Taliesin Beynon

Research engineer · London

contact@tali.link · tali.link
github.com/taliesinb · git.tali.link
@taliesinb · linkedin.com/in/taliesinb

I'm a curious generalist who rapidly assimilates and synthesizes ideas from disparate fields. I led the creation of Mathematica's deep learning framework at Wolfram, then left for independent research in subfields of mathematics and neuroscience that seemed ripe for neurosymbolic deep learning, co-authoring with various collaborators. In 2024 I joined a London startup with similar interests, where I'm now a staff engineer. I also build open-source infrastructure for long-horizon AI agents: statically verified and type-checked shell execution, flexible browser use, transcript analytics for agent evaluation, and a provenance-tracked research library. Above all, I'm drawn to machines that discover, not just recall.

EXPERIENCE

Symbolica · research engineer → research director → staff engineer — 2024 – present

- **typed RNNs**: developed a mathematically rigorous approach to generalizing RNNs to model distributions over terms of inductive families of types (e.g. algebraic data types)
- **torch kernel**: wrote a C++ PyTorch layer to reimplement a naive recursive Python syntax-tree position embedding as a packed parallel state machine employing GPU-friendly sparse matmul, giving a 1117× faster forward pass and 16× faster training step overall
- **theorem proving**: pioneered a multicategorical approach to representing Metamath proofs as terms of an inductive family of types
- **Lean**: led formalization and correctness proofs of the internal hypergraph rewriting engine
- **embodied RL**: built flatsim, a Rust/PyO3 simulator of agents with steerable eyes and 1D RGB retinas exploring a ray-traced 2D environment
- **experiment management**: wrote jobman, a multigraph job manager for parameterized sweeps and derived jobs, local or on a cluster; Merkle-hashed configs mean no job ever runs twice
- **Agentica SDK**: designed a low-level RPC protocol for virtualized "teleportation" of Python objects between agent harnesses, resulting in a 36% ARC-AGI-3 result in March 2026
- **Agentica Chat**: built a protocol for agents to live-code browser interfaces, significantly more effective and powerful at the time than Claude Artifacts

Independent research & teaching · deep learning, neuroscience, mathematics — 2018 – 2024

- xox: a schematic model of connectome neuroevolution (*Nat. Comm.*), with Dániel Barabási
- rainbow arrays: a diagrammatic calculus for classical and named-axis array programming
- quiver geometry: a discrete, computational framework for geometries with a "smallest scale"

Wolfram Research · lead developer, Deep Learning Group — 2016 – 2018

- led creation of Mathematica's built-in deep learning framework (`NetGraph`, `NetChain`, `NetTrain`)
- wrote & reviewed 50+ pages of reference docs and tutorials; upstream contributions to Apache MXNet (dynamic-shape ops)

Wolfram Research · manager & lead developer, Advanced Research Group — 2013 – 2015

- co-designed automated machine learning (`Classify`, `Predict`)
- designed dozens of new language features, including tabular data and associative arrays
- data science of the Facebook world: 1.26M "data donors", featured in *Wired*

Wolfram | Alpha · research programmer → manager, Special Projects — 2010 – 2013

- automated data & image analysis for W|A Pro; supervised Personal Analytics for Facebook

Earlier: Nimbula (cloud computing, 2009); CHPC and iThemba LABS internships (2008).

PAPERS

Complex computation from developmental priors
D. L. Barabási, **T. Beynon**, Á. Katona, N. Perez-Nieves · *Nature Communications 14, 2226* — 2023

Diagrammatic calculus and generalized associativity for higher-arity tensor operations
C. Zapata-Carratalá, X. D. Arsiwalla, **T. Beynon** · *Theoretical Computer Science 1020* — 2024

Hypermatrix algebra and irreducible arity in higher-order systems
C. Zapata-Carratalá, M. Schich, **T. Beynon**, X. D. Arsiwalla · *Advances in Complex Systems 26* — 2023

OPEN SOURCE

Agentica: Symbolica's Python & TypeScript SDK for code-mode agents that reason inside a persistent runtime; basis of Symbolica's open-source ARC-AGI-2 & 3 agents

DeepSea: native macOS coding-agent workbench over DeepSeek Harness: 36 plugins

- brow: real Safari & Chrome windows for agents: 8 browser tools via a Unix socket, CLI & MCP
- mash: a bash VM for agent harnesses: typechecks against sandbox policy

papyrus: local-first research library: provenance-tracked SQLite, local-model classification and extraction, full-text search, arXiv/OpenAlex matching, MCP server

tensatory: loss-landscape viewer: isosurfaces of metrics sampled around trained networks

klein: lays out and routes tensor networks and circuits from topology alone, as static SVG

nlab: pipeline turning the ~20k-page nLab category theory wiki into a single typst corpus, compiled to a modern static site with SVG diagrams, MathML and full-text search (live)

Earlier: hypernet (PyTorch hypernetworks), floatworld (evolving RNN-controlled agents), spieeltjie (multi-agent RL on differentiable games), quiver geometry.

RESEARCH INTERESTS

NEUROSCIENCE: Tolman-Eichenbaum machines, cognitive maps, grid & place cells, predictive coding

AI: machine discovery, representation learning, energy-based models & associative memory, reinforcement & meta-learning, interpretability, causality, vector symbolic architectures, distributional semantics, knowledge representation & reasoning, differentiable programming

MATHEMATICS: representation theory, applied category theory & categorical semantics, type & proof theory, algebraic geometry, string diagrams & graphical calculi, hypergraphs

EDUCATION

BSc (Hons), Mathematics — 2005 – 2008
University of Cape Town

SUMMER SCHOOLS

IBRO-Simons Comp. Neuroscience	2019
Reinforcement Learning (RLSS)	2019
Categories for AI	2022
Deep Learning Indaba	2017
Wolfram Summer School	2009

SELECTED TALKS

Deep learning vs the brain — 2019
IndabaX South Africa, w/ S. Bodenstein

Interpretable ML for fund. science — 2020
ICLR workshop panel

The power of named axes — 2023
ArrayCast podcast

Array algebra — 2022
Deep Learning Indaba, Tunisia

Intro to reinforcement learning — 2019
Deep Learning IndabaX Sudan

TEACHING & WRITING

Build your own TensorFlow — 2019
autodiff from scratch, DL Indaba

Wasserstein GANs — 2019
Depth First Learning curriculum

Rainbow array algebra — 2023
graphical calculus for arrays

Deep Learning Indaba — 2018 – 2022
practicals team: Kenya, S. Africa, Tunisia

Wolfram Summer School — 2010 – 2022
tutor & mentor

SKILLS

ML: PyTorch, JAX

LANGUAGES: Python, C++, TypeScript, Swift & Metal, Go, Mathematica, Rust, Zig, Haskell, Julia, Lean

PATENTS

co-inventor with Stephen Wolfram on 4 granted US patents