

Benjamin Dayan

✉ benjamin.l.dayan@gmail.com • 🌐 benjid.dev • 🗣️ [BenjiDayan](#)

Education

ETH Zurich

MSc Data Science, GPA: 5.6/6

Zurich, Switzerland

2021–2023

- **Thesis:** Generative Models for Social Networks (6/6)
- **Courses:** Reinforcement Learning, Natural Language Processing, Probabilistic AI, Big Data, Computer Vision, Computational Statistics, Neural Systems & Bio-inspired ML, Information Theory

University of Cambridge

BA (Hons) Mathematics, Upper Second Class (2:1)

Cambridge, UK

2017–2020

- **Courses:** Statistical Modelling, Linear Algebra, Numerical Analysis, Number Theory, Statistical Learning Theory, Markov Chains, Analysis, Graph Theory, Mathematics of Machine Learning

Experience

Listen Labs

Member of Technical Staff

San Francisco, California

Jan 2026–Present

- Full-stack AI engineer on an AI-moderated user-research platform (Next.js, TypeScript, Supabase, Temporal).
- Owned the speech-to-text pipeline & the LLM interviewer system, and worked on study creation functionality.
- Improved interviewer quality by addressing tone and leading-question issues through prompt engineering, tracing and observability, building dashboards in LangSmith and in-house tools.
- Integrated and A/B-tested multiple streaming & offline transcription providers; built QA and eval systems to mitigate mistranscriptions with keyword upweighting, ensemble models and LLM rewriting.
- Delivered product features end-to-end, including expanding platform localization to 120+ languages, per-language media variants, multi-speaker detection, interviewee feedback and rating, custom recruitment creation & billing, NDA form gating with PII isolation & embedded PDFs.
- Led a weekly coordination session with our researcher team to align on interviewer performance and desirable features. Supported urgent sales product requests and issues on live pilots. Made our recruiters a referral bot.

Centre for Human-Compatible AI, UC Berkeley

Research Intern (Bayesian Networks, GNNs, PyTorch, MCMC)

Berkeley, California

May 2025–Aug 2025

- Developed a meta-learning approach using GNNs to propose MCMC Gibbs blocks to speed up Bayesian network approximate inference.
- Delivered theoretical results on mixing time. Built a PyTorch pipeline for synthetic graph dataset generation, model training, and inference evaluation.
- [Paper](#) accepted to NeurIPS 2025 Frontiers of Probabilistic Inference workshop.

Finster AI (Startup founded 2023)

AI Engineer (Python, LLM API prompting, Backend & Data)

London, UK

Feb 2024–Mar 2025

- Built an LLM agent with RAG from the ground up. Wrote function-callable data source APIs, designed a multi-step agent architecture, and prompt-engineered LLMs to perform subtasks.
- Developed a novel citation algorithm for RAG textual and numerical data to mitigate LLM hallucinations.
- Built a data scraping, metadata extraction, chunking/embedding pipeline for RAG.
- Worked with tools such as: Gemini flash/pro, GPT series, Weaviate Vector DB, Cohere Rerank, GCP, MongoDB, Kubernetes, Flask/FastAPI, Docker, Cloud Run/Jobs.

ETH Zurich

Teaching Assistant (NLP) & Research Assistant (Complex Networks)

Zurich, Switzerland

Jul 2022–Jan 2024

- Created and taught a theory & coding module on dependency parsing, the Matrix-Tree theorem and Edmonds' algorithm.
- Worked on 3 papers (see publications). [NLP automata algorithms](#) (EMNLP 2022); [Fitting and evaluating a graph model of social networks](#) (1st author, Complex Networks 2024); [Benchmarking a graph algorithm on different networks](#).

Lyst LTD

Associate Software Engineer (Python, SQL, Backend API, Data)

London, UK

Mar 2021–Sep 2021

- Migrated a microservice to process and serve retailer data via a data lake architecture.
- Developed REST APIs and coordinated with downstream teams.
- Worked with Django, PostgreSQL, PySpark, pandas.

Panakeia AI (Startup founded 2018)

Software Engineer Product and Research (Python, Computer Vision, Data)

London, UK

Jun 2020–Jan 2021

- Researched and implemented a PCA-based image colour normalization preprocessing method for an ML training pipeline.
- Added model-parameter logging in a Postgres DB to make training runs reproducible.

BlackRock

Intern Financial Modelling Group (Python, R, Factor Models)

London, UK

Jul 2019–Sep 2019

- Ported an equity risk model to Python to aid future prototyping and research efforts.
- Documented and presented the model's methodology to colleagues.

Publications

Svete A, **Dayan B**, Cotterell R, Vieira T, Eisner J. [Algorithms for Acyclic Weighted Finite-State Automata with Failure Arcs](#). **EMNLP 2022 Proceedings**

Dayan B, Kaufmann M, Schaller U. [Expressivity of Geometric Inhomogeneous Random Graphs—Metric and Non-metric](#). **Complex Networks Conference 2024**

Cerf S, **Dayan B**, De Ambroggio U, Kaufmann M, Lengler J, Schaller U. [Balanced Bidirectional Breadth-First Search on Scale-Free Networks](#). Under submission to **Electronic Journal of Combinatorics**

Dayan B, Mishra B, Svegliato J, Russell S. [GNN-Guided Block Selection in Gibbs MCMC](#). Accepted to **SPIGM and FPI workshops at NeurIPS 2025**

Technical Skills

Programming Languages: **Expert:** Python **Familiar with:** TypeScript, R, SQL, Julia, C++, C#, Haskell, Perl, Rust, Lean

ML/AI: PyTorch, NumPy, Pandas, JAX, TensorFlow, Transformers, RNNs, CNNs, NLP, W&B

Software/Tools: Linux, Docker, Kubernetes, GCP/AWS, Git, Latex, PostHog, Datadog

Languages

English: Native **Mandarin:** Professional working proficiency **German, French:** B1 (intermediate proficiency)

Miscellaneous

VR Development: Led coding for student team-developed VR headset/PC games

- benjidayan.itch.io/cybercityvr (developer lead)
- kevponce4.itch.io/big-bang-blasters (art and code contributor)

Random Achievements: Sang hymns for Queen Elizabeth II, UK National Cipher Challenge gold medal (2017), Piano Grade 8, Undergrad college badminton captain, 3hr17 marathon (Paris 2025)

Interests (outside of tech/AI): Economics, Politics, Tennis, Badminton, Bouldering, Reading, Baking (mostly eating)