Antreas Antoniou

AI Research Scientist and Engineer

Informatics Forum
Edinburgh, UK EH89AB
✆ *+44 7516162339*
✉ *iam@antreas.io*
✈ *https://antreas.io/*
 GitHub
 Google Scholar

Education

2017–2021	**PhD in Machine Learning**, *The University of Edinburgh*.
2016–2017	**MScR in Data Science with Distinction**, *The University of Edinburgh*.
2014–2015	**MSc in Data Science with Distinction**, *Lancaster University*.
2011–2014	**BEng (Hons) in Computer Systems Engineering**, *Lancaster University*.

Research Projects

2021-Current	**TALI: Democratizing Multi-Modal Large Scale Machine Learning**, *Spearheaded an open-source quadra-modal dataset to democratize multi-modal machine learning*.
2021-Current	**GATE: Diversifying and Robustifying Representation Learning**, *Led the development of a multi-domain, multi-task, and multi-modal benchmark suite aimed at diversifying empirical evaluations and conclusions*.
2017-2021	**PhD Thesis: Meta-Learning for Few-Shot Learning**, *Conducted in-depth research on both supervised and unsupervised few-shot learning. Contributions include [6, 7, 8]*, Click for document.
2017	**MScR Thesis: Data Augmentation Generative Adversarial Networks**, *Developed a Generative Adversarial Network for data augmentation, resulting in improved generalization for machine learning models [9]*.
2014	**BEng Dissertation: Fault Tolerant, Self Monitoring Sensors**, *Researched a professional-grade sensing system capable of self-validating its own functionality using signal injection techniques and fault prediction*.

Employment

2021-Current	**Research Associate on Data-Efficient, Highly Transferable and Robust Generalization Learning**, *University of Edinburgh*, Gained recognition for contributions exceeding normal job expectations, including research support and team leadership. Rewrote codebases for efficiency and adaptability, significantly facilitating research projects. Acted as a mentor, enhancing the research capabilities of team members., Edinburgh, UK.
2020-2021	**Research Intern on Few-Shot Learning**, *Google*, Focused on enhancing the transferability of Google's few-shot learning systems, particularly under conditions of extreme domain shift., Mountain View, USA.
2017-2020	**Machine Learning Practical Lead Teaching Assistant, Group Tutor, Demonstrator and Piazza Instructor**, *University of Edinburgh*, Provided comprehensive support for the Machine Learning Practical course, including tutoring and course material development., Edinburgh, UK. `https://antreas.io/teaching/`
2016	**Speech-Scientist Intern**, *Amazon*, Worked on the extension and improvement of the capabilities of Amazon Echo., Cambridge, UK.

2015	**Research Associate**, *Lancaster University*, Participated in the Deep Online Cognition project, utilizing a new component-based programming language, Dana, for modular software development., Lancaster, UK.
2014	**Embedded Systems Research Intern**, *Lancaster University*, Designed, built, and programmed new hardware for Blackpool Illuminations, focusing on LED control via PWM and high-voltage frequency control., Lancaster, UK.
2013	**Software Developer Intern**, *Lancaster University*, Developed an Android app to facilitate real-time interaction between presenters and audiences., Lancaster, UK.

Awards and Nominations

2023	Sustained Excellence Contribution Reward as Salary Increment.
2020-2021	Staff Award for being the Machine Learning Practical Teaching Assistant – `https://antreas.io/awards`
2019	Received 5 Teaching Award Nominations: Best Practice in Inclusive Learning, Best Support Staff, and two for Best Student Who Tutors – `https://antreas.io/nominations`
2019	Finalist for the Best UK PhD Tutor Award.
2019	Top-3 Finalist in UK Open Source Awards for my MAML++ Framework.
2018	Nominated for the Best Student Who Tutors Award.
2015	Awarded the IBM Prize for Best Data Science Dissertation.
2014	Received MSc Data Science Scholarship.
2014	Secured 2nd Place in Lancaster University CS Hackathon 2014.

Teaching and Mentorship

2021-2023	**Mentorship and Team Leadership**, *Onboarded new team members and provided weekly teaching sessions on specific university modules (such as MLP) to enrich research capabilities. Conducted regular meetings to discuss research, engineering, and software development problems.*.
2021-2023	**One-on-One Research Guidance**, *Conducted one-on-one sessions to debug machine learning and software issues, unblocking stalled research directions.*.
2022	**Accelerated Deep Learning Fundamentals Course**, *Provided a 6-week accelerated course in deep learning fundamentals for a part-time student lacking the requisite background. Met once a week to ensure skill acquisition.*.
Sept. 2017 to Current	**Machine Learning Practical Course**, *Teaching Assistant, Group Tutor (Effectively Research Supervisor), Demonstrator and Piazza Instructor*, Full Description at `https://antreas.io/teaching/`.
Apr. 2015 to May 2015	**Digital Innovation**, *Teaching Assistant*.

Research Collaboration and Support

2023	**Community Support for EIDF A100 GPU Cluster**, *One of the early adopters and a key community support member. Created the Slack server, answered hundreds of questions, scheduled community meetings for key issues, and served as a bridge between users and developers.*.
2023	**Tooling and Infrastructure**, *Developed a Python package called kubejobs, simplifying Kubernetes job specifications. Authored early documentation for the university's EIDF A100 GPU cluster.*.

2022	**POEM Project**, *Directly supported the project by offering bi-weekly engineering and research support. Orchestrated 40 experiments in just two days to meet the paper submission deadline..*
2021	**HDR UK Medical Deep Learning Project**, *Provided significant contributions to the project by rewriting the codebase to improve quality and efficiency. Achieved a 35x speed-up in the training/evaluation pipelines. Provided weekly perspectives, suggestions, and direct support..*

Grants

2021-2023	Awarded $20K in Google Compute Platform Research Credits.
2021	Received TRC Compute Grant from Google, providing access to multiple Cloud TPUs for one month.
2017-2020	Managed GCP Compute Grant for the MLP Course, optimizing efficient resource utilization.

Programming Languages and Deep Learning Frameworks

Intermediate	C/C++, HTML, LaTeX, ASSEMBLY
Advanced	PYTHON, JAVA
Advanced	PYTORCH, TENSORFLOW, KERAS

Skills

Deep Learning

Development, *Very experienced in designing, implementing, debugging and tuning a large variety of end-to-end differentiable systems, a subset of which include 1) large-scale multi-modal transformers 2) meta-learning systems such as MAML, 3) GANs of all varieties, such as image-conditional GANs used for image translation, super-resolution, in-filling, domain-transfer, 4) classifiers incorporating any of the modern architecture building blocks, 5) adversarial attacks and defences, 6) state of the art machine translation systems utilizing LSTMs and transformers 7) Multi-sample, multi-parameter-set layers (https://github.com/pytorch/pytorch/issues/17983.*

Research, *Experienced deep learning researcher with a focus on large-scale multi-modal learning and meta-learning. I like to draw insights by actively working on different deep learning subfields, and then leveraging my across-task insights on task-specific projects. I have conducted/collaborated research on well over 65 separate projects in different subfields of deep learning. I have supervised 56 student groups working on deep learning projects, 3 of which were finalists (top-2 in the yearly MLP course competition for the IBM prize.*

Cloud Technologies and Large Scale Machine Learning, *I have 2 years of experience using Google cloud in conjunction with Docker and Kubernetes, to train large-scale multi-modal machine learning systems. I have generally been using Google Cloud infrastructure, including GPU machines, managing disks, images and snapshots, building, maintaining and using Kubernetes clusters, and building and maintaining docker images. Furthermore, I have experience with the nuances of getting models to train efficiently on such large-scale settings, both at the hardware and software levels, as well as the time and cost levels..*

Engineering

General Engineering Skills, *Control and Systems Engineering, Engineering Mathematics.*

Electronics Engineering Skills, *Digital Electronics Engineering, Advanced Electronics Theory Knowledge, Signal Processing, Hardware Design, Integrated Circuit Engineering.*

Software Engineering Skills, *Distributed Systems Development: Java RMI, JGroups, P2P, ReST, LoST, ChordNodes, Networks Programming Knowledge and Experience.*

Embedded Systems Engineering, *Experience programming low-level platforms such as Arduino, Raspberry Pi, PIC microcontrollers, ARM-based micro-controllers and Android*.

Operating Systems, *Windows 7, Windows 8, Windows 10, Ubuntu Linux, Mac OS, Unix*.

Languages

English	**Proficient**	
Greek	**Native**	
Japanese	**Basic**	*Passive usage mostly, picked up from watching 550+ Anime shows*

Teaching

2021-2023	**Supervisory Role for MSc Students**, *University of Edinburgh*, Main Supervisory Contact, Acted as the main supervisory contact for 4 MSc students, two of whom successfully completed their thesis with marks between 60-70%, and one of which was of quality and significance suitable for conference submission..
2021-2022	**Accelerated Deep Learning Fundamentals Course**, *University of Edinburgh*, Course Instructor, Designed and delivered a 6-week accelerated deep learning fundamentals course to aid a part-time student lacking the required background, meeting with them once a week..
Sept. 2017 to Current	**Machine Learning Practical Course**, *University of Edinburgh*, Teaching Assistant, Group Tutor (Effectively Research Supervisor), Demonstrator and Piazza Instructor, Full Description at https://antreas.io/teaching/.
Apr. 2015 to May 2015	**Digital Innovation**, *University of Edinburgh*, Teaching Assistant.

Proactive Technical Leadership

Onboarding and Mentorship:	Proactively **onboarded new team members**, ensured **smooth and low-friction remote development** on GPU machines, and provided **weekly teaching** on university modules to enrich research capabilities. Conducted **regular meetings** to discuss research, engineering, and software development challenges.
Research Quality Enhancement:	Conducted **one-on-one sessions** to debug machine learning model training and generalization behaviors and software issues, thereby **unblocking stalled research directions**.
HDR UK Medical Project:	Contributed significantly by **rewriting the project codebase** to increase code quality, adaptability, and efficiency. Achieved a **35x speed up** in the training/evaluation pipelines.
Research Collaborations:	Directly supported multiple projects like **POEM** by providing **bi-weekly engineering and research support**. **Orchestrated 40 experiments** in two days to meet project deadlines.
Infrastructure and Tooling	**Kubernetes Cluster**: Built and deployed a Kubernetes cluster for the research group. Developed Python tooling for cluster management and gave a tutorial to the group on Kubernetes usage. **Research Framework**: Authored a minimal machine learning research framework for the group, following best practices. **Deep Learning Server**: Successfully procured a £50K deep learning research server through comprehensive market analysis, vendor negotiations, and overseeing the setup process.
Community Engagement and Support:	**Early adopter** of the university's EIDF A100 GPU cluster. Authored **early documentation** and a Python package named **kubejobs** to facilitate user engagement. Founded the cluster's **Slack server**, answered hundreds of user questions, and organized **community meetings** for key issue discussions.

Research Orchestration:	Improved research orchestration by **learning how to deploy Kubernetes clusters**, subsequently setting up a cluster for the group. Conducted a **tutorial** on Kubernetes usage and provided **Python tooling** for easier management.
Knowledge Sharing:	**Built a comprehensive wiki page** for the research group to document best practices, tools, and other valuable resources.

Open Source Philosophy and Contributions

Philosophy: My conviction in the *democratization of Machine Learning research* stems from two core beliefs: the irreplaceable value of *individual expertise* and the *power of collective collaboration*. For me, open source is not merely a development model; it is a **fundamental necessity** for driving innovation and maintaining ethical standards in the field.

Contributions: I have been an active participant in the open-source ecosystem, contributing to various projects, educational resources, and frameworks. My contributions are publicly available on my GitHub profile and include:

- **How to Train Your MAML**: A user-friendly framework that simplifies the complexities of MAML meta-learning.

- **Minimal-ML-Template**: A scalable and adaptable machine learning template designed to minimize the initial setup overhead for researchers.

- **Kubejobs**: A Python package that simplifies working with Kubernetes, facilitating an easier and more efficient experience for users.

Impact: These projects reflect my ongoing commitment to *open science*, *efficiency*, and *community-driven research and development*.

References

[1] Alessandro Fontanella et al. "ACAT: Adversarial Counterfactual Attention for Classification and Detection in Medical Imaging". In: *ICML* (2023).

[2] Adam Jelley et al. "Contrastive Meta-Learning for Partially Observable Few-Shot Learning". In: *ICLR* (2023).

[3] Antreas Antoniou. "Meta learning for supervised and unsupervised few-shot learning". In: *The University of Edinburgh* (2021).

[4] Timothy Hospedales et al. "Meta-learning in neural networks: A survey". In: *IEEE TPAMI* (2021).

[5] Antreas Antoniou et al. "Defining benchmarks for continual few-shot learning". In: *Meta-Learrning Workshop, NeurIPS*. 2020.

[6] Antreas Antoniou, Harrison Edwards, and Amos Storkey. "How to train your MAML". In: *ICLR*. 2019.

[7] Antreas Antoniou and Amos Storkey. "Assume, Augment and Learn: Unsupervised Few-Shot Meta-Learning via Random Labels and Data Augmentation". In: *arXiv preprint* (2019).

[8] Antreas Antoniou and Amos J Storkey. "Learning to Learn by Self-Critique". In: *NeurIPS*. 2019.

[9] Antreas Antoniou, Amos Storkey, and Harrison Edwards. "Data Augmentation Generative Adversarial Networks". In: *ICANN* (2018).

[10] Antreas Antoniou et al. "Dilated Densenets for Relational Reasoning". In: *arXiv preprint* (2018).

[11] Luke N Darlow et al. "CINIC-10 is not ImageNet or CIFAR-10". In: *arXiv preprint* (2018).

[12] Antreas Antoniou and Plamen Angelov. "A general purpose intelligent surveillance system for mobile devices using deep learning". In: *IJCNN*. IEEE. 2016.