Luke N. Darlow, PhD.

SENIOR MACHINE LEARNING RESEARCHER & ENGINEER

+44 (0)738 4272869 lukedarlow@yahoo.com London, UK
uk.linkedin.com/in/luke-darlow-6457a0150
https://scholar.google.com/citations?user=qxgq8n8AAAAJ

PERSONAL SUMMARY

I have a passion for navigating difficult technical problems using **creative machine learning solutions**, and developing ways of expressing my solutions and findings in easy-to-understand, intuitive, and **demonstrable** ways. I have worked as a ML Research Scientist and ML Engineer, with expert level experience across a broad variety of domains and using many modalities. After obtaining a CS and Physics degree I applied **Computer Vision** and ML to fingerprint-related problems and **published 10 papers** during my first CS MSc. I then obtained my second MScR. in **Data Science** and a CS PhD from the University of Edinburgh, focussing on representation learning in complex unsupervised settings where **neural networks** fail to generalise. During my PhD I took part in the 3 minute thesis competition, where I learned to excel at communicating my technical research to a lay audience, was a research and teachers assistant, and co-supervised a MSc. student who extended the work in my first MSc. I worked in speech synthesis at Altered, where I became an expert at **transformer** architectures before joining Huawei R&D in Edinburgh as a Senior ML Engineer and ML Researcher. Additionally, I act as the **ML Team** Lead for Systems Infrastructure Research, where I mentor junior colleagues, provide advice and direction to my line manager regarding our broader ML work and growing expertise, conduct interviews, and supervise interns. I am an **expert at time series forecasting**, and have developed a foundational model in this space, resulting in an upcoming publication at **ICLR2024**. I have been writing software for almost 20 years, and always leverage my technical prowess to first understand new problems from low- to high-level before designing and implementing solutions that **add value** to my team, the company, and the ML community.

KEY SKILLS

Model Design	Deep Learning	Transformers	Time Series Forecasting	Computer Vision
Information Theory	Audio Modelling	Generative Modelling	Representation Learning	Data/Model Visualisation
Supervision and Leadership	Creativity	Big Picture Understanding	Problem Scoping	Technical Communication
PyTorch	Tensorflow	SKLearn	Git	Software Development

PROFESSIONAL EXPERIENCE

Senior Engineer & ML Team Lead August 2022 – Present
Huawei R&D, Systems Infrastructure Research Edinburgh, UK

Systems Infrastructure Research is a team led by Prof. Adam Barker with the mandate to improve Huawei's cloud resource management

MODEL DESIGN IMPLEMENTATION ANALYSIS

- Designed a SoTA **foundational forecasting model,** resulting in a paper to be presented at ICLR2024
  *Forecasting time series is under-explored for foundational models, and my paper is the **first on this topic** to be accepted at a high-level conference.*
- Created and trained models to **forecast high-resolution** in-house cloud data & published at EuroSys2023
  *The model I designed (**FoldFormer**) is currently **deployed** and forecasting cloud resource demand for Huawei's backend infrastructure, and has been submitted for patent.*
- Data **analytics and model testing** for in-house datasets; published at ACM-SoCC 2023
  *I added value to this paper by applying several forecasting models, and encouraged the lead author (a junior colleague) to write this paper as his **first published work**.*

SUPERVISION MENTORSHIP LEADERSHIP

- Supervised interns and junior colleagues to ensure their **consistent growth** in core ML skills and research
  *I have mentored 3 interns, where my goal has been to grow their ML engineering, research, and technical writing skills, while applying them to **real-world** problems.*
- Responsible for **interviewing** interns and contractors
- Provided **project-specific advice** for data analytics and modelling based on broad experience
  *I consistently sit in meetings to ensure any **data-related work** is well-justified and well-structured, and where any ML modelling approaches are **achievable** and interesting.*
- Proposed and **led new projects** and directions for the wider team
  *Since joining Huawei I have proposed and/or or led 4 projects, all of which have yielded published papers and/or **deployment**. I proposed '**Universal Forecasting**' as a fundamental aspect of our current 18 month technical charter, and have since made significant progress thereon. I have also proposed and led several intern projects.*

Deep Learning Research Scientist April 2022 – July 2022
ALTERED UK

Altered is a speech synthesis company that enables users to alter their voice using curated counterparts via deep learning models

- Implemented and tested SoTA **text-to-speech** and **speech-to-text transformer** models
  *I added **robust code contributions** to their repository, all of which required PyTorch or Tensorflow implementations of existing Transformer-based audio/speech models.*
- Determined via extensive experiments hyper-parameter configurations for **optimising transformer models**
  *This yielded extensive documentation detailing hyper-parameter sensitivity for transformer-based models, and provided **novel insight** into their learning rate schedulers.*
- **Advised** junior colleagues and software development team on **viable directions** for model improvements
  *Using my broad ML skillset, I helped colleagues by explaining **potential anomalous behaviour** of their in-house models, and proposed new directions.*

Teachers Assistant, Tutor & Research Assistant September 2020 – October 2022
University of Edinburgh UK

Worked with various lecturers and professors as a tutor, marker, research assistant & teachers assistant

- Teachers Assistant
  *Assisted lecturers and **managed the Q&A platform** for machine learning practical MSc. level course, providing help and advice to students.*
- Cancer Detection Research Assistant
  *Worked with Prof. Robert Fisher on **Computer Vision models** for cancer detection in blood samples, providing novel insights for biology partners into cancer predictability.*

PROFESSIONAL EXPERIENCE (continued)

Biometrics Researcher December 2015 – August 2017
Council for Scientific and Industrial Research Pretoria, South Africa

CSIR sits between universities and industry as the government's lead research institution and aims to solve a wide range of problems

- Designed and trained a **computer vision model** to **detect features from fingerprints,** for downstream use.

Designed and trained a **computer vision model** to **detect features from fingerprints**, created by aggregating votes from many off-the-shelf deterministic detection algorithms. This work was **published at the top biometrics conference**, IJCB.

- Optical Coherence Tomography for 3D subsurface fingerprint acquisition
  - Designed **algorithms to extract usable subsurface fingerprints** from 3D OCT scans

*CSIR was developing a novel 3D **subsurface fingerprint scanner** using OCT technology. My algorithms were highly efficient and ran in real-time to extract fingerprints from **large volumetric scans** in **real-time**. I also wrote the software (both backend and frontend) to perform these scans, extract fingerprints, and detect minutiae therein. My work was **presented to the president** of South Africa, the **minister of science and technology**, and at various trade shows.*

  - Published **10 papers** (8 of which I was 1st author) on this topic in the span of 2 years while completing my first MSc. by research, including 2 journal papers.

Miscellaneous 2011 – 2019

- Software development

*I developed a health and safety application C-sharp for AFFSAF such that **health and safety professionals** could record, analyse, and track inspection data digitally. This was packaged and distributed to dozens of large companies alongside a best-selling health and safety book.*

- Supervision

*I co-supervised an MSc. student with Dr. Patrick Marais, extending my own work on 3D fingerprint acquisition. His **thesis title** was: "matching 3D fingerprints with a modified toroidal iterative closest point algorithm"*

EDUCATION

PhD. Computer Science, University of Edinburgh July 2022
Supervised by Prof. Amos Storkey as part of the Data Science CDT

***Thesis title**: Learning reliable representations when proxy objectives fail*

MScR. Computer Science (with distinction), University of Edinburgh October 2018

***Thesis title**: Information Theoretic Analysis of Deep Neural Networks. **Coursework**: Machine Learning Practical; Machine Learning and Pattern Recognition; Probabilistic Modelling and Reasoning; Introduction to Research in Data Science; Extreme Computing.*
*Also developed a **web tool** to understand NNs: https://www.bayeswatch.com/assets/ginn/good3.html*

MSc. Computer Science (with distinction), Rhodes University, South Africa December 2015

***Thesis title**: Internal Fingerprint Extraction. **Awards**: best masters studentship (from CSIR) and best student researcher (from Rhodes University)*

BSc. (Hons) Computer Science and Physics, Rhodes University, South Africa December 2013

PUBLICATION HIGHLIGHTS

After PhD

- DAM: Towards a foundation model for time series forecasting, ICLR2024.

The DAM uses a novel mechanism to sample time non-linearly and forecasts by way of representation-parameterised basis function composition.

- How Does It Function? Characterizing Long-term Trends in Production Serverless Workloads, A Joosen, A Hassan, M Asenov, R Singh, L Darlow, J Wang, A Barker. Proceedings of the 2023 ACM Symposium on Cloud Computing, 443-458.

Assisted in time series forecasting using cloud resource data (e.g., CPU utilisation) as a baseline estimate for this data release. https://github.com/sir-lab/data-release

- FoldFormer: sequence folding and seasonal attention for fine-grained long-term FaaS forecasting, LN Darlow, A Joosen, M Asenov, Q Deng, A Barker. EuroMLSys 2023: Proceedings of the 3rd Workshop on Machine Learning and Systems.

Efficient transformer model to forecast for very high-resolution cloud resource data, using 3 novel architectural components

- TSMix: time series data augmentation by mixing sources, LN Darlow, A Joosen, M Asenov, Q Deng, J Wang, A Barker. EuroMLSys 2023: Proceedings of the 3rd Workshop on Machine Learning and Systems.

A technique to improve generalisation of time series models by mixing samples with scale invariance.

During PhD

- Deep Decision Tree Layer for Learning a High-coverage Semantic Hash Function, LN Darlow, AJ Storkey, 2022.

A technique to extend the categorisation capability of a deep classifier, employing contrastive learning on a deep decision tree representation.

- Latent Adversarial Debiasing: Mitigating Collider Bias in Deep Neural Networks, LN Darlow, S Jastrzębski, AJ Storkey, 2020.

Gradient-based method for removing latent features of data that are informative of class but that do not generalise outside of the training set.

- DHOG: Deep Hierarchical Object Grouping, LN Darlow, AJ Storkey, 2020.

A deep clustering method developed to learn novel clusterings by minimising the mutual information between cluster allocations across multiple heads.

- What Information Does a ResNet Compress?, LN Darlow, AJ Storkey, 2019.

Used model-based mutual information to determine whether a ResNet compresses information about input images when learning to classify.

- CINIC-10 is not ImageNet or CIFAR-10, LN Darlow, EJ Crowley, A Antoniou, AJ Storkey, 2018

Dataset constructed by combining downscaled ImageNet images of the same classes from CIFAR-10, released open-source to the community.

Before PhD

- Fingerprint minutiae extraction using deep learning, LN Darlow, B Rosman - IEEE International Joint Conference on Biometrics, 2017.

A CNN trained to emulate the popular vote from several off-the-shelf minutiae detection algorithms, thereby surpassing them owing to an ensemble-like student-teacher setup

- Efficient internal and surface fingerprint extraction and blending using optical coherence tomography, LN Darlow, J Connan - Journal of Applied Optics, 2015.

OCT extracts multiple misaligned fingerprints with different quality in different regions. This work took quality into account in order to realign and combine these.