Ben Thorne

ben0x1@icloud.com | b-thorne | ben.kfq.io | San Francisco, CA

EXPERIENCE

Senior Machine Learning Engineer — 2023 - present

Atomic Industries — Remote / San Francisco, CA

Physics-based automated design of plastic injection molds — Julia · Python

- Led the end-to-end design and implementation of an internal tool for automated cooling circuit design in injection molds, combining advanced physics, computational geometry, and optimization.
- Sole architect of the tool; designed core data structures, parallel computing architecture, and simulation pipelines in Julia. These included:
  - A custom solver for elliptic PDEs, enabling scalable thermal simulation in highly detailed 3D geometries.
  - A geometry package for 3D spatial queries and acceleration structures, used to evaluate physical constraints and manufacturability in real time.
  - A distributed evolutionary algorithm framework tailored to the discrete, constrained nature of the design space.
- Acted as technical lead for a small team, setting direction, running sprint planning, and reviewing code.
- Collaborated closely with manufacturing engineers to validate simulation outputs against physical behavior, and worked with product leadership to support customer PoCs and Series A fundraising with design outputs.

Machine Learning Engineer — 2023

National Energy Research Scientific Computing Center (NERSC), Berkeley Lab — Berkeley, CA

Machine learning benchmarks for high-energy physics (NeurIPS'24) [PDF] [Project Page] — Python · PyTorch

- Coordinated the FAIR Universe project: a cross-institutional collaboration of 10+ scientists and engineers from Berkeley Lab, University of Washington, Paris-Saclay, and ChaLearn to build a ML challenge platform for scientific computing.
- Contributed to the design of uncertainty-aware ML benchmarks and scoring metrics for high-energy physics datasets, shaping the first public challenge launched at NeurIPS 2024.
- Worked closely with engineers at NERSC to support the deployment of containerized ML payloads on the Perlmutter supercomputer, using Kubernetes and Rancher for orchestration.

Postdoctoral Researcher — 2019 - 2022

University of California, Davis — Davis, CA

Generative modeling of astronomical images using VAEs in PyTorch [PDF] [GitHub] — Python · pytorch

- Curated novel training dataset of 1,000 images of the interstellar medium from public data.
- Designed and trained convolutional variational autoencoder in PyTorch.
- Applied trained model to Bayesian inverse problems: data imputation, denoising, inference.

Differentiable physical models for cosmology with CUDA and auto-diff [PDF] — Julia · CUDA

- Developed automatically-differentiable foreground model extension to CMBLensing.jl.
- Implemented GPU-acceleration with CUDA.jl.
- Developed sparse approximations & preconditioners to speed up log likelihood evaluation by $\sim 100\times$.
- Distributed Bayesian inference pipeline across 10's of A100 GPU nodes on *Perlmutter*.
- Used to analyze 10^6-pixel 3-channel images from *BICEP/Keck-South Pole Telescope* joint analysis.

Python Sky Model [PDF] [GitHub] [Project Page] Python · numba

- Original author of the Python Sky Model package, *pysm*, for simulating microwave sky maps. This package has become the de facto standard for simulating cosmic microwave background surveys, with over 300 citations.
- Uses numba and MPI for execution in high-performance computing environments.
- Frequently runs across 100's of nodes at NERSC in large-scale simulation campaigns for the *Simons Observatory* and *Stage-IV* experiments.
- Due to widespread community reliance on the package, since 2020 it has been maintained and developed by the *Pan-Experiment Galactic Science Group*, an academic consortium.

Other duties

- Co-organized the weekly cosmology seminar from January 2020 to December 2022.
- Co-supervised PhD students on various projects in machine learning and physics.
- Regularly delivered seminars and conference talks.

EDUCATION

PhD in Astrophysics 2015 - 2019

University of Oxford, Princeton University & Kavli IPMU Oxford, Princeton, Tokyo

Supervisors: Prof. Jo Dunkley (Princeton & Oxford), Prof. Nobuhiko Katayama (Kavli IPMU, University of Tokyo)

Performance forecasts for the Simons Observatory (SO) [PDF_1] [PDF_2] [GitHub] Python · C

- Wrote maximum likelihood estimation algorithm for pixelized satellite data using Numba.
- Performed Monte Carlo simulations to forecast constraints on primordial gravitational waves from SO data.

Data analysis for the Atacama Cosmology Telescope [PDF] Python · numba

- Computed Bayesian priors for 2018 likelihood analysis by calculating and fitting power spectra of public sky maps.

Axion-SU(2) inflation and chiral gravitational waves [PDF] [GitHub] Python · C

- Developed analytical and numerical predictions for chiral primordial gravitational wave signals in parity-violating cosmologies.
- Modified C code, *CLASS*, to compute parity-violating CMB signatures, and quantified the sensitivies of the upcoming laser interferometer and CMB satellites *LISA* and *LiteBIRD* to the derived signals.
- Resulted in a *high-impact paper* (100+ citations).

M.S. & B.S. in Physics, first class honours (4.0 GPA) 2011 - 2015

University of Oxford, New College Oxford, UK

Photometric decomposition of barred and double-barred galaxies [PDF] Python · IDL

- Collected dataset of galaxy images from the *Sloan Digital Sky Survey*.
- Decomposed galaxy images into components using least-squares fitting, constraining bar and spiral structure.

SKILLS

Python: pytorch, jax, scipy, dask, mpi4py, numba, numpy
Julia: CUDA.jl, StaticArrays.jl, Distributed.jl, Flux.jl, Zygote.jl
C: MPI, CLASS
Other: Docker, Nomad, Git, DVC, AWS [S3, EC2], Quarto, Javascript, Slurm, MPI