

CĂTĂLINA CANGEA

S2, 8 Handyside Street, London N1C 4DJ, United Kingdom

Email ◊ Personal website ◊ LinkedIn

PERSONAL PROFILE

Staff Research Scientist and former **Generative Music co-lead** at Google DeepMind. Passionate about solving real-world problems with ML. Obtained my PhD from the University of Cambridge.

PREVIOUS PROFESSIONAL EXPERIENCE

Quantitative Researcher—Qube Research & Technologies October 2024–April 2025

Senior Research Scientist—Google DeepMind May 2023–June 2024

As 1/3 Generative Music **co-leads** (team of ~40), I solicited ideas, planned projects, enabled communication between leadership & team and research & product workstreams, helped maintain momentum, ran syncs. **IC work** on Lyria and Music AI tools, model controls, finetuning. Led **demos** for stakeholders.

Research Scientist—DeepMind October 2021–April 2023

Multimodal learning and generative methods for long-range sequential data. Co-led several research projects, published at ICML 2022 and hosted an RS intern. External mentoring at EEML and LOGML.

Research Scientist Internship—DeepMind July 2020–November 2020

Hosted by Piotr Mirowski in the Robotics, Embodied Agents and Lifelong learning (REAL) team led by Raia Hadsell. I executed my self-proposed project fully remote, which led to a full-time conversion offer.

ML Research Consultancy—Relation Therapeutics June 2020–July 2020

Developed (graph-)ML solutions to aid in drug development and repurposing efforts.

AI Residency—X, the moonshot factory May 2019–August 2019

Worked on an early-stage project, adapting SotA techniques to track changes in code. Patent issued.

ML Research Internship—Mila July 2018–September 2018

Collaboration with Aaron Courville on a visual reasoning project which resulted in a novel benchmark and alternative perspective on EQA-style tasks. BMVC publication and NeurIPS ViGIL spotlight talk.

Software Engineer Internship—Facebook June 2016–September 2016

Worked on the LogDevice team, making client operations on a distributed RocksDB data store more efficient and flexible, while leading to fewer system failures.

Software Engineer Internship—Facebook July 2015–September 2015

Worked on iOS Product Infrastructure to reduce the time taken by the Facebook iOS app to load content close to the screen current view. Improved the infrastructure and network request prioritisation system.

Student Training in Engineering Program—Google June 2014–September 2014

Added processing progress for video uploads on YouTube. Developed a JavaScript client implementation that requests processing information from the server and thumbnail rendering of processing videos.

SELECTED PROJECTS

Lyria, Dream Track and Music AI tools (core team, co-led previous foundational project)

Contributed to both fan engagement and creative exploration avenues: GDM tech lead for the Dream Track Quality Workstream; found new data signals, finetuned models and designed human evals.

Active Acquisition for Multimodal Temporal Data: A ... Decision-Making Task (second)

Active feature acquisition for high-dimensional, multimodal inputs. We solved a synthetic scenario and learned cost-reactive acquisition on Kinetics-700 and AudioSet. *TMLR 2023, NeurIPS-W FMDM 2022*.

General-purpose, long-context autoregressive modeling with Perceiver AR (second)

Autoregressive architecture that attends to over 100k tokens, bypassing the usual Transformer $O(N^2)$ complexity. Strong performance/state-of-art on image, language and music. *ICML 2022*. Patent issued.

VideoNavQA: Bridging the Gap between Visual and Embodied QA (first)

Novel task that studies QA performance in EQA-like settings with nearly-ideal navigation paths. Generalised VQA models for temporal reasoning. *BMVC 2019, NeurIPS-W 2019 ViGIL* spotlight.

XFlow: Cross-modal Deep Neural Networks for Audiovisual Classification (first)

Cross-modal SotA architectures for temporally-aligned data. *IEEE Transactions on Neural Networks and Learning Systems 2019, ARM Research Summit 2017, ICDL-EPIROB-W 2017 CMCML*.

EDUCATION

Department of Computer Science and Technology, University of Cambridge, UK

PhD in Machine Learning (10/2017–03/2021) *No corrections (Nic Lane, Xavier Bresson) (Thesis)*

MPhil in Advanced Computer Science (10/2016–07/2017) *Distinction (Publication)*

BA in Computer Science (10/2013–06/2016) *First Class (Thesis, Code)*

Colegiul Național "I.L.Caragiale", Ploiești, România

Computer Science & Mathematics (09/2009–05/2013) *Valedictorian*

ACADEMIC TEACHING / SUPERVISING

Lectures: *Graph Generation and Probabilistic Methods, Graph Generation Methods*

Selected supervised projects: *Structure-aware Generation of Molecules in Protein Pockets, Goal-conditioned Reinforcement Learning in the Presence of an Adversary (NeurIPS-W DeepRL), Representation Learning for Spatio-Temporal Graphs (ICLR-W RLGM), Dynamic Temporal Analysis for Graph Structured Data (ICLR-W RLGM), Benchmarking Graph Neural Networks using Wikipedia (ICML-W GRL+ spotlight), The PlayStation Reinforcement Learning Environment (NeurIPS-W DeepRL)*

Courses: AI, Databases, Discrete Maths, Foundations of CS, Logic & Proof, ML & Real-World Data

PROFESSIONAL SKILLS

Programming Languages Python (since '15), C/C++ (competitive programming, '16 internship)
Objective-C ('15 internship), JavaScript ('14 internship)

Frameworks Standard ML, Java, C# (university)
PyTorch, JAX, Keras, TF1, PyG, Haiku, Graph Nets, Sonnet, CUDA

INTERESTS

Languages English (fluent), French (beginner), German (beginner), Romanian (native)

Music Guitar, piano, vocals; live performances, recording, writing lyrics

Sports Rowing (Darwin/King's College 2nd/1st Women's VIII), cycling, yoga