Cambridge Yang

camyang@csail.mit.edu
people.csail.mit.edu/camyang/
thisiscam

Education

2020–2025 *Ph.D. in EECS (Minor in Computational Physics)*, *MIT, Cambridge*, MA.
Thesis: *On the Learnability of Reinforcement-learning Objectives*
Advisor: Michael Carbin

2017–2020 *S.M. in EECS*, *MIT*, Cambridge.
Thesis: *Simplifying Dependent Reductions in the Polyhedral Model*
Advisor: Michael Carbin

2013–2017 *B.S. in EECS*, *University of California*, Berkeley.
Research Mentors: Michael Lustig, Koushik Sen, Sanjit Seshia

Research Interests

AI safety and alignment; design of safe and reliable AI systems. My past research includes theoretical work on reinforcement-learning objectives, programming systems for intelligent agents, and probabilistic programming. I am especially interested in bridging these theoreticalfoundations with empirical research on frontier models and agents — developing frameworks and experiments that connect computability, objective design, and emergent behaviors in frontier AI systems.

Experience

2025-Current *Research Scientist in AI*, *Basis Research Institute*, Cambridge, MA.

2019 *Research Intern*, *DeepMind*, London, UK.

2014 - 2015 *Software Researcher*, *Sellegit Inc. (Venture Startup)*, Berkeley, CA.

Research Projects

Theory of Reinforcement-learning Objectives

2025-Current *World-Model Learning*, *Industry Research (Basis)*.
Built a novel benchmark for world-model-learning agents. Developed a reward-free, representation-agnostic evaluation framework.

2020-2023 *Tractable Reinforcement-learning Objectives*, *Ph.D. Research*.
Proved formal theorems on the learnability of linear temporal logic (LTL) objectives and showed a dichotomy of learnability in the LTL hierarchy. Proved the link between computability and learnability of general objectives.

2020-2021 *Geometric Linear Temporal Logic*, *Ph.D. Research*.
Formalized the semantics of geometric linear temporal logic (GLTL) and proved various properties, including convergence to LTL.

Probabilistic Programming

2021-2022 *Affine Probabilistic Graphical Model*, *UROP Research (Mentor)*.
Formally defined an affine probabilistic graphical model in the polyhedral model framework. Explored compiling efficient inference procedures for probabilistic affine graphical models.

2018-2020 *Complexity Optimization of Dependent Reductions*, *Master's Research*.
Developed an algorithm that reduces the complexity of dependent reductions in the polyhedral model, achieving orders-of-magnitude speedups in wall-clock time. These reductions are prevalent in probabilistic programming and scientific computing.

Machine Learning and Compilers

2020-2021 *Neural Compilation*, *Graduate Research*.
Applied deep reinforcement-learning algorithms to compiler tasks including auto-vectorization, program-size minimization, and register allocation.

2019 *Graph Conditional Variational Autoencoder*, *Research Internship at DeepMind*.
Designed a graph conditional variational autoencoder to generate probabilistic samplers for XLA graphs (the intermediate representation of JAX programs).

Machine Learning and Computational Physics

2024 *Simulated Fluid Sensing from Temperature Data*, *Computational Physics Research*.
Simulated the heat partial differential equation to infer fluid velocity and thermal diffusivity using two temperature sensors; used probabilistic modeling and a two-layer neural network to infer fluid mixtures from the simulation.

2023 *Chaos in Classical Mechanics Simulation*, *Computational Physics Research*.
Studied the boundary of chaos and orbits around equilibria in classical mechanical systems.

2022 *Gravitational Wave Detection*, *Computational Physics Research*.
Confirmed gravitational waves using the NSF–LIGO open dataset.

Programming Languages and Compilers

2016-2017 *Compiler for Actor-based Programming*, *Undergraduate Research*.
Researched a compiler for the P language, an actor-based concurrent programming language. Developed a model checker that uses binary decision diagrams and symbolic execution to verify program correctness.

2016-2017 *Optimizing Singular Value Decomposition on GPU*, *Undergraduate Research*.
Developed an efficient algorithm for batch rolling-window SVD on GPU.

2014 - 2015 *Java-to-Objective-C Transpiler*, *Research Internship at Berkeley Startup*.
Researched dead code elimination in a Java-to-Objective-C transpiler, reducing deployed app size by over 50%.

Teaching and Mentoring

2025 *Research Mentor*, *Basis Research Institute*.
Mentored a research trainee at Basis.

2022 *Teaching Assistant*, *MIT*.
6.4110/16.420: Representation, Inference, and Reasoning in AI

2022 *UROP Mentor*, *MIT*.
Mentored two MIT students on affine probabilistic graphical models research

2015-2017 *Undergraduate Student Instructor*, *UC Berkeley*.
Programming Languages and Compilers, Discrete Mathematics and Probability Theory

Service

Program Committee, *AAAI*.

Reviewer, *POPL, PLDI, NeurIPS, IJCAI, TOPLAS*.

Artifact Evaluation Committee, *OOPSLA, PLDI*.

Manuscripts in Preparation

[1] Yichao Liang, Thanh Dat Nguyen, Cambridge Yang, Tianyang Li, Joshua B. Tenenbaum, Carl Edward Rasmussen, Adrian Weller, Zenna Tavares, Tom Silver, and Kevin Ellis. ExoPredicator: Learning Abstract Models of Dynamic Worlds for Robot Planning. In: arXiv eprint. 2025. **In Review**

[2] Archana Warrier, Thanh Dat Nguyen, Michelangelo Naim, Moksh Jain, Yichao Liang, Karen Schroeder, Cambridge Yang, Joshua B. Tenenbaum, Sebastian Josef Vollmer, Kevin Ellis, and Zenna Tavares. Benchmarking World-model Learning. In: ICML Workshop. 2025. **In Review**

[3] Cambridge Yang and Dongdong Li. Taylor-mode Automatic Differentiation for Efficient Maximum Likelihood Estimation in High-dimensional Copulas. In: (2025). **Draft**

Publications

[4] Cambridge Yang, Michael L Littman, and Michael Carbin. Computably Continuous Reinforcement Learning Objectives are PAC-learnable. In: AAAI. 2023. **Long Oral**

[5] Cambridge Yang, Michael L Littman, and Michael Carbin. On the (In)Tractability of Reinforcement Learning for LTL Objectives. In: IJCAI. 2022. **Long Oral**

[6] Cambridge Yang, Eric Atkinson, and Michael Carbin. Simplifying Dependent Reductions in the Polyhedral Model. In: POPL. 2021. **Talk**

[7] Charith Mendis, Cambridge Yang, Yewen Pu, Saman Amarasinghe, and Michael Carbin. Compiler Auto-vectorization with Imitation Learning. In: NeurIPS. 2019.

[8] Eric Atkinson, Cambridge Yang, and Michael Carbin. Verifying Handcoded Probabilistic Inference Procedures. In: arXiv eprint. 2018.

[9] Ankush Desai, Indranil Saha, Jianqiao Yang, Shaz Qadeer, and Sanjit A Seshia. Drona: A Framework for Safe Distributed Mobile Robotics. In: ICCPS. 2017.

Skills

Languages.
Python, C++, Haskell, Java, JavaScript, C#, Mathematica

ML Frameworks.
JAX, TensorFlow, PyTorch, NumPy

Other

Programmer *Please visit my GitHub page for more details.*.
I code for fun. Among many projects, here are some fun highlights:
- AlphaZero-style Chinese dark chess AI.
- Software that uses computer vision to help me complete a 10000-piece jigsaw puzzle.
- Slack plugin for displaying LaTeX equations.

I also regularly contribute to various open-source projects. Some samples:
- JAX implementation of Taylor-mode automatic differentiation (jet).
- Happy-Hare, a 3D printer firmware for multi-material printing.

Maker .
I build 3D printers and use them to print for various projects, including a complete indoor hydroponics plant growing system in my apartment. I wrote software that fully automates the growing system. It has been producing fruits and veggies for my home for a few years.

Member of HKN .
Eta Kappa Nu, the honor society of IEEE