

Yatong Bai

Deep Learning, AI Audio Generation, Adversarial Robustness, Optimization

yatong_bai (at) berkeley.edu

Website: bai-yt.github.io

LinkedIn: linkedin.com/in/yatong-bai

EDUCATION

Ph. D. and M. S., University of California, Berkeley

Aug 2020 – Aug 2025 (Expected)

- Master's degree conferred in May 2022. GPA: 4.00 / 4.00
- Areas: Deep learning (especially audio generative AI and adversarial robustness), Optimization, and Control.
- Advisor: Somayeh Sojoudi.
- Courses: Convex optimization and algorithms, Deep neural networks, Statistical learning theory, Deep reinforcement learning, Advanced control systems, Theoretical statistics.

B. S., Georgia Institute of Technology

Aug 2016 – Aug 2020

- Double major in Computer Engineering and Mechanical Engineering. GPA: 4.00 / 4.00
- Courses: Machine learning, Computer vision, Signals and systems, Embedded systems, Computer architecture.

GRADUATE-LEVEL WORK EXPERIENCE (For Berkeley experiences please see publications)

Microsoft Applied Science Group, Research Intern

Redmond, WA, May 2023 – Aug 2023

- Accelerate diffusion-model-based text-to-audio generation hundreds of times with minimal performance decrease via distillation.
- Innovatively combine the “consistency distillation” framework and “classifier-free guidance” to distill high-performance models.
- The distilled model needs only one neural net query (diffusion models need 400), enabling improving audio semantics by optimizing audio-space losses. We use the CLAP score loss as an example and use objective and subjective evaluation to show its effectiveness.
- Project website consistency-tta.github.io. Preprint paper on ArXiv (2309.10740).

Scale AI, Machine Learning Research Intern

San Francisco, CA, May 2022 – Dec 2022

- Researched proposing a dataset with 15 million image-caption pairs and processing its captions with various language models.
- Applied supervised and self-supervised image classification, object detection, image reconstruction, and generation methods (in PyTorch) to provide benchmarks on the dataset. Applied dimension reduction (UMAP) to visualize the embedding clustering.
- Used the above results to characterize the distribution shift of our data from existing datasets. Preprint paper on ArXiv (2401.04575).

PUBLICATIONS

Accelerating Diffusion-Based Text-to-Audio Generation with Consistency Distillation

Yatong Bai, T Dang, D Tran, K Koishida, and S Sojoudi. *Preprint*, 2023. arxiv.org/abs/2309.10740

- See the internship at Microsoft (listed above) for details.

Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off

Yatong Bai, BG Anderson, and S Sojoudi. In *IEEE Conference on Control Technology and Applications*, 2023. arxiv.org/abs/2311.15165

- We show that by mixing the output probabilities of an accurate (often non-robust) classifier and a robust classifier, the accuracy- (adversarial) robustness trade-off can be greatly alleviated, and certified robustness can be achieved. We show that the robust base classifier's prediction confidence difference between correct and incorrect examples is the main source of this improvement.

Improving the Accuracy-Robustness Trade-Off of Classifiers via Local Adaptive Smoothing

Yatong Bai, BG Anderson, A Kim, and S Sojoudi. *Preprint*, 2023. arxiv.org/abs/2301.12554

- Based on the work above, we further introduce a “mixing network” to adjust the mixing strengths for benign and attacked inputs differently, further improving the accuracy-robustness trade-off. On the CIFAR-10 and CIFAR-100 datasets, adaptive smoothing is the second most robust method listed on RobustBench, while noticeably improving the clean accuracy over all other works.
- Project code available at github.com/Bai-YT/AdaptiveSmoothing.

Efficient Global Optimization of Two-Layer ReLU Networks: Adversarial Training and Quadratic-Time Algorithms

Yatong Bai, T Gautam, and S Sojoudi. In *SIAM Journal on Mathematics of Data Science*, 2022. arxiv.org/abs/2201.01965

- 2021 INFORMS Data Mining Best Paper Competition (Student Track) Runner-up (2nd out of 48 papers).
- We develop efficient ADMM algorithms for the “convex training” formulation, which trains one-hidden-layer neural networks via convex optimization. We prove that the proposed algorithms polynomially improve the computational complexity.

Initial State Interventions for Deconfounded Imitation Learning

S Pfrommer, Yatong Bai, H Lee, and S Sojoudi. In *IEEE Conference on Decision and Control*, 2023. arxiv.org/abs/2307.15980

- Imitation learning agents suffer from causal confusion. We use a beta-VAE neural net to obtain disentangled latent representations underlying the observations, and use a statistical test to mask confounding latent variables so that the agent performs significantly better when the observations are confounded.

Practical Convex Formulation of Robust One-Hidden-Layer Neural Network Training

Yatong Bai, T Gautam, Y Gai, and S Sojoudi. In *American Control Conference*, 2022. arxiv.org/abs/2105.12237

- We leverage the duality theory and robust optimization techniques to develop efficient convex optimization formulations that train robust one-hidden-layer ReLU neural networks via adversarial training.
- Our method demonstrates improved adversarial robustness on common datasets, including CIFAR-10.

Let's Go Shopping (LGS) – Web-Scale Image-Text Dataset for Visual Concept Understanding

Yatong Bai, U Garg, A Shanker, H Zhang, S Parajuli, E Bas, I Filipovic, AN Chu, ED Fomitcheva, E Branson, A Kim, S Sojoudi, K Cho. *Preprint*, 2024. arxiv.org/abs/2401.04575

- See the internship at Scale AI (listed above) for details.

UNDERGRADUATE EXPERIENCE

Georgia Institute of Technology

Undergraduate Student Researcher

TINKER Group, RoboMed Group, Meaud Research Group, GT Off-road

Jan 2018 – Jan 2020

- Compile the SPEC 2017 computer architecture benchmark into ARM binary programs using GCC-ARM; Used the Gem5 computer architecture simulator (in C++) to convert the binary programs into debug trace files.
- Built Graphical User Interfaces (GUIs) for a cochlear dynamics simulator in MATLAB. The GUIs controlled simulations, logged and processed experiment data, and visualized the simulation results.

Senior design project: Avionics and test stand controller for a “Monocopter” aircraft

- Implemented the avionics system of a novel unmanned “Monocopter” and a PID-controlled testbed using C++. The avionics filtered noisy magnetometer readings to accurately recover aircraft heading and control the actuators. Also developed a Windows C# GUI.

Honda Aircraft Company, Intern

Greensboro, NC, May 2019 – Aug 2019

- Conducted dynamic simulations for flap linkages in MSC ADAMS, and evaluated the stress, deflection, and kinematics in CATIA via Finite Element Analyses (FEA).
- Defined the flap skew & asymmetry warning thresholds and designed a flap control logic in MATLAB.

Tesla, Inc., Intern

Palo Alto, CA, May 2018 – Aug 2018

- Implement scripts that convert simulation models between different tolerance stack-up (GD&T) simulators.

ACADEMIC ACTIVITIES

- **Reviewing:** CDC 2022, ICML 2023, CCTA 2023, CDC 2023, NeurIPS 2023, and ICLR 2024 conferences.
- **Teaching:** Graduate Student Instructor (TA) of Spring, Fall 22, and Fall 23 “IEOR 160: Nonlinear and Discrete Optimization”.
- **Presentation:** Presenter at ACC 2022, CCTA 2023, INFORMS 2021, and MOPTA 2021 conferences.

AWARDS

INFORMS Data Mining Best Paper Competition (student track) Runner-up

Oct 2021

UC Berkeley Graduate Division Block Grant Fellowship

April 2021

Georgia Tech School of ECE Roger P. Webb ECE Senior Scholar Awards

April 2021

CODING SKILLS: Python (PyTorch, CvxPy), MATLAB, LaTeX, C, C++, R, Java, cloud virtual machine.
