP and Spectral Optimal Transport. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.
- Richard, G., Chouteau, P. and **Torres, B.**, A Fully Differentiable Model for Unsupervised Singing Voice Separation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.
- **Torres, B.**, Lattner, S. and Richard, G., Singer Identity Representation Learning using Self-Supervised Techniques. *International Society for Music Information Retrieval Conference (ISMIR)*, 2023.

Teaching Experience

Telecom Paris, Institut Polytechnique de Paris
Paris, France

Teaching Assistant
2023–Present

Deep Learning I — IP Paris Data Science M.S. program
TSIA 201 Signal Processing
TSIA 206 Speech and Audio Processing
TSIA 203 Introduction to Deep Learning

Awards & Scholarships

CAPES Foundation (Brazil)
BRAFITEC Excellence Double-Degree Scholarship
2020–2022

Skills

Deep Learning: PyTorch, TensorFlow; self-supervised learning, contrastive learning, generative models (diffusion, VAE, flow-matching)
Audio & Signal Processing: torchaudio, librosa, SciPy, differentiable DSP, spectral analysis, music source separation, music information retrieval (MIR)
Scientific Computing: Python (NumPy, SciPy, Matplotlib, Jupyter, Pandas), C, C++
Research Infrastructure: Git, Docker, Slurm, Hydra, Weights & Biases, Linux
Spoken Languages: Portuguese (native), English (fluent, C2), French (fluent), Spanish (B2)
Interests: Music, sound synthesis, bioacoustics, electronic music production and mixing