Paula Peixoto Moledo, PhD

+34 637 092 889 · paulapeixotomoledo@gmail.com · www.linkedin.com/in/paula-peixoto-moledo-phd-0107ba15a · github.com/paulapeixotomoledo · scholar.google.com/citations?user=OU_nTaIAAAAJ · Madrid, Spain — willing to relocate to San Francisco

Computational neuroscientist who reverse-engineers the internal mechanisms of systems that were never designed to be understood. I build probabilistic and deep generative models to recover latent structure from high-dimensional neural and behavioral recordings, then use causal interventions and attribution analysis to establish which internal components actually *drive* observed behavior rather than merely correlate with it. I work across the full stack: experimental design, GPU/cluster inference pipelines, and the visualization infrastructure that makes results legible to collaborators.

Keywords: Python | PyTorch | Causal intervention & attribution | Simulation-based inference | Normalizing flows | Latent-variable models | Neural decoding | GPU/HPC pipelines | Research infrastructure

SELECTED RESEARCH PROJECTS

Causal attribution in a biophysical circuit model

Simulation-based inference to identify the mechanism behind a pharmacological effect — manuscript in preparation

- Built an 8-parameter per-cell inference pipeline (sequential neural posterior estimation with a neural spline flow, 50,000 simulations) and ran it on ~25,000 neurons across two paradigms, nine time points and two treatment groups, on GPU and an HPC cluster.
- **Designed a counterfactual intervention framework to get from correlation to cause:** shift one inferred parameter by the group-level delta, re-simulate, and measure how much of the observed group difference it recovers. Shapley-decomposed the contributions to separate individual effects from interactions.
- Showed that the attribution **inverts** with the choice of counterfactual: parameter replacement and parameter shift assign the effect to opposite mechanisms. Diagnosed why, and kept the failing version in the report as a documented pitfall. The same choice governs ablation methodology in network interpretability.
- Built the supporting infrastructure: resumable GPU inference with checkpointing, cached simulation and posterior artifacts, and an interactive HTML report collecting the analyses, figures and methodological decisions in one place.

Mechanistic interpretability of transformer attention heads

Personal project, August 2026 — github.com/paulapeixotomoledo/interp-induction-shapley

- Replicated induction behaviour in GPT-2 small, then ablated the induction heads individually and across all 32 subsets under three ablation counterfactuals, decomposing the result with exact Shapley values.
- Marginal effects recover only a small fraction of the joint effect; most of it is interaction that single-head ablation cannot observe by construction. Extended the same treatment to the previous-token heads that feed the circuit, and to a cumulative sweep with a matched control at every step.
- **Three controls changed the conclusions.** One showed a head's importance is contextual rather than intrinsic: the same head costs twenty times more ablated last than ablated first. One showed the headline interaction figure was substantially an artefact of measuring in log space. A seed sweep overturned a finding I had dismissed as noise.
- Published with the corrections in the commit history rather than rewritten out of it.

Latent-state decomposition of behavior

Hierarchical GLM-HMM recovering interpretable strategies from navigation data

- Extended standard GLM-HMM architectures with custom observation models for circular data, plus optimization procedures to improve inference stability.
- Performed systematic model selection over K = 1–10 latent states using K-fold cross-validation, leave-one-animal-out validation, and stratified/balanced resampling with warm-start initialization.
- Recovered latent states that each map onto an identifiable behavioural strategy rather than remaining an uninterpretable mixture, decomposing a complex system into named, individually inspectable components.
- Built the figure-generation and reporting pipeline used to communicate the model and its diagnostics to collaborators.

Year-scale population decoding from durable single-neuron recordings

Second author, IEEE MetroXRAINE 2026 (accepted) — 7 animals, 30–42 chronic imaging sessions each

- Built the end-to-end decoding pipeline from raw imaging data: non-rigid motion correction, cross-session single-cell registration tracking the same identified neurons across the full recording period, neuron-count normalization by bootstrap subsampling, and stratified cross-validation.
- Benchmarked multiple decoder architectures under matched evaluation (ridge regression, population-vector decoding, gradient-boosted trees) and diagnosed *why* the weaker performers failed rather than only ranking them.
- **Designed the adversarial control suite written in response to review.** A reviewer argued the accuracy could reflect temporal autocorrelation rather than genuine coding, so I built four independent controls to break my own result: block cross-validation with enforced temporal gaps, a circular time-shift lag sweep, per-neuron permutation, and label shuffling.
- The controls settled more than the objection: they separated population-level coding from single-unit tuning, established temporal specificity, and identified which of two candidate null distributions was the correct one. They also surfaced a confound in my own analysis, which I reported: it invalidated one class of comparison while leaving the main result intact.

RESEARCH EXPERIENCE

Post-Doctoral Researcher — Systems Neuroscience

Computational Neuroscience of Memory Lab, Cajal Neuroscience Center – CSIC, Madrid, Spain
November 2025 – Present

- Develop and maintain scalable Python pipelines for multimodal neural/behavioral analysis (reproducible preprocessing, model training, evaluation, visualization).
- Implement and benchmark latent-state inference frameworks to quantify strategy dynamics and learning.
- Collaborate across experimental and computational teams; contribute to manuscript preparation, figures, and method write-ups.

PhD Researcher — Systems Neuroscience

Laboratory of Functional and Systems Neurobiology, Instituto Cajal – CSIC, Madrid, Spain **Laboratory for the study of the pathogenesis of immune-mediated neuronal disorders, IDIBAPS – Hospital Clínic**, Barcelona, Spain *September 2020 – November 2025*

- Led end-to-end research on memory interference and adaptive behavior, integrating experimental neuroscience with probabilistic and machine-learning–based modeling.
- Developed probabilistic models to infer latent behavioral states, and simulation-based inference for parameter estimation, model comparison, and predictive evaluation, in PyTorch and TensorFlow.
- Fine-tuned and executed established high-throughput behavioral paradigms, and designed the experiments run on them: treatment groups, pharmacological manipulations, controls and timepoints, ensuring reproducible data collection under controlled conditions.
- Acquired and analyzed in vivo calcium imaging data. Separately, cleaned and analyzed preliminary fiber photometry recordings acquired by colleagues, including synchronisation with the behavioral record.
- Applied pose estimation (DeepLabCut) and automated feature extraction to convert high-dimensional behavioral recordings into model-ready representations.
- Built integrated hardware–software systems (Arduino) to synchronize behavioral task events with imaging acquisition via hardware triggers and timestamped event logs.
- Mentored Bachelor's and Master's students on thesis projects: experimental design, analysis pipelines, and scientific writing.

Research Technician — Systems Neuroscience

IDIBAPS – Hospital Clínic, Barcelona, Spain *June 2019 – September 2020*

- Experimental and analytical research on immune-mediated synaptopathies affecting memory and behavior, including behavioral experiments, pharmacological interventions, and histological analysis.

Earlier Research Positions

- **Research Intern**, Laboratory of Paediatric Neurology, Vall d'Hebron Research Institute — MSc thesis research *(2017–2018)*
- **Research Intern**, Cellular Biology, Physiology and Immunology, UAB — BSc thesis research *(2016–2017)*
- **Research Intern**, Group of Nanocomposites, IMDEA Nanoscience, Madrid *(2015)*

EDUCATION

PhD in Biomedicine (Neuroscience) — Universitat de Barcelona, *October 2020 – November 2025* Thesis: "Interference-Based Forgetting in Spatial Navigation: Latent Behavioral Signatures and the Impact of Synaptic Plasticity Disruption."

MSc Translational Biomedical Research — Vall d'Hebron Campus / UAB, *2017 – 2018*

BSc Nanoscience and Nanotechnology — UAB, *2013 – 2017*

PUBLICATIONS AND CONFERENCES

- P. Jercog, **P. Peixoto-Moledo**, T. Rogerson, V. Minces. *Many Neurons, Many Months: Durable Single-Neuron Optical Measurement Supporting Year-Scale Spatial Decoding*. 2026 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (**IEEE MetroXRAINE 2026**), Chemnitz, Germany, October 2026. **Accepted; to appear in IEEE Xplore.** *(Second author.)*
- **P. Peixoto-Moledo**, Y. Ahmadian, & P.E. Jercog (2026). Poster, FENS Forum meeting — *"Strategy-dependent manifestation of memory recall under interference."*
- **P. Peixoto-Moledo**, Y. Ahmadian, & P.E. Jercog (2026). Poster, Systems-IN-Action meeting — *"Interference-Based Forgetting in Spatial Navigation: Latent Behavioral Signatures and the Impact of Synaptic Plasticity Disruption."*
- **P. Peixoto-Moledo** (2025). PhD Thesis, Universitat de Barcelona.
- A.P. Zamani*, **P. Peixoto-Moledo***, D.P. Tomàs, H.G. Rotstein, J. Dalmau, P.E. Jercog. *Anti-NMDAR encephalitis antibodies cause long-lasting degradation of the hippocampal neural representation of memory*. Submitted. https://www.biorxiv.org/content/10.1101/2022.11.25.517901v2 (*co-first author)

- Maudes E, Planagumà J, … **Peixoto-Moledo P**, … Dalmau J. *Neuro-immunobiology and treatment assessment in a mouse model of anti-NMDAR encephalitis.* **Brain.** 2025 Jun 3;148(6):2023–2037.
- **P. Peixoto-Moledo**, A.P. Zamani, J. Dalmau, Y. Ahmadian, P.E. Jercog. *Behavioral impairments in short- and long-term spatial memory in an amnesic animal model of inducible anti-NMDA encephalitis.* In preparation.
- **Peixoto-Moledo, P.**, Basar, S., Dalmau, J., Ahmadian, Y., & Jercog, P. (2023). *Synaptopathy auto-antibodies reveal how long-term memories are formed, recalled, and used during behavior*. IBRO Neuroscience Reports, 15, S794.
- Chaari M, Gaztelumendi N, … **Peixoto-Moledo P**, … Núñez R. *Fluorescent BODIPY-Anionic Boron Cluster Conjugates as Potential Agents for Cell Tracking*. **Bioconjugate Chem.** 2018, 29, 1763–1773.
- Additional posters: FENS Forum (2022); Sociedad Española de Neurociencia (2021); Barcelona Computational, Cognitive and Systems Neuroscience (2021).

Award: Best Poster Presentation, Barcelona Computational, Cognitive and Systems Neuroscience Conference (2021).

TECHNICAL SKILLS

Modeling & inference: Simulation-based inference (SNPE), normalizing flows, latent-variable and state-space models (GLM-HMM), Bayesian inference, causal intervention design, Shapley attribution, posterior predictive checks, model selection and cross-validation.

Programming: Python (NumPy, Pandas, SciPy, scikit-learn), PyTorch, TensorFlow. GPU training and inference; HPC cluster job submission; incremental/resumable pipelines with checkpointing; reproducible caching.

Data: High-dimensional neural and behavioral data; in vivo calcium imaging; fiber photometry; pose estimation (DeepLabCut); cross-session registration; large-scale HDF5 handling.

Communication & infrastructure: Self-contained interactive HTML research reports; scripted figure pipelines; manuscript and methods writing; mentoring.

In vivo & experimental: Experimental design (treatment groups, controls, timepoints, statistical design); chronic one-photon calcium imaging acquisition in freely behaving mice; high-throughput rodent behavioral paradigms; pharmacological manipulation; stereotaxic surgery (GRIN lens implantation, baseplating). Spanish animal experimentation accreditation, función d) — design of procedures and projects (FELASA C equivalent).

Hardware: Arduino-based experimental control, sensor/actuator integration, hardware-triggered synchronization.

ADDITIONAL

MODEP — Took part as one of three team members in MODEP, a proposed AI platform to decode and predict parkinsonian motor states from continuous movement data for therapy optimization, presented at ACTIVA-T, the CSIC programme for knowledge-based companies (Vigo, May 2026).

Co-founder, WoMars (womars.org) — Geneva-based non-profit organizing all-female space analogue missions and training. Responsible for mission logistics, team coordination, partnerships, and fundraising.

Advanced training: EMBO Practical Course — Advanced Modeling of Behavior (2022, Barcelona).

Languages: Spanish, Galician (native); English, Catalan (fluent); basic Italian, French, Portuguese.