# Thomas Heap

thomas.heap@bristol.ac.uk | github.com/ThomasHeap | Brussels, Belgium

## Technical Skills

**Tools:** Python, PyTorch, JAX, NumPy, SciPy, distributed training, AWS, Git, Linux

## Education

**University of Bristol** Bristol, UK
PhD Computer Science *March 2022 – September 2026*

- Self-directed research on generative modelling, scalable approximate inference, and evaluation and interpretability of deep learning models.
- Independently trained protein language models from scratch on amino acid sequence data using a modified nanoGPT setup, and am currently building a sequential Monte Carlo pipeline to guide a protein generator toward binders for a specific target, applying probabilistic inference methods from my core research to biological sequence design.
- Developed massively parallel importance-weighting algorithms for Bayesian inference and co-built an open-source probabilistic programming framework (Alan PPL).
- 5 publications at top ML venues including ICLR 2026 and a NeurIPS 2025 Workshop Spotlight.

**University of Edinburgh** Edinburgh, UK
MSc Artificial Intelligence – Distinction *September 2018 – July 2019*

- Thesis: *Inferring the Parameters of a Systems Biology Model with Sequential Neural Likelihood* — applied simulation-based inference and neural density estimation to infer the parameters of a biological dynamical system from sparse, noisy experimental data, working directly with biology domain experts. Rated outstanding.

**Warwick University** Coventry, UK
BSc Mathematics – 2.1 *August 2014 – July 2017*

## Professional Experience

**Applied Science Intern** *Amazon AGI-F*
*2025–2026*

- Designed and validated entropy-based proxy metrics for evaluating 7B parameter language models, testing correlation with downstream benchmark performance using partial correlations and confound analysis.
- Built data processing and evaluation pipelines on AWS, handling dataset curation, transformation, and systematic evaluation at scale on 64+ GPUs.
- Communicated findings to cross-functional research, engineering, and product teams; work is being prepared for publication.

**Research Engineer III** *Smarsh Inc.*
*2021–2022*

- Maintained and improved production speech recognition models in PyTorch; improved transcription accuracy by 10% via data augmentation for low-resource languages.
- Built a speech diarisation model and pipeline end-to-end, from research prototype through to production deployment.

**Data Scientist (Contract)** *Warner Media*
*2019–2021*

- Built data processing pipelines enriching a video database with NLP-derived metadata at scale.
- Deployed time-series forecasting models into production dashboards used directly by business stakeholders.

## Selected Publications

1. **Thomas Heap**, Tim Lawson, Lucy Farnik, Laurence Aitchison. Automated Interpretability Metrics Do Not Distinguish Trained and Random Transformers. *ICLR 2026.*
2. **Thomas Heap**, Laurence Aitchison, Emma Cahill, Adriana Casado Rodriguez. Rodent-Bench. *NeurIPS 2025 AI for Science Workshop (Spotlight).*
3. **Thomas Heap**, Sam Bowyer, Laurence Aitchison. Massively Parallel Expectation Maximization For Approximate Posteriors. *AABI 2025.*
4. Sam Bowyer, **Thomas Heap**, Laurence Aitchison. Using Autodiff to Estimate Posterior Moments, Marginals and Samples. *UAI 2024.*
5. **Thomas Heap**, Gavin Leech, Laurence Aitchison. Massively Parallel Reweighted Wake-Sleep. *UAI 2023.*