

Jeremy Warner jeremywnr.com [Linkedin](#) jeremy.warner@berkeley.edu

Apple — Machine Learning Research Scientist

- ◇ Model Optimization Lead, Video Computer Vision (Science). Serve as the bridge between foundation model research teams and product implementation, overseeing optimization workflows for Apple Intelligence on-device LLMs ([research post: optimization](#)).
- ◇ Technical Gatekeeper for model releases, applying QLoRA, quantization, and cross-format model conversions to enable on-device deployment. Established verification systems ensuring consistent performance across multilingual and multimodal use cases from research to product.

Education

PhD in EECS, UC Berkeley (2023)

Active member of BAIR ([Berkeley AI Research](#)) and BiD ([Berkeley Institute of Design](#)) labs. Research: human-computer interaction, AIML, developer tools, computer vision, graphics.

BS & MS in ECE/CS, University of Rochester Cum laude, high distinction.

Experience

Apple (2023 →) — ML Research Scientist, SIML - Apple Intelligence Model Optimization

Apple (2020, 2021) — Research Scientist Intern, Artificial Intelligence/Machine Learning

Adobe (2019) — Research Scientist Intern, Creative Intelligence Lab

Autodesk (2018) — Research Scientist Intern, User Interface Group

Technical Skills

Machine Learning & AI — PyTorch, JAX, Adapter finetuning (QLoRA), Quantization techniques, Multi-modal model evaluation, Apple Neural Engine on-device optimization

ML Operation — Distributed training/evaluation, Kubeflow pipelines, ML telemetry & observability, Model visualization tools, CI/CD for robust ML workflows and model delivery

Languages & Tools — Python (expert), Rust, TypeScript, Ruby, \LaTeX , Unix, Docker

Teaching

2020, 2017 Head GSI for CSC 160, *User Interface Design and Development*

2019 Head Graduate Student Instructor for CSC 169, *Software Engineering*

Research

See the full publication list (20+ papers) in [Google Scholar](#).

CHI 2024 [Generating Automatic Feedback on UI Mockups with Large Language Models](#)

CHI 2024 [Rambler: Supporting Writing With Speech via LLM-Assisted Gist Manipulation](#)

ICML 2023 [Interactively Optimizing Layout Transfer for Vector Graphics](#)

UIST 2023 [Interactive Flexible Style Transfer for Vector Graphics](#)

IUI 2023 [SlideSpecs: Automatic and Interactive Presentation Feedback Collation](#)

Special Recognition

2023 Google-Berkeley AI Research (BAIR) Commons Grant

2022 University of California, Berkeley EECS Departmental Fellowship

2016 NSF Graduate Student Fellowship Program, *Honorable Mention*