Laurence Aitchison — Curriculum Vitae

Education and Employment

Lecturer (US Assistant Professor): University of Bristol 2019–
Machine Learning

Scientific Consultant: Srini Turaga (Janelia Research Campus) 2018–2019
Computational Neuroscience and Machine Learning

Postdoc: Máté Lengyel (University of Cambridge) 2016–2018
Computational Neuroscience and Machine Learning

Management Consultancy: Oliver Wyman 2015-2016
Consultant

PhD: Gatsby Computational Neuroscience Unit (UCL) 2011–2015
Computational Neuroscience and Machine Learning

MSci: University of Cambridge 2010–2011
Systems Biology **(First)**

BA: University of Cambridge 2007–2010
Experimental and Theoretical Physics **(Double First)**

Recent preprints and workshop papers

My recent work is on various aspects of LLMs, including developing a self-improving coding agent, LLM interpretability with sparse autoencoders, understanding pretraining dynamics and hyperparameter choices, evals, and Bayesian finetuning.

Robeyns M, Szummer M, and **Aitchison L**, "A self-improving coding agent," *ICLR Workshop on Scaling Self-Improving Foundation Models without Human Supervision*, **(Oral)** 2025

Farnik L, Lawson T, Houghton C, and **Aitchison L**, "Jacobian sparse autoencoders: Sparsify computations, not just activations," *arXiv:2502.18147*, (ICML 2025 Submission)

Milsom E, Anson B, and **Aitchison L**, "Function-space learning rates," *arXiv:2502.17405*, (ICML 2025 Submission)

Bowyer S, **Aitchison L**, and Ivanova DR, "Position: Don't use the CLT in LLM evals with fewer than a few hundred datapoints," *arXiv:2503.01747*, (ICML 2025 Submission)

Heap T, Lawson T, Farnik L, and **Aitchison L**, "Sparse autoencoders can interpret randomly initialized transformers," *arXiv:2501.17727*, (ICML 2025 Submission)

Wang X and **Aitchison L**, "How to set AdamW's weight decay as you scale model and dataset size," *arXiv:2405.13698*, (ICML 2025 Submission)

Leech G, Vazquez JJ, Kupper N, Yagudin M, and **Aitchison L**, "Questionable practices in machine learning," *arXiv:2407.12220*, 2024

Yang AX, Robeyns M, Coste T, Wang J, Bou-Ammar H, and **Aitchison L**, "Bayesian reward models for LLM alignment," *Workshop on Secure and Trustworthy Large Language Models*, 2024

Wang X and **Aitchison L**, "Batch size invariant Adam," *arXiv:2402.18824*, 2024

Wang X, **Aitchison L**, and Rudolph M, "LoRA ensembles for large language model fine-tuning," *NeurIPS Workshop: Efficient Natural Language and Speech Processing*, 2023

First/last author publications

Lawson T, Farnik L, Houghton C, and **Aitchison L**, "Residual stream analysis with multi-layer SAEs," *ICLR*, 2025

Milsom E, Anson B, and **Aitchison L**, "Stochastic kernel regularisation improves generalisation in deep kernel machines," *NeurIPS*, 2024

Comi M, Tonioni A, Yang M, Tremblay J, Blukis V, Lin Y, Lepora NF, and **Aitchison L**, "Snap-it, tap-it, splat-it: Tactile-informed 3D Gaussian splatting for reconstructing challenging surfaces," *3DV*, 2024

Garibbo M, Ludwig CJH, Lepora NF, and **Aitchison L**, "Relating human error–based learning to modern deep RL algorithms," *Neural Computation*, 2024

Yang AX, Robeyns M, Wang X, and **Aitchison L**, "Bayesian low-rank adaptation for large language models," *ICML*, 2024

Milsom E, Anson B, and **Aitchison L**, "Convolutional deep kernel machines," *ICML*, 2024

Malkin J, O'Donnell C, Houghton C, and **Aitchison L**, "Signatures of Bayesian inference emerge from energy efficient synapses," *ELife*, 2024

Bowyer S, Heap T, and **Aitchison L**, "Using autodiff to estimate posterior moments, marginals and samples," *UAI*, 2024

Aitchison L and Ganev S, "InfoNCE is variational inference in a recognition parameterised model," *TMLR*, 2024

Garibbo M, Robeyns M, and **Aitchison L**, "Taylor TD-learning," *NeurIPS*, 2023

Yang XY, Robeyns M, Milsom E, Anson, B, Schoots N, and **Aitchison L**, "A theory of representation learning gives a deep generalisation of kernel methods," *ICML*, 2023

Ganev, S and **Aitchison L**, "Semi-supervised learning with a principled likelihood from a generative model of data curation," *ICLR*, 2023

Wang X and **Aitchison L**, "Robustness to corruption in pre-trained Bayesian neural networks," *ICLR*, 2023

Heap T, Leech G, and **Aitchison L**, "Massively parallel reweighted wake-sleep," *UAI*, 2023

Ober S, Anson B, Milsom E, and **Aitchison L**, "An improved variational approximate posterior for the deep Wishart process," *UAI*, 2023

- 25 minute oral presentation.

Leech, G, Rogers-Smith, C, Sandbrink, JB, Snodin, B, Zinkov, R, Rader, B, Brownstein, JS, Gal, Y, Bhatt, S, Sharma, M, Mindermann, S, Brauner, JM, and **Aitchison, L**, "Mass mask-wearing notably reduces COVID-19 transmission," *PNAS*, 2022

Fortuin V, Garriga-Alonso A, Wenzel F, Rätsch G, Turner R, van der Wilk M, and **Aitchison L**, "Bayesian neural network priors revisited," *ICLR*, 2022

Nabarro S, Ganev S, Garriga-Alonso A, Fortuin V, van der Wilk M, and **Aitchison L**, "Data augmentation in Bayesian neural networks and the cold posterior effect," *UAI*, 2022

Ober SW and **Aitchison L**, "A variational approximate posterior for the deep Wishart process," *NeurIPS*, 2021

Aitchison L, Yang AX, and Ober SW, "Deep kernel processes," *ICML*, 2021

Ober SW and **Aitchison L**, "Global inducing point variational posteriors for Bayesian neural networks and deep Gaussian processes," *ICML*, 2021

Aitchison L, "A statistical theory of cold posteriors in deep neural networks," *ICLR*, 2021

Aitchison L, Jegminat J, Menendez J, Pfister JP, Pouget A, and Latham PE, "Synaptic plasticity as Bayesian inference," *Nature Neuroscience*, 2021

Unlu A and **Aitchison L**, "Gradient regularisation as approximate variational inference," *Entropy*, 2021

Fortuin V, Garriga-Alonso A, van der Wilk M, and **Aitchison L**, "BNNpriors: A library for Bayesian neural network inference with different prior distributions," *Software Impacts*, 2021

Aitchison L, "Bayesian filtering unifies adaptive and non-adaptive neural network optimization methods," *NeurIPS*, 2020

Aitchison L, "Why bigger is not always better: on finite and infinite neural networks," *ICML*, 2020

Aitchison L, "Tensor Monte Carlo: particle methods for the GPU era," *NeurIPS*, 2019

- Built into Pyro (a toolbox for probabilistic programming) by Uber (link)

Garriga-Alonso A, Rasmussen CE, and **Aitchison L**, "Deep convolutional networks as shallow Gaussian processes," *ICLR*, 2019

- Foundational work in the field of infinite neural networks, with $>$300 citations.

Aitchison L, Russell L, Packer A, Yan J, Castonguay P, Häusser M, and Turaga SC, "Model-based Bayesian inference of neural activity and connectivity from all-optical interrogation of a neural circuit," *NeurIPS*, 2017

Aitchison L and Lengyel M, "With or without you: predictive coding and Bayesian inference in the brain," *Current Opinion in Neurobiology*, 2017

Bang D, **Aitchison L**, Moran R, Castanon SH, Rafiee B, Mahmoodi A, Lau JYF, Latham PE, Bahrami B, and Summerfield C, "Confidence matching in group decision-making," *Nature Human Behaviour*, 2017

Aitchison L and Lengyel M, "The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics," *PLoS Computational Biology*, 2016

Aitchison L, Corradi N, and Latham PE, "Zipf's law arises naturally when there are underlying, unobserved variables," *PLoS Computational Biology*, 2016

Aitchison L, Bang D, Bahrami B, and Latham PE, "Doubly Bayesian analysis of confidence in perceptual decision-making," *PLoS Computational Biology*, 2015

Mid author publications

Antonio B, McRae ATT, MacLeod D, Cooper FC, Marsham J, **Aitchison L**, Palmer TN, and Watson PAG, "Postprocessing east african rainfall forecasts using a generative machine learning model," *Journal of Advances in Modeling Earth Systems*, 2025

Papamarkou T, Skoularidou M, Palla K, **Aitchison L**, and et al., "Position paper: Bayesian deep learning in the age of large-scale AI," *ICML*, 2024

Shi Z, Yang AX, Wu B, **Aitchison L**, Yilmaz E, and Lipani A, "Instruction tuning with loss over instructions," *NeurIPS*, 2024

Mauro C, Yijiong L, Church A, Alessio T, **Aitchison L**, and Lepora NF, "TouchSDF: A DeepSDF approach for 3D shape reconstruction using vision-based tactile sensing," *IEEE Robotics and Automation Letters*, 2024

Vosper E, Watson P, Harris L, McRae A, Santos-Rodriguez R, **Aitchison L**, and Mitchell D, "Deep learning for downscaling tropical cyclone rainfall to hazard-relevant spatial scales," *Journal of Geophysical Research: Atmospheres*, 2023

Gupta AK, **Aitchison L**, and Lepora NF, "Tactile image-to-image disentanglement of contact geometry from motion-induced shear," *CoRL*, 2021

Sharma M, Mindermann S, Rogers-Smith C, Leech G, Snodin B, Ahuja J, Sandbrink JB, Monrad JT, Altman G, Dhaliwal G, Finnveden L, Norman AJ, Oehm SB, Sandkühler JF, **Aitchison L**, Gavenciak T, Mellan T, Kulveit J, Chindelevitch L, Flaxman S, Gal Y, Mishra S, Bhatt S, and Brauner JM, "Understanding the effectiveness of government interventions against the resurgence of COVID-19 in Europe," *Nature Communications*, 2021

Echeveste R, **Aitchison L**, Hennequin G, and Lengyel M, "Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference," *Nature Neuroscience*, vol. 23, 2020

Schmidt-Hieber C, Toleikyte G, **Aitchison L**, Roth A, Clark BA, Branco T, and Häusser, M, "Active dendritic integration as a mechanism for robust and precise grid cell firing," *Nature Neuroscience*, 2017

Hennequin G, **Aitchison L**, and Lengyel M, "Fast sampling-based inference in balanced neuronal networks," *NeurIPS*, 2014

Other preprints and workshop papers

Aitchison L, "Why you don't overfit, and don't need Bayes if you only train for one epoch," *arXiv:2411.14478*, 2024

He Z, Lin W, Zheng H, Zhang F, Jones M, **Aitchison L**, Xu X, Liu M, Kristensson PO, and Shen J, "Human-inspired perspectives: A survey on AI long-term memory," *arXiv:2411.00489*, 2024

Yang AX, **Aitchison L**, and Moss, HB, "MONGOOSE: Path-wise smooth Bayesian optimisation via meta-learning," *ICML 2024 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2024

Addison H, Kendon E, Ravuri S, **Aitchison L**, and Watson P, "Machine learning emulation of a local-scale UK climate model," *NeurIPS 2022 Workshop: Tackling Climate Change with Machine Learning*, 2022

Hanslope J and **Aitchison L**, "Imitating careful experts to avoid catastrophic events," *NeurIPS 2022 5th Robot Learning Workshop: Trustworthy Robotics*, 2022

Benzing F, Schug S, Meier R, von Oswald J, Akram Y, Zucchet N, **Aitchison L**, and Steger A, "Random initialisations performing above chance and how to find them," *NeurIPS Workshop: Optimization for Machine Learning*, 2022

Robeyns M, Fotopolou S, Walmsley M, and **L Aitchison**, "Fast estimation of physical galaxy properties using simulation-based inference," *ICML Workshop: Machine Learning for Astrophysics*, 2022

Wang, X and **Aitchison L**, "Undefined class-label detection vs out-of-distribution detection," *AABI*, 2022

Aitchison L, Hennequin G, and Lengyel M, "Sampling-based probabilistic inference emerges from learning in neural circuits with a cost on reliability," *arXiv: 1807.08952*, 2018

Aitchison L, Adam V, and Turaga SC, "Discrete flow posteriors for variational inference in discrete dynamical systems," *arXiv: 1805.10958*, 2018

Successful funding

$32,000 Bayesian modelling of LLM capabilities from evals (2025)

£1,070,000 CoI: NERC - Future Rainfall and Flood Extremes (2024)

£27,000 CoI: NERC (internal) - ML-based African weather forecasting (2023)

£100,000 PI: Bid for a private donation for GPUs for Bristol's Computational Neuroscience Unit (2021)

£100,000 PI: Bid for internal equipment funding for GPUs for the School (2021)

£53,500 CoI: David Telling One Year Fellowship Proposal - Longitudinal Variability and Progression of Sleep Disturbance in Mild Cognitive Impairment (2020)

£20,000 PI: Royal Society Equipment grant (2020)

Public impact

- Live interviews on BBC News Channel (14/07/2021) and BBC Radio Bristol (07/07/2021). Quoted in ∼175 publications [data from Bristol's Press Office] around the COVID pandemic, including: The New York Times, i news, Sky News Live [scroll down to 14:44], Wired, The Herald, Bristol Live, My London.
- Advised HM Government Cabinet Office, COVID-19 Taskforce on their Winter 2021 "Plan B" non-pharmaceutical policies, based on Leech et al. 2021 (I was senior author; 23/9/21)

Invited Talks

Functional Inference and Machine Intelligence Workshop (University of Bristol) "Deep kernel processes and machines" (2024)

Bristol statistics seminars "Massively parallel probabilistic inference" (2024)

Oxford Computational Statistics and Machine Learning (OxCSML) seminar. "Deep kernel processes and machines" (2024)

Deep Learning Classics and Trends "(Convolutional) Deep Kernel Machines" (2023)

Workshop on generative models and uncertainty quantification (GenU) "Deep kernel processes" (2023)

MILA "Deep kernel processes" (2023)

Gatsby Computational Neuroscience Unit "Deep kernel processes" (2023)

Bayesian World Meeting (ISBA) "Deep kernel processes" (2022; cancelled due to travel issues).

Finnish Center for Artificial Intelligence (A flagship of the Academy of Finland combining many research groups from the University of Helsinki, Aalto University, and VTT Technical Research Centre of Finland.) "Deep kernel machines" (2022)

Max Planck Institute for Biological Cybernetics "A statistical theory of cold posteriors, semi-supervised learning and OOD detection" (2021)

University of Amsterdam "A statistical theory of cold posteriors, semi-supervised learning and OOD detection" (2021)

UCL; CSML "Deep kernel processes" (2021)

SecondMind "Deep kernel processes" (2021)

Bernstein Workshop: Neural sampling "Variability + learning = sampling-based probabilistic inference in the brain" (2017)

EITN Workshop: Learning and Synaptic Plasticity "Probabilistic synapses" (2015)

Bernstein "Sparks" Workshop: Neural models of perceptual decision making in natural environments: from theory to experiment "Probabilistic synapses" (2015)

Posters

Yan J, Kerlin A, **Aitchison L**, Mohar B, Svoboda K, and Turaga SC, "Deep Dendrite: Bayesian inference of synaptic inputs from dendritic calcium imaging," *COSYNE*, 2019

Aitchison L, Hennequin G, and Lengyel M, "Probabilistic inference emerges from learning in neural circuits with a cost on reliability," *COSYNE*, 2018

Aitchison L, Castonguay P, Yan J, Packer A, Russell L, Häusser M, and Turaga SC, "All-optical mapping of local circuit connectivity *in vivo*," *COSYNE*, 2017

Orbán G, **Aitchison L**, and Lengyel M, "Efficient probabilistic inference with oscillatory, excitatory-inhibitory neural circuit dynamics," *COSYNE*, 2015

Aitchison L, Corradi N, and Latham PE, "Latent variable models (almost) always display signatures of criticality," *COSYNE*, 2014

Lengyel M, Hennequin G, and **Aitchison L**, "Fast sampling in recurrent neural circuits," *COSYNE*, 2014

Aitchison L and Latham PE, "The synaptic sampling hypothesis," *COSYNE*, 2013

Orbán G, **Aitchison L**, and Lengyel M, "Hamiltonian Monte Carlo sampling and oscillatory activity in V1," *COSYNE*, 2012

Reviews

Reviewed extensively in machine learning: NeurIPS, ICLR, ICML, AISTATS, AABI, and in computational neuroscience: eLife, Nature Neuroscience, Nature Communications, PLoS Computational Biology, Neural Computation.

Programming and Technical Skills

Deep Learning frameworks: PyTorch, Jax, TensorFlow.

Programming Languages: Python, Julia, MATLAB, R, ML, C.

Operating Systems: Linux, MacOS, Windows.

Typesetting: LaTeX.

Teaching Experience

I am a Fellow of the HEA.

Course Director (2023-25) Introduction to Artificial Intelligence
University of Bristol, Computer Science, compulsory second year module on deep learning with 150 students.

Course Director (2020-23) Data Driven Computer Science
University of Bristol, Computer Science, compulsory second year module with 250 students.

Course Director (2020-23) Introduction to Computational Neuroscience
University of Bristol, Computer Science, optional third year module with around 150 students from many different backgrounds.

Graduated PhD students

- Adam Yang (2025; viva forthcoming, Research Scientist, Mistral AI)
- Gavin Leech (2024; Founder, Arb Research)
- Michele Garribo (2024; Deep Learning Scientist, Centre for Genomic Regulation in Barcelona)

PhD examinations

- Will Greedy (Internal; 2024)
- Yingxuan Yi (Internal; 2024)
- Frederik Benzing (ETH Zurich; 2022)
- Will Price (Internal; 2021)
- Yu Chen (Internal; 2021)
- Tim Pearce (University of Cambridge; 2020)

Grant review

- Israel Science Foundation Research Grant (Research Grant; 2024)
- BBSRC (Supporting research in cognitive computational neuroscience; 2022)
- Dutch Research Council (Veni research proposal; 2021)

Selected Events

- Dagstuhl Seminar 24461: Rethinking the Role of Bayesianism in the Age of Modern AI (2024)

Selected Professional Memberships

- ELLIS Member

Consultancy

Oliver Wyman: Co-authored the Expert Advisory Panel's response to the Treasury's Financial Advice Markets Review.

Oliver Wyman: Supported project management for a large IT transformation programme in a major UK Bank.