Brandon Amos

brandon.amos.cs@gmail.com • bamos.github.io • bamos

Last updated on October 27, 2026

Education

Ph.D. in Computer Science, *Carnegie Mellon University* (0.00/0.00) 2014–2019
Thesis: *Differentiable Optimization-Based Modeling for Machine Learning*
Advisor: J. Zico Kolter

B.S. in Computer Science, *Virginia Tech* (3.99/4.00) 2011–2014

Positions

Member of Technical Staff, *Reflection AI*, New York City 2025–
Research Scientist in FAIR, *Meta Superintelligence Labs*, New York City 2019–2025
Visiting Lecturer of Computer Science, *Cornell Tech*, New York City 2024
Research Assistant, *Carnegie Mellon University* (with J. Zico Kolter on ML and optimization) 2016–2019
Research Intern, *Intel Labs*, Santa Clara (with Vladlen Koltun on computer vision) 2018
Research Intern, *Google DeepMind*, London (with Nando de Freitas and Misha Denil on RL) 2017
Research Assistant, *Carnegie Mellon University* (with Mahadev Satyanarayanan on mobile systems) 2014–2016
Research Intern, *Adobe Research*, San Jose (with David Tompkins on distributed systems) 2014
Research Assistant, *Virginia Tech* (with Layne Watson and David Easterling on optimization) 2013–2014
Research Assistant, *Virginia Tech* (with Jules White and Hamilton Turner on mobile systems) 2012–2014
Research Assistant, *Virginia Tech* (with Binoy Ravindran and Alastair Murray on compilers) 2012–2014
Software Intern, *Snowplow* (Scala development) 2013–2014
Software Intern, *Qualcomm*, San Diego (Python and C++ development) 2013
Software Intern, *Phoenix Integration*, Virginia (C++, C#, and Java development) 2012
Network Administrator Intern, *Sunapsys*, Virginia 2011

Honors & Awards

INFORMS Computing Society Prize, Honorable Mention 2026
AISTATS Best Reviewer 2025
Outstanding Paper Award at the ICML Theoretical Foundations Workshop 2024
NeurIPS Top Reviewer 2022
ICML Outstanding Reviewer 2022
ICLR Outstanding Reviewer 2019
Best Paper Award at ACM MMSys 2017
NSF Graduate Research Fellowship 2016–2019
Nine undergraduate scholarships 2011–2014
Roanoke County Public Schools Engineering, Salem–Roanoke County Chamber of Commerce, Papa John's, Scottish Rite of Freemasonry, VT Intelligence Community Center for Academic Excellence, VT Pamplin Leader, VT Benjamin F. Bock, VT Gay B. Shober, VT I. Luck Gravett

Publications [Google Scholar: 15k+ citations and an h-index of 47]

Selected publications I am a primary author on are highlighted.

2026

1. *Cultivating Pluralism In Algorithmic Monoculture: The Community Alignment Dataset*
   Lily H Zhang, Smitha Milli, Karen Long Jusko, Jonathan Smith, **Brandon Amos**, Wassim Bouaziz, Jack Kussman, Manon Revel, Lisa Titus, Bhaktipriya Radharapu, Jane Yu, Vidya Sarma, Kristopher Rose, and Maximilian Nickel
   ICLR 2026

2. *The Surprising Difficulty of Search in Model-Based Reinforcement Learning*
Wei-Di Chang, Mikael Henaff, **Brandon Amos**, Gregory Dudek, and Scott Fujimoto
ICML 2026

3. *Safety Alignment of LMs via Non-cooperative Games* code
Anselm Paulus, Ilia Kulikov, **Brandon Amos**, Rémi Munos, Ivan Evtimov, Kamalika Chaudhuri, and Arman Zharmagambetov
ICML 2026

4. *A Fully First-Order Layer for Differentiable Optimization* code
Zihao Zhao, Kai-Chia Mo, Shing-Hei Ho, **Brandon Amos**, and Kai Wang
ICML 2026

5. *Score Function Gradient Estimation to Widen the Applicability of Decision-Focused Learning*
Mattia Silvestri, Senne Berden, Jayanta Mandi, Ali İrfan Mahmutoğulları, **Brandon Amos**, Tias Guns, and Michele Lombardi
JAIR 2026

6. *Social Choice Foundations for Simulation-Augmented Generation*
Sonja Kraiczy, Smitha Milli, Ratip Emin Berker, Avinandan Bose, **Brandon Amos**, Jamelle Watson-Daniels, Maximilian Nickel, Edith Elkind, and Ariel D. Procaccia
NeurIPS 2026 2026

2025

7. *Online Intrinsic Rewards for Decision Making Agents from Large Language Model Feedback* code
Qinqing Zheng, Mikael Henaff, Amy Zhang, Aditya Grover, and **Brandon Amos**
RLC 2025

8. *AdvPrompter: Fast Adaptive Adversarial Prompting for LLMs* code slides
Anselm Paulus*, Arman Zharmagambetov*, Chuan Guo, **Brandon Amos**†, and Yuandong Tian†
ICML 2025

9. *Wasserstein Flow Matching: Generative modeling over families of distributions* code
Doron Haviv, Aram-Alexandre Pooladian, Dana Pe'er, and **Brandon Amos**
ICML 2025

10. *Adjoint Sampling: Highly Scalable Diffusion Samplers via Adjoint Matching* code
Aaron Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Brandon Wood, Daniel Levine, Bin Hu, **Brandon Amos**, Brian Karrer, Xiang Fu, Guan-Horng Liu, and Ricky T. Q. Chen
ICML 2025

11. *Exact Byte-Level Probabilities from Tokenized Language Models for FIM-Tasks and Model Ensembles* code
Buu Phan, **Brandon Amos**, Itai Gat, Marton Havasi, Matthew Muckley, and Karen Ullrich
ICLR 2025

12. *Meta Flow Matching: Integrating Vector Fields on the Wasserstein Manifold* code
Lazar Atanackovic, Xi Zhang, **Brandon Amos**, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov
ICLR 2025

13. *AdvPrefix: An Objective for Nuanced LLM Jailbreaks*
Sicheng Zhu, **Brandon Amos**, Yuandong Tian, Chuan Guo, and Ivan Evtimov
NeurIPS 2025

14. *AlgoTune: Can Language Models Speed Up General-Purpose Numerical Programs?* code
Ori Press, **Brandon Amos**, Haoyu Zhao, Yikai Wu, Samuel K. Ainsworth, Dominik Krupke, Patrick Kidger, Touqir Sajed, Bartolomeo Stellato, Jisun Park, Nathanael Bosch, Eli Meril, Albert Steppi, Arman Zharmagambetov, Fangzhao Zhang, David Pérez-Piñeiro, Alberto Mercurio, Ni Zhan, Talor Abramovich, Kilian Lieret, Hanlin Zhang, Shirley Huang, Matthias Bethge, and Ofir Press
NeurIPS Datasets and Benchmarks Track 2025

15. *BaNEL: Exploration Posteriors for Generative Modeling Using Only Negative Rewards*
Sangyun Lee, **Brandon Amos**, and Giulia Fanti
arXiv 2025

2024

16. *Neural Optimal Transport with Lagrangian Costs* code
Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky T. Q. Chen, and **Brandon Amos**
UAI 2024

17. *Learning to Warm-Start Fixed-Point Optimization Algorithms* code
Rajiv Sambharya, Georgina Hall, **Brandon Amos**, and Bartolomeo Stellato
JMLR 2024
INFORMS Computing Society Prize, Honorable Mention

18. *Unlocking Tokens as Data Points for Generalization Bounds on Larger Language Models* code
Sanae Lotfi, Yilun Kuang, Marc Anton Finzi, **Brandon Amos**, Micah Goldblum, and Andrew Gordon Wilson
NeurIPS 2024
Outstanding Paper Award at the ICML Theoretical Foundations Workshop

19. *Stochastic Optimal Control Matching* code
Carles Domingo-Enrich, Jiequn Han, **Brandon Amos**, Joan Bruna, and Ricky T. Q. Chen
NeurIPS 2024

20. *To the Globe (TTG): Towards Language-Driven Guaranteed Travel Planning*
Da JU, Song Jiang, Andrew Cohen, Aaron Foss, Sasha Mitts, Arman Zharmagambetov, **Brandon Amos**, Xian Li, Justine T Kao, Maryam Fazel-Zarandi, and Yuandong Tian
EMNLP Demo 2024

2023

21. *Tutorial on amortized optimization* code
Brandon Amos
Foundations and Trends in Machine Learning 2023

22. *On amortizing convex conjugates for optimal transport* code
Brandon Amos
ICLR 2023

23. *End-to-End Learning to Warm-Start for Real-Time Quadratic Optimization* code
Rajiv Sambharya, Georgina Hall, **Brandon Amos**, and Bartolomeo Stellato
L4DC 2023

24. *Meta Optimal Transport* code
Brandon Amos, Samuel Cohen, Giulia Luise, and Ievgen Redko
ICML 2023

25. *Multisample Flow Matching: Straightening Flows with Minibatch Couplings*
Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, **Brandon Amos**, Yaron Lipman, and Ricky T. Q. Chen
ICML 2023

26. *Semi-Supervised Offline Reinforcement Learning with Action-Free Trajectories* code
Qinqing Zheng, Mikael Henaff, **Brandon Amos**, and Aditya Grover
ICML 2023

27. *TaskMet: Task-Driven Metric Learning for Model Learning* code
Dishank Bansal, Ricky T. Q. Chen, Mustafa Mukadam, and **Brandon Amos**
NeurIPS 2023

28. *Landscape Surrogate: Learning Decision Losses for Mathematical Optimization Under Partial Information* code
Arman Zharmagambetov, **Brandon Amos**, Aaron Ferber, Taoan Huang, Bistra Dilkina, and Yuandong Tian
NeurIPS 2023

29. *Koopman Constrained Policy Optimization: A Koopman operator theoretic method for differentiable optimal control in robotics* code
Matthew Retchin, **Brandon Amos**, Steven Brunton, and Shuran Song
ICML Differentiable Almost Everything Workshop 2023

2022

30. *Cross-Domain Imitation Learning via Optimal Transport* code
Arnaud Fickinger, Samuel Cohen, Stuart Russell, and **Brandon Amos**
ICLR 2022

31. *Matching Normalizing Flows and Probability Paths on Manifolds*
Heli Ben-Hamu*, Samuel Cohen*, Joey Bose, **Brandon Amos**, Aditya Grover, Maximilian Nickel, Ricky T. Q. Chen, and Yaron Lipman
ICML 2022

32. *Semi-Discrete Normalizing Flows through Differentiable Tessellation*
Ricky T. Q. Chen, **Brandon Amos**, and Maximilian Nickel
NeurIPS 2022

33. *Theseus: A Library for Differentiable Nonlinear Optimization* code
Luis Pineda, Taosha Fan, Maurizio Monge, Shobha Venkataraman, Paloma Sodhi, Ricky Chen, Joseph Ortiz, Daniel DeTone, Austin Wang, Stuart Anderson, Jing Dong, **Brandon Amos**, and Mustafa Mukadam
NeurIPS 2022

34. *Nocturne: a driving benchmark for multi-agent learning* code
Eugene Vinitsky, Nathan Lichtlé, Xiaomeng Yang, **Brandon Amos**, and Jakob Foerster
NeurIPS Datasets and Benchmarks Track 2022

2021

35. *On the model-based stochastic value gradient for continuous reinforcement learning* code slides
Brandon Amos, Samuel Stanton, Denis Yarats, and Andrew Gordon Wilson
L4DC 2021 (Oral)

36. *Riemannian Convex Potential Maps* code slides
Samuel Cohen*, **Brandon Amos***, and Yaron Lipman
ICML 2021

37. *CombOptNet: Fit the Right NP-Hard Problem by Learning Integer Programming Constraints* code
Anselm Paulus, Michal Rolínek, Vít Musil, **Brandon Amos**, and Georg Martius
ICML 2021

38. *Scalable Online Planning via Reinforcement Learning Fine-Tuning*
Arnaud Fickinger, Hengyuan Hu, **Brandon Amos**, Stuart Russell, and Noam Brown
NeurIPS 2021

39. *Aligning Time Series on Incomparable Spaces* code slides
Samuel Cohen, Giulia Luise, Alexander Terenin, **Brandon Amos**, and Marc Peter Deisenroth
AISTATS 2021

40. *Learning Neural Event Functions for Ordinary Differential Equations* code
Ricky T. Q. Chen, **Brandon Amos**, and Maximilian Nickel
ICLR 2021

41. *Neural Spatio-Temporal Point Processes* code
Ricky T. Q. Chen, **Brandon Amos**, and Maximilian Nickel
ICLR 2021

42. *Improving Sample Efficiency in Model-Free Reinforcement Learning from Images* code
Denis Yarats, Amy Zhang, Ilya Kostrikov, **Brandon Amos**, Joelle Pineau, and Rob Fergus
AAAI 2021

43. *Neural Fixed-Point Acceleration for Convex Optimization* code
Shobha Venkataraman* and **Brandon Amos***
ICML AutoML Workshop 2021

44. *Sliced Multi-Marginal Optimal Transport*
Samuel Cohen, Alexander Terenin, Yannik Pitcan, **Brandon Amos**, Marc Peter Deisenroth, and K S Sesh Kumar
NeurIPS OTML Workshop 2021

45. *Input Convex Gradient Networks*
Jack Richter-Powell, Jonathan Lorraine, and **Brandon Amos**
NeurIPS OTML Workshop 2021

46. *Imitation Learning from Pixel Observations for Continuous Control*
Samuel Cohen, **Brandon Amos**, Marc Peter Deisenroth, Mikael Henaff, Eugene Vinitsky, and Denis Yarats
NeurIPS DeepRL Workshop 2021

47. *MBRL-Lib: A Modular Library for Model-based Reinforcement Learning* code
Luis Pineda, **Brandon Amos**, Amy Zhang, Nathan Lambert, and Roberto Calandra
arXiv 2021

2020

48. *The Differentiable Cross-Entropy Method* code slides
Brandon Amos and Denis Yarats
ICML 2020

49. *Objective Mismatch in Model-based Reinforcement Learning*
Nathan Lambert, **Brandon Amos**, Omry Yadan, and Roberto Calandra
L4DC 2020

50. *QNSTOP: Quasi-Newton Algorithm for Stochastic Optimization* code
Brandon Amos, David Easterling, Layne T. Watson, William Thacker, Brent Castle, and Michael Trosset
ACM TOMS 2020

51. *Neural Potts Model*
Tom Sercu, Robert Verkuil, Joshua Meier, **Brandon Amos**, Zeming Lin, Caroline Chen, Jason Liu, Yann LeCun, and Alexander Rives
MLCB 2020

52. *Deep Riemannian Manifold Learning*
Aaron Lou, Maximilian Nickel, and **Brandon Amos**
NeurIPS Geo4dl Workshop 2020

2019

53. *Differentiable Optimization-Based Modeling for Machine Learning* code
Brandon Amos
Ph.D. Thesis 2019

54. *Differentiable Convex Optimization Layers* code
Akshay Agrawal*, **Brandon Amos***, Shane Barratt*, Stephen Boyd*, Steven Diamond*, and J. Zico Kolter*
NeurIPS 2019

55. *The Limited Multi-Label Projection Layer* code
Brandon Amos, Vladlen Koltun, and J. Zico Kolter
arXiv 2019

56. *Generalized Inner Loop Meta-Learning* code
Edward Grefenstette, **Brandon Amos**, Denis Yarats, Phu Mon Htut, Artem Molchanov, Franziska Meier, Douwe Kiela, Kyunghyun Cho, and Soumith Chintala
arXiv 2019

2018

57. *Learning Awareness Models*
Brandon Amos, Laurent Dinh, Serkan Cabi, Thomas Rothörl, Sergio Gómez Colmenarejo, Alistair Muldal, Tom Erez, Yuval Tassa, Nando de Freitas, and Misha Denil
ICLR 2018

58. *Differentiable MPC for End-to-end Planning and Control* code
Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J. Zico Kolter
NeurIPS 2018

59. *Depth-Limited Solving for Imperfect-Information Games*
Noam Brown, Tuomas Sandholm, and **Brandon Amos**
NeurIPS 2018

60. *Enabling Live Video Analytics with a Scalable and Privacy-Aware Framework*
Junjue Wang, **Brandon Amos**, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan
ACM TOMM 2018

2017

61. *OptNet: Differentiable Optimization as a Layer in Neural Networks* code slides
Brandon Amos and J. Zico Kolter
ICML 2017

62. *Input Convex Neural Networks* code slides
Brandon Amos, Lei Xu, and J. Zico Kolter
ICML 2017

63. *Task-based End-to-end Model Learning* code
Priya L. Donti, **Brandon Amos**, and J. Zico Kolter
NeurIPS 2017

64. *Quasi-Newton Stochastic Optimization Algorithm for Parameter Estimation of a Stochastic Model of the Budding Yeast Cell Cycle*
Minghan Chen, **Brandon Amos**, Layne T. Watson, John Tyson, Yang Cao, Cliff Shaffer, Michael Trosset, Cihan Oguz, and Gisella Kakoti
IEEE/ACM TCBB 2017

65. *You can teach elephants to dance: agile VM handoff for edge computing*
Kiryong Ha, Yoshihisa Abe, Thomas Eiszler, Zhuo Chen, Wenlu Hu, **Brandon Amos**, Rohit Upadhyaya, Padmanabhan Pillai, and Mahadev Satyanarayanan
SEC 2017

66. *An Empirical Study of Latency in an Emerging Class of Edge Computing Applications for Wearable Cognitive Assistance*
Zhuo Chen, Wenlu Hu, Junjue Wang, Siyan Zhao, **Brandon Amos**, Guanhang Wu, Kiryong Ha, Khalid Elgazzar, Padmanabhan Pillai, Roberta Klatzky, Daniel Siewiorek, and Mahadev Satyanarayanan
SEC 2017

67. *A Scalable and Privacy-Aware IoT Service for Live Video Analytics* code
Junjue Wang, **Brandon Amos**, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan
ACM MMSys 2017
Best Paper Award

2016

68. *OpenFace: A general-purpose face recognition library with mobile applications* code
Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan
CMU 2016

69. *Image Completion with Deep Learning in TensorFlow* code
Brandon Amos
Blog 2016

70. *Collapsed Variational Inference for Sum-Product Networks*
Han Zhao, Tameem Adel, Geoff Gordon, and **Brandon Amos**
ICML 2016

71. *Quantifying the impact of edge computing on mobile applications*
Wenlu Hu, Ying Gao, Kiryong Ha, Junjue Wang, **Brandon Amos**, Zhuo Chen, Padmanabhan Pillai, and Mahadev Satyanarayanan
ACM SIGOPS 2016

72. *Privacy mediators: helping IoT cross the chasm*
Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and **Brandon Amos**
HotMobile 2016

2015 and earlier

73. *Edge Analytics in the Internet of Things*
Mahadev Satyanarayanan, Pieter Simoens, Yu Xiao, Padmanabhan Pillai, Zhuo Chen, Kiryong Ha, Wenlu Hu, and **Brandon Amos**
IEEE Pervasive Computing 2015

74. *Bad Parts: Are Our Manufacturing Systems at Risk of Silent Cyberattacks?*
Hamilton Turner, Jules White, Jaime A. Camelio, Christopher Williams, **Brandon Amos**, and Robert Parker
IEEE Security & Privacy 2015

75. *Early Implementation Experience with Wearable Cognitive Assistance Applications*
Zhuo Chen, Lu Jiang, Wenlu Hu, Kiryong Ha, **Brandon Amos**, Padmanabhan Pillai, Alex Hauptmann, and Mahadev Satyanarayanan
WearSys 2015

76. *The Case for Offload Shaping*
Wenlu Hu, **Brandon Amos**, Zhuo Chen, Kiryong Ha, Wolfgang Richter, Padmanabhan Pillai, Benjamin Gilbert, Jan Harkes, and Mahadev Satyanarayanan
HotMobile 2015

77. *Are Cloudlets Necessary?*
Ying Gao, Wenlu Hu, Kiryong Ha, **Brandon Amos**, Padmanabhan Pillai, and Mahadev Satyanarayanan
CMU 2015

78. *Adaptive VM handoff across cloudlets*
Kiryong Ha, Yoshihisa Abe, Zhuo Chen, Wenlu Hu, **Brandon Amos**, Padmanabhan Pillai, and Mahadev Satyanarayanan
CMU 2015

79. *Global Parameter Estimation for a Eukaryotic Cell Cycle Model in Systems Biology*
Tricity Andrew, **Brandon Amos**, David Easterling, Cihan Oguz, William Baumann, John Tyson, and Layne T. Watson
SummerSim 2014

80. *Applying machine learning classifiers to dynamic Android malware detection at scale* code
Brandon Amos, Hamilton Turner, and Jules White
IWCMC 2013

Open Source Repositories

40k+ GitHub stars across all repositories.

1. oripress/AlgoTune | ★119 2025
2. facebookresearch/adjoint_sampling | ★139 2025
3. facebookresearch/oni | ★45 | *Online LLM intrinsic rewards for NetHack* 2025

4. facebookresearch/advprompter | ★183 | *Fast Adaptive Adversarial Prompting for LLMs* 2024
5. facebookresearch/lagrangian-ot | ★62 2024
6. lazaratan/meta-flow-matching | ★77 2024
7. facebookresearch/soc-matching | ★43 | *Stochastic Optimal Control Matching* 2024
8. kuleshov/cornell-cs5785-2024-applied-ml | ★538 | *Slides for our applied ML course* 2024
9. facebookresearch/amortized-optimization-tutorial | ★267 2023
10. facebookresearch/taskmet | ★21 | *Task-Driven Metric Learning for Model Learning* 2023
11. facebookresearch/w2ot | ★46 | *Wasserstein-2 optimal transport* 2023
12. facebookresearch/LANCER | ★36 | *Landscape Surrogate Learning Decision Losses* 2023
13. facebookresearch/theseus | ★2062 | *Differentiable non-linear optimization library* 2022
14. facebookresearch/meta-ot | ★108 2022
15. bamos/presentations | ★140 2022
16. facebookresearch/gwil | ★27 | *Gromov-Wasserstein Cross Domain Imitation Learning* 2022
17. facebookresearch/nocturne | ★299 | *A partially-observable multi-agent driving simulator* 2022
18. facebookresearch/rcpm | ★68 | *Riemannian Convex Potential Maps* 2021
19. facebookresearch/svg | ★58 | *Model-based stochastic value gradient* 2021
20. facebookresearch/mbrl-lib | ★1062 | *Model-based reinforcement learning library* 2021
21. martius-lab/CombOptNet | ★73 2021
22. samcohen16/Aligning-Time-Series | ★51 | *Aligning time series on incomparable spaces* 2021
23. facebookresearch/neural_stpp | ★108 | *Neural Spatio-Temporal Point Processes* 2021
24. facebookresearch/neural-scs | ★30 | *Neural Fixed-Point Acceleration for SCS* 2021
25. rtqichen/torchdiffeq | ★6492 | *PyTorch Differentiable ODE Solvers (differentiable event handling)* 2021
26. facebookresearch/dcem | ★124 | *The Differentiable Cross-Entropy Method* 2020
27. facebookresearch/higher | ★1624 | *PyTorch higher-order gradient and optimization library* 2019
28. bamos/thesis | ★351 | *Ph.D. Thesis LaTeX source code* 2019
29. cvxgrp/cvxpylayers | ★2146 | *Differentiable Convex Optimization Layers* 2019
30. locuslab/lml | ★59 | *The Limited Multi-Label Projection Layer* 2019
31. locuslab/mpc.pytorch | ★1080 | *Differentiable PyTorch Model Predictive Control library* 2018
32. locuslab/differentiable-mpc | ★341 2018
33. locuslab/icnn | ★318 | *Input Convex Neural Network experiments* 2017
34. locuslab/optnet | ★594 2017
35. locuslab/qpth | ★803 | *Differentiable PyTorch QP solver* 2017
36. bamos/densenet.pytorch | ★840 2017
37. bamos/block | ★307 | *Intelligent block matrix constructions* 2017
38. bamos/setGPU | ★112 | *Automatically use the least-loaded GPU* 2017
39. bamos/dcgan-completion.tensorflow | ★1317 | *Image completion with GANs* 2016
40. cmusatyalab/openface | ★15446 | *Face recognition with deep neural networks* 2015
41. bamos/girl | ★68 | *GitHub README link checker* 2015
42. bamos/conference-tracker | ★71 2015
43. vtopt/qnstop | ★11 | *Fortran quasi-Newton stochastic optimization library* 2014
44. bamos/snowglobe | ★28 | *Haskell-driven, self-hosted web analytics* 2014
45. bamos/zsh-history-analysis | ★242 2014
46. bamos/beamer-snippets | ★112 2014
47. bamos/latex-templates | ★369 2013
48. cparse/cparse | ★368 | *C++ expression parser using Dijkstra's shunting-yard algorithm* 2013
49. bamos/cv | ★420 | *Source for this CV: creates LaTeX/Markdown from YAML/BibTeX* 2013
50. bamos/parsec-benchmark | ★134 | *PARSEC benchmark support for Arch Linux* 2013
51. bamos/python-scripts | ★194 2013
52. bamos/reading-list | ★184 | *YAML reading list and notes system* 2013
53. bamos/dotfiles | ★239 | ♥ *Linux, xmonad, emacs, vim, zsh, tmux* 2012

Invited Talks

Slides for my major presentations are available here under a CC-BY license.

1. *On prompt optimization and coding agents* • Columbia NLP Seminar 2026
2. *On prompt optimization and coding agents* • Cornell Tech ML Reading Group 2026
3. *On prompt optimization and coding agents* • Tel Aviv University 2026
4. *On prompt optimization and coding agents* • Princeton PLI Seminar 2026
5. *On prompt optimization and coding agents* • Netflix 2026
6. *On prompt optimization and coding agents* • NeurIPS ScaleOPT Workshop 2025
7. *Language model alignment and tool use* • Cornell Tech CS5785 Guest Lecture 2025
8. *On transport, flows, and physics* • The Physics of John Hopfield Workshop at Princeton University 2025
9. *Demystifying language models* • ICT Dumaguete Seminar 2025
10. *On prompt optimization and coding agents* • Flatiron ML Seminar 2025
11. *AdvPrompter: Fast Adaptive Adversarial Prompting for LLMs* • USC 2025
12. *On amortized optimization for RL, Bayesian optimization, and biology* • ai4b.io 2025
13. *Transport and flows between distributions over distributions* • Columbia University 2025
14. *Transport and flows between distributions over distributions* • Genesis Therapeutics 2024
15. *Transport and flows between distributions over distributions* • UT Austin 2024
16. *On LLM prompt optimization and amortization* • Dagstuhl Seminar on ML for CO 2024
17. *Amortized optimization for optimal transport and LLM attacks* • ISMP 2024
18. *Differentiable optimization for robotics* • RSS Optimization for Robotics Workshop 2024
19. *Amortized optimization-based reasoning for AI* • University of Amsterdam 2024
20. *End-to-end learning geometries for graphs, dynamical systems, and regression* • LoG New York 2024
21. *Amortized optimization for optimal transport* • NeurIPS Optimal Transport and ML Workshop 2023
22. *On optimal control and machine learning* • ICML Control and Dynamical Systems Workshop 2023
23. *Tutorial on amortized optimization* • Brown University 2023
24. *Learning with differentiable and amortized optimization* • NYU AI Seminar 2023
25. *Learning with differentiable and amortized optimization* • Vanderbilt ML Seminar 2023
26. *Learning with differentiable and amortized optimization* • Microsoft Research 2022
27. *Amortized optimization for computing optimal transport maps* • Flatiron Workshop 2022
28. *Learning with differentiable and amortized optimization* • Cornell AI Seminar 2022
29. *Learning with differentiable and amortized optimization* • Cornell Tech Seminar 2022
30. *Learning with differentiable and amortized optimization* • Argonne National Laboratory 2022
31. *Theseus: A library for differentiable nonlinear optimization* • NYU 2022
32. *Theseus: A library for differentiable nonlinear optimization* • University of Zurich 2022
33. *Differentiable optimization-based modeling for machine learning* • Colorado Mines AMS Colloquium 2022
34. *Differentiable optimization* • IJCAI Tutorial 2022
35. *Differentiable optimization for control and RL* • ICML Workshop on Decision Awareness in RL 2022
36. *Differentiable optimization-based modeling for machine learning* • CPAIOR Master Class 2022
37. *Tutorial on amortized optimization* • ICCOPT 2022
38. *Differentiable optimization for control and RL* • Gridmatic 2022
39. *Learning for control with differentiable optimization and ODEs* • Columbia University 2021
40. *Differentiable optimization-based modeling for machine learning* • IBM Research 2021
41. *Differentiable optimization for control* • Max Planck Institute (Tübingen) 2020
42. *Differentiable optimization-based modeling for machine learning* • Mila Seminar 2020
43. *Deep Declarative Networks* • ECCV Tutorial 2020
44. *On differentiable optimization for control and vision* • CVPR Deep Declarative Networks Workshop 2020
45. *Differentiable optimization-based modeling for machine learning* • Caltech CS159 Guest Lecture 2020
46. *Unrolled optimization for learning deep energy models* • SIAM MDS Minisymposium 2020
47. *Differentiable optimization-based modeling for machine learning* • NYU CILVR Seminar 2019
48. *Differentiable optimization-based modeling for machine learning* • INFORMS 2019

49. *Differentiable optimization-based modeling for machine learning* • Facebook AI Research 2019
50. *Differentiable optimization-based modeling for machine learning* • ISMP 2018
51. *Differentiable optimization-based modeling for machine learning* • Google Brain 2018
52. *Differentiable optimization-based modeling for machine learning* • Bosch Center for AI 2018
53. *Differentiable optimization-based modeling for machine learning* • Waymo Research 2018
54. *Differentiable optimization-based modeling for machine learning* • Tesla AI 2018
55. *Differentiable optimization-based modeling for machine learning* • NVIDIA Robotics 2018
56. *Differentiable optimization-based modeling for machine learning* • Salesforce Research 2018
57. *Differentiable optimization-based modeling for machine learning* • OpenAI 2018
58. *Differentiable optimization-based modeling for machine learning* • NNAISENSE 2018
59. *Differentiable optimization and control* • UC Berkeley 2018

Interns and Students

Kathy Zhou (visiting FAIR from NYU) 2025
Ulyana Piterbarg (visiting FAIR from NYU, now at Google DeepMind) 2025
Ollie Liu (visiting FAIR from USC, now at Reflection) 2025
Doron Haviv (MSKCC PhD committee, now at Genentech) 2025
Aaron Havens (visiting FAIR from UIUC, now postdoc at FAIR) 2024–2025
Aram-Alexandre Pooladian (visiting FAIR from NYU, now at Yale) 2022–2024
Carles Domingo-Enrich (visiting FAIR from NYU, now at MSR) 2022–2024
Anselm Paulus (visiting FAIR from Max Planck Institute, Tübingen) 2023–2024
Matthew Retchin (Columbia MS thesis committee, now at Harvard) 2023
Sanae Lotfi (visiting FAIR from NYU, now scientist at FAIR) 2022–2023
Dishank Bansal (AI resident at FAIR, now at Google DeepMind) 2022–2023
Arnaud Fickinger (visiting FAIR from Berkeley) 2021–2022
Aaron Lou (visiting FAIR from Cornell and Stanford, now scientist at OpenAI) 2020–2022
Eugene Vinitsky (visiting FAIR from Berkeley, now professor at NYU) 2021–2022
Samuel Cohen (visiting FAIR from UCL, now CEO at FairGen) 2021–2022
Ricky Chen (visiting FAIR from Toronto, now scientist at FAIR) 2020
Paul Liang (visiting FAIR from CMU, now professor at MIT) 2020
Phillip Wang (at CMU, now CEO at Gather) 2018

Professional Activities

NeurIPS Learning Meets Combinatorial Optimization Workshop Organizer 2020
CVPR Deep Declarative Networks Workshop Organizer 2020
ECCV Deep Declarative Networks Tutorial Organizer 2020
CMU CSD MS Admissions 2014–2015

Area Chair

Artificial Intelligence and Statistics (AISTATS)
Association for the Advancement of Artificial Intelligence (AAAI)
International Conference on Learning Representations (ICLR)
International Conference on Machine Learning (ICML)
Neural Information Processing Systems (NeurIPS)
Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks

Reviewer

American Controls Conference (ACC)
Artificial Intelligence and Statistics (AISTATS)
Association for the Advancement of Artificial Intelligence (AAAI)
IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
IEEE Conference on Decision and Control (CDC)
IEEE Control Systems Letters (L-CSS)
IEEE International Conference on Computer Vision (ICCV)
IEEE International Conference on Intelligent Robots and Systems (IROS)
IEEE International Conference on Robotics and Automation (ICRA)
International Conference on Learning Representations (ICLR)
International Conference on Learning Representations (ICLR) Blog Posts
International Conference on Machine Learning (ICML)
International Conference on Machine Learning (ICML) SODS Workshop
International Conference on the Constraint Programming, AI, and Operations Research (CPAIOR)
Journal of Machine Learning Research (JMLR)
Learning for Dynamics and Control (L4DC)
Mathematical Programming Computation (MPC)
Neural Information Processing Systems (NeurIPS)
Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track
Neural Information Processing Systems (NeurIPS) Deep RL Workshop
Neural Information Processing Systems (NeurIPS) DiffCVGP Workshop
Neural Information Processing Systems (NeurIPS) OPT Workshop
Optimization Letters
Transactions on Machine Learning Research (TMLR)
Uncertainty in Artificial Intelligence (UAI)

Teaching

Applied Machine Learning, *Cornell Tech CS5785*, Co-instructor F2024
Graduate AI, *CMU 15-780*, TA S2017
Distributed Systems, *CMU 15-640*, TA S2016
Software Design and Data Structures, *VT CS2114*, TA S2013

Skills

Programming C, C++, Fortran, Haskell, Java, Lua, Make, Mathematica, Python, R, Scala
Frameworks JAX, NumPy, Pandas, PyTorch, SciPy, TensorFlow, Torch7
Toolbox Linux, emacs, vim, evil, org, mu4e, xmonad, git, tmux, zsh