

Zhoujie Ding

Berkeley, CA | +1 (510)717-1017 | zhoujie.ding@berkeley.edu | <https://jasonding0401.github.io>

EDUCATION

University of California, Berkeley

Expected Graduation: 05/2023

B.A. in Computer Science and Applied Mathematics

GPA: 3.96/4.00

Honors: Phi Beta Kappa (top 10% L&S students); Upsilon Pi Epsilon (top 35% declared CS students); Dean's List

Core Coursework: Operating Systems, Database Systems, Algorithms, Optimization Models, Machine Learning, Deep Learning, Computer Vision, Reinforcement Learning, Theoretical Statistics, Real Analysis, Complex Analysis

RESEARCH EXPERIENCE

Undergraduate Researcher

04/2022 – Present

Advisors: Profs. Joseph Gonzalez, Ion Stoica, at UC Berkeley SkyLab (formerly RISELab)

- Researched on [Skyplane](#) project, a tool for fast inter-cloud data transfers. Gained over 600 GitHub stars.
- Benchmarked data transfer speed and cost against AWS DataSync: up to 110x faster and 84% lower cost.
- Implemented continuous profiling with privacy-preserving to improve core APIs and prioritize bug fixes.
- Familiarized with AWS, GCP, Azure CLIs, Python APIs, object storage, and object transfer services.

Undergraduate Researcher

02/2022 – Present

Advisors: Profs. Yizheng Chen, David Wagner, at Berkeley Security Group

- Researched deep learning for vulnerable program detection and android malware detection.
- Curated ~300,000 C/C++ functions from large-scale popular real-world projects.
- Benchmarked vulnerability detection on GitHub code changes, commits, and issues among 3 SOTA models.
- Now focus on fine-tuning large programming language models with curated code and text information.

Undergraduate Researcher

03/2021 – 12/2021

Advisor: Prof. Joseph Gonzalez, Dr. Nathan Pemberton at UC Berkeley RISELab

- Researched on Kernel-as-a-Service project, which manages shareable GPU memory and schedules user kernels across this pool of available GPUs.
- Implemented KaaS by Ray, and integrated Python API-based interfaces for users to easily deploy it.
- Benchmarked on online and offline workloads against Ray actors: 12x lower tail latency, 6x higher endurance.
- Paper published on arXiv here: <https://arxiv.org/abs/2212.08146>.

WORK EXPERIENCE

Software Developer Intern

06/2021 – 08/2021

Manager: Huahua Chen, at Hundsun Technologies Inc., Hangzhou, China

- Devised Word and pdf content information-extraction, and tested over 200 fund contract documents.
- Resolved ~3000 conflict data by Python program for NLP model training, saving 30% time in hand labeling.
- Deployed a Python API with Tornado for document info-extraction and tested 100 HTTP requests with Postman.

PROJECTS

[Facial Keypoint Detection](#)

10/2021 – 12/2021

Class Individual Project

- Modified ResNet-50 to train on iBug dataset of 6666 images of varying image sizes with data augmentation.
- Turned the ground truth keypoint coordinates into pixel-aligned heatmaps and modified FCN-ResNet to train.
- Reached top 10 in class Kaggle competition among over 150 students.

[Pintos Operating System](#)

01/2021 – 05/2021

Class Group Project

- Designed an efficient alarm clock and strict priority scheduler to support multithreaded user programs.
- Added a buffer cache for Pintos file system, support for extending files, and syscalls for directory manipulations.

SKILLS

Programming Languages: Python, C/C++, Java, SQL, RISC-V, Matlab;

Libraries: PyTorch, NumPy, Skimage, Sklearn, Ray, Matplotlib; **Toolkits:** AWS, Docker, Postman, Git, Grafana