IAN GRIFFITH

PhD Candidate @ Harvard-MIT | AI Research Scientist

Cambridge, MA iangriffith@g.harvard.edu
github.com/IanMGriff

PROFESSIONAL SUMMARY

PhD computational neuroscientist developing human-level AI systems for audio perception and attention. Created the first model to achieve human-level performance on the cocktail party problem (Nature Human Behavior). Expertise in self-supervised learning, ML, human-AI alignment, and large-scale model evaluation. Invited speaker at Meta, MIT, and Flatiron Institute.

EXPERIENCE

Research Scientist Intern, Machine Learning

Meta Reality Labs, Redmond, WA — Sep 2025 - Present

Researching spatial-audio LLMs and distillation objectives for on-device systems

- Created ASR-based loss function for on-device speech enhancement, yielding 5% intelligibility and 30% preference gains on-par with 10x larger models
- Introduced robust evaluation infrastructure to benchmark speech enhancement systems across 11+ real-world use cases
- Invited by senior researchers to translate findings into production-viable technology and contribute to patents

Research Affiliate

Flatiron Institute Center for Computational Neuroscience, New York, NY — Sep 2024 - Present

Focus on self-supervised representation learning algorithms for perception models

- Developed equivariant contrastive learning objectives for audio perception models achieving 10% better generalization than supervised baselines
- Optimized distributed transfer learning workflows across 128 GPU nodes, reducing experimental runtime by 40%
- Achieved top 1% human neural alignment in model representations, bridging biological perception and machine learning

PhD Candidate

MIT Laboratory for Computational Audition, Cambridge, MA — Sep 2020 - Present

Researching human-centric audio models for attentional selection and music LLMs

- Devised brain-inspired attention mechanism and learning algorithm for human-level selective listening (in press at Nature Human Behavior, first author)
- Created novel audio simulation dataset (4M+ scenarios) enabling sim-to-real transfer that predicted 2 previously unknown human behaviors
- Built first human-AI alignment benchmark for auditory attention with 12+ matched human and model experiments
- Designed audio-symbolic music generation LLMs, improving human-model alignment metrics by 20%

EDUCATION

PhD in Speech and Hearing Biosciences and Technology

Harvard University, Cambridge, MA — Aug 2020 - Present

Specialization in Perception, Machine Learning, and Artificial Intelligence

BA in Cognitive Science

University of California Berkeley, Berkeley, CA — Aug 2016 - May 2018

Concentration in Artificial Intelligence and Computational Neuroscience

SKILLS

Python

PyTorch

TensorFlow

LLMs

Self-Supervised Learning

Generative Modeling

Computer Vision

Automatic Speech Recognition

Natural Language Processing

Human-Model Alignment

Distributed Computing

Psychophysics

CORE GRADUATE COURSES

Machine learning MIT 6.790

Spoken Language Processing MIT 6.862

REFERENCES

Antje Ihlefeld
Meta

Josh McDermott
MIT BCS

Malcolm Slaney
Stanford | ICSI

James Glass
MIT CSAIL