

## Ratneshwaran (Ratnesh) Maheswaran

London, UK • [r.maheswaran@imperial.ac.uk](mailto:r.maheswaran@imperial.ac.uk) • +44 (0)7824 388253

[LinkedIn](#) [GitHub](#) [Portfolio](#)

**Research Profile:** AI researcher with experience in reliable, evaluation-led LLM systems for high-stakes settings, mechanistic interpretability, and reproducible ML pipelines for complex real-world datasets, including large-scale biomedical cohorts and time-series biosignals. Strong in experimental design, benchmarking, error analysis, data harmonisation, and building research software that supports transparent iteration. Interested in reliable and safe AI systems, evaluation and benchmarking, agentic AI, mechanistic interpretability, and AI for scientific and biomedical data.

## Education

---

**University College London (UCL), UK** Sep 2024 – Dec 2025  
MSc Artificial Intelligence for Biomedicine and Healthcare  
*Grade: Distinction*

**The University of Nottingham, UK** Sep 2021 – Jul 2024  
BEng (Hons) Electronic and Computer Engineering  
*Grade: Upper Second-Class Honours*

## Research Experience

---

**UK Dementia Research Institute, Imperial College London** Jun 2026 – Present  
Research Assistant, Sandor Lab (*supervised by Dr Cynthia Sandor*)

- Lead data access, harmonisation, and integration of large-scale, multimodal biomedical datasets for Parkinson's disease research, spanning longitudinal clinical cohorts including PPMI, OPDC, UK Biobank, and All of Us.
- Build reproducible preprocessing pipelines for wearable time-series data (wrist accelerometry and IMU), producing clean, analysis-ready inputs for a downstream self-supervised learning model.
- Apply signal processing methods (filtering, resampling, artefact handling) to derive harmonised wearable features across heterogeneous cohorts, with an emphasis on reproducibility and documentation.
- Contribute to self-supervised learning model development for wearable-based prodromal Parkinson's detection.

**UCL Centre for Artificial Intelligence** Oct 2025 – Present  
Research Assistant (*collaboration with Freedom from Torture, UK charity*)

- Prototyped multilingual, trauma-informed conversational agents using LLMs and agentic frameworks, including LangGraph and CrewAI.
- Built structured synthetic clinician-style dialogue datasets and evaluation harnesses to assess reliability and safety-related behaviours.
- Implemented conservative risk monitoring and escalation logic for high-risk content, with audit-friendly logging.
- Developing evaluation methods for LLM-derived workflow graphs using novel metrics.

**Secier Lab, UCL Genetics Institute** Mar 2025 – Sep 2025  
MSc AI Researcher

- Built reproducible pipelines (PyTorch, Scanpy, scvi-tools) for processing and analysing large-scale single-cell datasets across multiple cohorts.
- Benchmarked foundation model approaches for classification and state prediction using Accuracy, F1, AUROC, and AUPRC.
- Studied generalisation under batch effects and dataset shift using systematic diagnostics and stratified error analysis.
- Maintained clean research code, experiment tracking, and documentation to support reuse and preprint preparation.

## Teaching Experience

---

### UCL Computer Science

Mar 2026 – Jun 2026

Teaching Assistant, COMP0190: AI for Domain Specific Applications Project (*part-time*)

- Support delivery of a project-based AI module, helping students develop domain-specific AI applications from problem definition to implementation.
- Advise on project design, technical choices, experimentation, and effective presentation of methods and results.
- Provide practical support on debugging, model development, and communicating AI work clearly in reports and discussions.

## Selected Research Outputs and Research Software

---

### Grounded By Design: Safety-aware LLM Voice Assistant (Prototype)

2025 – 2026

- Research prototype exploring reliability and safety mechanisms for multilingual voice-based LLM support in sensitive contexts.
- Implemented monitoring and guardrails to detect high-risk content and route to conservative handover pathways.
- Built an end-to-end STT and TTS pipeline with instrumentation for auditability, including logging, transcript capture, and outcome signals.
- Manuscript in preparation.

### sae-interp: Sparse Autoencoders for Mechanistic Interpretability

2026

- Built a full pipeline for training sparse autoencoders on GPT-2 activations to extract interpretable features, building on Anthropic's monosemanticity research.
- Implemented activation extraction and caching, SAE training with both TopK and L1 sparsity, dead neuron resampling, and feature analysis dashboards (top activating examples, decoder similarity, dead feature diagnostics).
- [GitHub](#)

### Multimodal Time-Series Preprocessing Pipeline (HAR, EEG, ECG)

2026

- Built a reproducible pipeline for downloading, preprocessing, and validating five public datasets (PAMAP2, WISDM, mHealth, EEGMMIDB, PTB-XL) into a unified channels-first format for downstream self-supervised learning.
- Implemented signal cleaning (bandpass, notch, common average reference), resampling, label harmonisation, and per-modality windowing across HAR, EEG, and ECG modalities.
- Wrote 84 automated validation checks and 113 unit tests covering array integrity, harmonisation, leakage control, and resource estimates.
- Profiled the sequential pipeline and built a parallel ECG loader using Python multiprocessing, achieving roughly 20x speedup with bit-identical output.
- [GitHub](#)

### Cuff-less Blood Pressure Prediction using ECG and PPG (CNN-LSTM)

2024

- Undergraduate dissertation on time-series biosignal preprocessing and modelling for continuous, non-invasive blood pressure estimation.
- Built an end-to-end workflow for ECG and PPG processing, including filtering, noise and motion artefact handling, and segment preparation for modelling.
- Implemented classical regression baselines with engineered physiological features and a CNN-LSTM hybrid model using raw ECG and PPG sequences.
- Evaluated on held-out data using MAE and RMSE alongside agreement diagnostics (Bland-Altman) to characterise error distributions in wearable-style settings.

### Zero-Shot Cell Annotation with Foundation Models

2025

- Implemented zero-shot pipelines for cell classification and multi-class state prediction across multiple datasets.
- Integrated evaluation using Accuracy, F1, AUROC, and AUPRC, alongside diagnostic visualisations to assess generalisation and dataset shift.
- [Blog Post](#), [LangCell GitHub](#), [Geneformer GitHub](#)

## **Record2Table: Multi-document Processing Pipeline for Structured Outputs** 2026

- Pipeline to ingest multiple documents per case and produce consolidated structured JSON, tables, and concise summaries.
- Implemented schema-based extraction, normalisation utilities, provenance-friendly outputs with source snippets, and evaluation scripts.

## **Enhancing Aspect-Based Sentiment Analysis with Adversarial Training and Hierarchical Aggregation (BERT)** 2025

- Improved aspect extraction and sentiment classification by combining adversarial training with hierarchical aggregation.

## **LitAgent: AI Agent for Literature Search and Synthesis** 2025

- LangGraph-driven multi-agent system for automated search and synthesis across *arXiv* and *PubMed*, with Crossref enrichment.
- Exports structured outputs in Markdown, JSON, and CSV, supports CLI and REST API workflows, and includes critique checks to flag overclaims.
- [GitHub](#)

*Additional projects and technical write-ups are available on my [website](#) and [GitHub](#).*

## **Research Interests**

---

Reliable and Safe LLM Systems, AI Evaluation and Benchmarking, Agentic AI and Workflow Reasoning, Mechanistic Interpretability, Reproducible Machine Learning, Time-series and Biosignal Modelling, Responsible AI, AI for Scientific and Biomedical Data

## **Skills**

---

**Programming:** Python, R, SQL, C/C++, MATLAB, Java, TypeScript

**ML and Data:** PyTorch, TensorFlow, scikit-learn, NumPy, Pandas; evaluation (Accuracy, F1, AUROC, AUPRC), error analysis

**Time-series and Signals:** filtering (bandpass, notch), resampling, artefact handling, peak-based features, basic frequency-domain analysis (ECG, PPG, EEG, IMU)

**Scientific ML:** Scanpy, scvi-tools; single-cell data preprocessing and QC, batch-aware benchmarking, embedding diagnostics

**LLM Systems:** LangChain, LangGraph, CrewAI; RAG and embeddings; OpenAI APIs; AWS Bedrock; monitoring and structured evaluation

**Interpretability:** Sparse autoencoders (TopK, L1), activation extraction and caching, feature dashboards, transformer hook-based analysis

**Compute and Reproducibility:** Linux, bash, HPC (SLURM), multiprocessing, Snakemake or Nextflow, conda or poetry, CI (GitHub Actions), Docker

**Tools and Platforms:** Git, FastAPI, PostgreSQL, Node.js, React, Streamlit, LaTeX, AWS

**Additional:** Computer vision (segmentation, IoU, Dice), geospatial tooling (GeoPandas, Rasterio, QGIS basic)

## **Leadership and Service**

---

### **Ensemble, London**

2026 – Present

Founder

- Setting up a London-based research and community organisation working at the intersection of AI safety, governance, and the philosophy of AI.
- Designing programmes including an ambassador network across London universities, an invitation-only fellowship, and a recurring AI seminar series.
- Leading on founding manifesto, recruitment of ambassadors, partnerships with researchers and institutions, and event programming across London.

## **Open-Source Assistive Devices (OPAD), University of Nottingham**

Sep 2022 – Jul 2024

President, Faculty of Engineering Society

- Directed 12+ student-led assistive and medical device projects, coordinating cross-disciplinary teams across hardware, software, and data workflows.
- Set requirements, milestones, and QA checks for reliable prototypes aligned to user needs; secured £1500 in grant funding.

## **40seconds.org (Ratneshwaran Foundation)**

Jan 2020 – Present

Founder and CEO (*currently on hiatus*)

- Founded a global non-profit promoting mental health awareness and human rights among youth; led initiatives across 19 countries.
- Coordinated 120+ volunteers, delivered initiatives and articles, and raised £2000 for COVID-19 relief.

## **Honours**

---

- **Vice-Chancellor's Medal (Top 0.1%), University of Nottingham** 2024
- **EEE Undergraduate Studies Award, University of Nottingham** 2024
- **Nottingham Advantage Awards, University of Nottingham** 2023 – 2024
- **Union Prize and Student Representative of the Year, University of Nottingham** 2022
- **Top 6 Teen Heroes, India** 2021
- **Speaker Invitations (selected): World Summit AI, TEDx, university conferences**

**References:** Available upon request.