

CIAN EASTWOOD

London, United Kingdom

✉ cian.eastwood@gmail.com | 🏠 Personal website | 🐙 Github | 🔗 LinkedIn | 🎓 Google scholar

Summary

A machine-learning researcher whose PhD focused on preparing models for new contexts, with proven expertise in robustness, adaptation, generalization, fine-tuning, and self-supervised learning.

Education

PhD in Machine Learning

University of Edinburgh & Max Planck Institute for Intelligent Systems, Tübingen

Sept 2018 – Jan 2024

UK & Germany

- Supervisors: Prof. Chris Williams and Prof. Bernhard Schölkopf.
- Thesis: “Shift happens: How can we best prepare machine learning systems?”
- Research: Robustness, out-of-distribution generalization, domain adaptation, fine-tuning, in-context learning, self-supervised learning, causality, representation learning and disentanglement.
- Awards: Enlightenment Scholarship, NUI Travelling Doctoral Studentship.

MSc in Artificial Intelligence – Distinction (84%)

University of Edinburgh

Sept 2016 – Sept 2017

UK

- Thesis: “Experiments with information-maximising generative adversarial networks”
- Courses: Various courses in AI, e.g. probabilistic modelling and reasoning, pattern recognition and natural language processing.
- Awards: **Best Thesis**, The Informatics Scholarship, UK/EU Masters Scholarship.

Exchange Program

University of Toronto

Aug 2014 – Dec 2014

Canada

- Courses: Programming languages, computer networks, compilers, and formal methods of software design.

BSc in Computer Science – First-Class Honours (88%)

National University of Ireland (NUI), Maynooth

Sept 2012 – June 2016

Ireland

- Thesis: Minimising Volatility, Maximising Diversification.
- Courses: Broad range of courses in computer science and mathematics.
- Awards: Intel Medal—**graduated top of class**, STEM Scholarship, Entrance Scholarship.

Professional Experience

Student Researcher (PhD)

Google DeepMind

Sept – Dec 2023

London, UK

- Large-scale generative transformers for continuous data (see [1]).

Research Scientist Intern

Spotify

Jun 2023 – Aug 2023

London, UK

- Learning representations of high-dimensional treatments (e.g., playlists).
- Part of the Causal Inference Lab.

Research Intern (AI)

Meta

Aug – Dec 2022; Jan – April 2023

New York, USA; London, UK

- Self-supervised representation learning for improved downstream performance (see [2]).
- Part of FAIR Labs.

Research Assistant

University of Edinburgh

Nov 2017 – Sept 2018

Edinburgh, UK

- Deep generative models for human motion synthesis.
- Collaborative research environment, with Prof. Taku Komura.

Intern Analyst

Accenture

Feb 2015 – Aug 2015

Dublin, Ireland

- Large-scale professional software development.

Publications

- [1] “GIVT: Generative Infinite-Vocabulary Transformers”
M Tschannen, **C Eastwood**, F Mentzer
Preprint, Under Review (2023)
- [2] “Self-Supervised Disentanglement by Leveraging Structure in Data Augmentations”
C Eastwood, J von Kügelgen, L Ericsson, D Bouchacourt, P Vincent, B Schölkopf, M Ibrahim
NeurIPS 2023 Workshops on Self-Supervised Learning and Causal Representation Learning
- [3] “Spuriousity Didn’t Kill the Classifier: Using Invariant Predictions to Harness Spurious Features”
C Eastwood*, S Singh*, A Nicolicioiu, M Vlastelica, J von Kügelgen, B Schölkopf
37th Conference on Neural Information Processing Systems (NeurIPS 2023)
- [4] “DCI-ES: An Extended Disentanglement Framework with Connections to Identifiability”
C Eastwood*, A Nicolicioiu*, J von Kügelgen*, A Kekić, F Träuble, A Dittadi, B Schölkopf
11th International Conference on Learning Representations (ICLR 2023)
- [5] “Probable Domain Generalization via Quantile Risk Minimization”
C Eastwood*, A Robey*, S Singh, J von Kügelgen, H Hassani, G J Pappas, B Schölkopf
36th Conference on Neural Information Processing Systems (NeurIPS 2022)
- [6] “Align-Deform-Subtract: An Interventional Framework for Explaining Object Differences”
C Eastwood*, N Li*, C K I Williams
ICLR 2022 Workshop on Objects, Structure and Causality
- [7] “On the DCI Framework for Evaluating Disentangled Representations: Extensions and Connections to Identifiability”
C Eastwood*, A Nicolicioiu*, J von Kügelgen*, A Kekić, F Träuble, A Dittadi, B Schölkopf
UAI 2022 Workshop on Causal Representation Learning
- [8] “Source-Free Adaptation to Measurement Shift via Bottom-Up Feature Restoration”
C Eastwood*, I Mason*, C K I Williams, B Schölkopf
10th International Conference on Learning Representations (ICLR 2022, **Spotlight**)
- [9] “Unit-Level Surprise in Neural Networks”
C Eastwood*, I Mason*, C K I Williams
NeurIPS 2021 Workshop “I Can’t Believe it’s Not Better” and PMLR 163:33-40 (**Spotlight & Didactic Award**)
- [10] “Learning Object-Centric Representations of Multi-Object Scenes from Multiple Views”
N Li*, **C Eastwood***, R Fisher
34th Conference on Neural Information Processing Systems (NeurIPS 2020, **Spotlight**)
- [11] “A Framework for the Quantitative Evaluation of Disentangled Representations”
C Eastwood*, C K I Williams
6th International Conference on Learning Representations (ICLR 2018)

Awards

- 2022 NeurIPS Top Reviewer
- 2022 ICLR Highlighted Reviewer
- 2019 NUI Travelling Doctoral Studentship in Artificial Intelligence
- 2018 University of Edinburgh Enlightenment PhD Scholarship
- 2017 Informatics Dissertation Prize (Award for best thesis in the MSc Artificial Intelligence)
- 2016 Informatics Masters Scholarship
- 2016 UK/EU Masters Scholarship
- 2016 The Intel Medal (Award for best results in the BSc Computer Science)
- 2012 NUI Undergraduate STEM Scholarship
- 2012 NUI Undergraduate Entrance Scholarship
- 2012 600 points in The Leaving Certificate (Final secondary-school exams, 99.7th percentile nationally)

Invited Talks

- [A] “Distribution shift and causal/disentangled representations”
Computational Intelligence, Vision, and Robotics Lab, New York University, December 2022
- [B] “Probable domain generalization via quantile risk minimization”
Copenhagen Causality Lab, University of Copenhagen, November 2022 (virtual)
- [C] “Shift happens: How can we best prepare?”
(Neuro)Science of Deep Learning Group, Massachusetts Institute of Technology, November 2022 (virtual)
- [D] “Tackling distribution shift and out-of-distribution generalization”
Seminar on Out-of-Distribution Generalization, Saarland University, November 2022 (virtual)

Community Service/Engagement

- 2023 Reviewer: ICLR, ICML; Organizing Committee: Conference on Causal Learning and Reasoning (CLeaR); Co-Organizer: NeurIPS Causal Representation Learning workshop.
- 2022 Reviewer: NeurIPS, ICLR.
- 2021 Reviewer: NeurIPS.

Skills

- Programming Python (PyTorch, JAX, Pandas, NumPy, Scikit-learn. etc.)
- Miscellaneous Linux, Shell (Bash/Zsh), \LaTeX , Git

References available upon request.