Meghdad Kurmanji

Cambridge, UK – Global Talent Visa Holder
+44 7949 726 826 • mk2296@cam.ac.uk • meghdadk.github.io
meghdadk • meghdadk

Professional Summary

AI Scientist with 8+ years of experience designing and deploying scalable (generative) AI systems across research and production settings. Expertise spans the full lifecycle of machine learning systems, from prototyping novel models to building robust, production-grade pipelines for enterprise applications. Proven ability in delivering impactful solutions in areas such as large-scale model training, and multimodal learning, and AI Safety. Published in top-tier venues (NeurIPS, ICLR, etc) and secured competitive research fundings (£560k SPRIN-D, $150K Foresight) to advance privacy-preserving and safety-critical AI systems.

Experience

IQVIA **London, UK**
AI Research Scientist *2026–Present*

- Designed and trained transformer-based foundation models on longitudinal clinical data (patient timelines), enabling structured reasoning over long contexts.
- Developed agentic LLM systems for automated knowledge extraction and verification from large document collections (TMF), leveraging LangGraph and retrieval pipelines.

University of Cambridge **Cambridge, UK**
Postdoctoral Research Associate *2024–2026*

- Secured a **£150k Foresight AI Safety grant** to advance research on mechanistic machine unlearning and AI safety.
- Co-led a **€530k-funded** research project on scalable decentralized LLM pre-training, focusing on privacy-preserving and robust training pipelines.
- Developed novel machine unlearning techniques for efficient and secure data removal, addressing regulatory requirements such as GDPR compliance.
- Designed and evaluated distributed LLM training systems using PyTorch, Transformers, and MosaicML in large-scale environments.
- Built scalable evaluation pipelines for LLM fine-tuning across downstream tasks, improving reproducibility and benchmarking.
- Published research in top-tier venues including NeurIPS and ICLR.

University of Warwick **Coventry, UK**
Graduate Research Assistant *2020–2024*

- Pioneered machine unlearning algorithms addressing data deletion and privacy requirements in large-scale ML systems.
- Designed a continual learning framework achieving over **10× throughput improvement** compared to prior approaches.
- Developed a machine unlearning algorithm outperforming prior SOTA methods by up to **10%** across benchmarks.
- Secured **£150k** research funding from Huawei for ML-based database indexing systems.
- Led collaborations with Google DeepMind, including co-organizing the NeurIPS 2023 Machine Unlearning Challenge.
- Co-authored **7+ publications** in NeurIPS, SIGMOD, and CIDR.

Iran Telecommunication Research Center (ITRC) **Tehran, Iran**
Data Engineer *2019–2020*

- Built an end-to-end big data pipeline (crawling → Hadoop → OLAP), reducing data onboarding time by **5**×.
- Implemented scalable ETL workflows, achieving **5× query speed-up**.
- Integrated Elasticsearch with PowerBI dashboards, reducing reporting delays by **60%**.

Refah Retail Chain Stores Co. **Tehran, Iran**
Machine Learning Engineer *2017–2019*

- Led development and deployment of a real-time computer vision system for customer footfall analysis, achieving **81% accuracy** across 20 stores.
- Developed in-store heatmap analytics to optimize staffing and store layout decisions.
- Built a multi-modal recommendation system (LSTM + CNN), increasing customer engagement by **15%**.
- Implemented time-series regression models for customer behavior prediction.

Sensifai **Belgium (Remote)**
Deep Learning Engineer *2016–2017*
- Improved production audio classification accuracy by **9%** using multi-modal transfer learning techniques.
- Built an **88% accurate** music mood classifier using spectrogram-based CNN models.
- Optimized distributed video crawling pipelines, achieving **1.8× throughput**.

Education

University of Warwick **Coventry, UK**
Ph.D. in Computer Science *2020–2024*
Thesis: *Adaptability of ML-Based Database Systems* (**SIGMOD Honorable Mention Award**)
- Conducted the first study of data deletion (unlearning) in learned database systems (SIGMOD 2024).
- Developed **SCRUB**, a state-of-the-art unlearning algorithm for large-scale deep models (NeurIPS 2023).
- Proposed **DDUp**, a framework for efficient data insertion in learned database systems (SIGMOD 2023).
- Collaborated with Google DeepMind to launch the NeurIPS Machine Unlearning Challenge.

Tarbiat Modares University **Tehran, Iran**
M.Sc. in Computer Science, GPA: 3.67/4 *2014–2017*
Dissertation: Hand Gesture Recognition Using 2D and 3D Convolutional Neural Networks from Video

Isfahan University of Technology **Isfahan, Iran**
B.Sc. in Computer Engineering, GPA: 3.65/4 *2010–2014*

Skills

Core Expertise: LLM Pre-training & Fine-tuning, Decentralized ML, Machine Unlearning, AI Safety
Programming: Python, C++, SQL, Bash
ML Frameworks: PyTorch, TensorFlow, Hugging Face, Transformers
Distributed Systems: PyTorch DDP/FSDP, Ray, MosaicML, Slurm
MLOps & Cloud: Docker, CI/CD, AWS, Azure, Databricks, MLFlow, Weights & Biases
Data Systems: SQL, NoSQL, OLAP, Hadoop, Learned Indices

Honors & Awards

2026: Secured £150k Foresight AI Safety Grant
2025: NeurIPS Top Reviewer Recognition
2025: SIGMOD Jim Gray Doctoral Dissertation Honorable Mention
2024–2025: €530k SPRIN-D Grant (Co-lead, CambridgeFlower)
2023: Organizer, NeurIPS Machine Unlearning Workshop
2021: Best Presentation Award, WPCCS, University of Warwick
2020–2024: Graduate Scholarship (£25k/year), University of Warwick
2020–2024: Research Grant (£15k/year), Huawei

Selected Publications

ICLR 2025: DEPT: Decoupled Embeddings for Pre-training Language Models (top 1%)
NeurIPS 2025: Bridge the Gaps between Machine Unlearning and AI Regulation

NeurIPS 2024: What Makes Unlearning Hard and What to Do About It

SIGMOD 2024: Machine Unlearning in Learned Database Systems

NeurIPS 2023: Towards Unbounded Machine Unlearning

SIGMOD 2023: Detect, Distill and Update: Learned DB Systems Facing OOD Data

Full list: `https://scholar.google.com/citations?user=7t9HbecAAAAJ`

References

Available upon request.