Shuo Yuan

shuoyuan@berkeley.edu ◇ 510-701-4942 ◇ shuo-yuan.github.io

Education

University of California, Berkeley Aug. 2026 – Present
Master of Science in Electrical Engineering and Computer Sciences GPA: 4.0/4.0

University of California, Berkeley Aug. 2022 – May 2026
Bachelor of Arts in Computer Science, Bachelor of Arts in Applied Mathematics GPA: 4.0/4.0

- **Honors:** Graduated with ***Highest Distinction in General Scholarship***; five semesters on the Dean's List.
- **Relevant Coursework:** Deep Learning, Machine Learning, Deep Reinforcement Learning, Natural Language Processing, Computer Vision, Probability and Random Processes, Optimization, Algorithms, Forecasting, Advanced Matrix Computation, Advanced LLM Agents, Database Systems, Robotics, Data Structures, Real Analysis, Abstract Linear Algebra, Abstract Algebra, and Numerical Analysis.

Publications and Preprints *Equal contribution · †Main contributor

- **Shuo Yuan**, Dacheng Li, Joseph E. Gonzalez. *D-Skill: Distilling Teacher Knowledge through Agent Skills*. [Paper]
- **Shuo Yuan***, Chengkun Cao*, Huanzhi Mao*, Joseph E. Gonzalez. *I-PTC: Interactive Programmatic Tool Calling for Stateful Tool-Augmented Agents*. **NeurIPS 2026**.
- Dacheng Li, Shiyi Cao, Fangzhou Zhao, **Shuo Yuan**, Sumanth R. Hegde, Joseph E. Gonzalez, Ion Stoica. *RLEM: Reinforcement Learning for External Memory*. **COLM 2026 Efficient Reasoning Workshop**. [Paper]
- Shiyi Cao*, Dacheng Li*, Fangzhou Zhao†, **Shuo Yuan**†, Sumanth R. Hegde†, Connor Chen†, Charlie Ruan†, Tyler Griggs, Shu Liu, Eric Tang, Richard Liaw, Philipp Moritz, Matei Zaharia, Joseph E. Gonzalez, Ion Stoica. *SkyRL-Agent: Efficient RL Training for Multi-turn LLM Agent*. [Paper]
- Shu Liu, Alan Zhu, Sumanth Hegde, Shiyi Cao, **Shuo Yuan**, Samion Suwito, Tyler Griggs, Matei Zaharia, Joseph E. Gonzalez, Ion Stoica. *SkyRL-SQL: Multi-turn SQL Data Agents via RL*. **NeurIPS 2025 MTI-LLM Workshop**. [Paper]
- Shu Liu, Soujanya Ponnapalli, Shreya Shankar, Sepanta Zeighami, Alan Zhu, Shubham Agarwal, Ruiqi Chen, Samion Suwito, **Shuo Yuan**, Ion Stoica, Matei Zaharia, Alvin Cheung, Natacha Crooks, Joseph E Gonzalez, Aditya G Parameswaran. *Supporting Our AI Overlords: Redesigning Data Systems to be Agent-first*. **CIDR 2026**. [Paper]

Research Projects

GRPO for Continuous Rewards Sky Computing Lab, UC Berkeley
Project Lead mentored by Prof. Joseph Gonzalez 09/2026 – Present

- Proposed adaptive reward quantization for GRPO with continuous rewards to improve learning efficiency as the policy evolves.
- Analyzed advantage amplification and information loss, and compared continuous versus quantized rewards in preliminary training experiments.

D-Skill: Distilling Teacher Knowledge through Agent Skills Sky Computing Lab, UC Berkeley
Project Lead mentored by Prof. Joseph Gonzalez, in submission to ICLR 2027 06/2026 – 09/2026

- Led the development of D-Skill, an offline framework that distills paired black-box teacher–student trajectories into reusable natural-language skills without labels or weight updates.
- Designed a difference-analysis and consolidation pipeline that extracts student-specific lessons and supports multiple skill-update backends.
- Evaluated D-Skill across SpreadsheetBench-Verified, ALFWorld, and chess with multiple model families, closing **82–95%** of the teacher–student performance gap.
- Achieved performance comparable to multi-epoch SFT after one offline pass while maintaining **5.7×–8.2× lower inference cost** than the teacher.

Interactive Programatic Tool Calling Sky Computing Lab, UC Berkeley
Project Lead mentored by Prof. Joseph Gonzalez, NeurIPS 2026. 03/2026 – 05/2026

- Developed **I-PTC**, a stateful programmatic tool-calling framework that lets LLM agents execute code in persistent notebook-style environments, reuse intermediate state, inspect outputs, and recover from errors.
- Built **PTC-Bench**, a 240-task benchmark for long-horizon tool-use workflows with simulated APIs, deterministic verifiers, multi-turn follow-ups, and transient failure recovery.
- Implemented **self-evolving API documentation** and **on-demand tool inspection** to reduce repeated API misuse and minimize upfront tool-schema token overhead.
- Evaluated I-PTC across multiple LLM families, improving GPT-5.4 Mini multi-turn pass rate from **69.0%** to **77.1%** and error-recovery pass rate from **61.3%** to **72.9%**, while reducing token usage and redundant API calls.

Reinforcement Learning for External Memory Sky Computing Lab, UC Berkeley
Undergraduate Researcher mentored by Prof. Ion Stoica and Prof. Joseph Gonzalez 09/2025 – 01/2026

- Designed a trainable external-memory + RL framework that lets LLMs handle contexts up to around 1M tokens.
- Implemented a step-wise GRPO RL scheme that gives fine-grained rewards on each read/write step, teaching models to use external memory effectively.
- Achieved ~10% absolute accuracy gains over an 8B baseline on RULER-HotpotQA across context lengths up to 1M tokens with a 4B model.

SkyRL-Agents, Computer Use Agents Post-Training Sky Computing Lab, UC Berkeley
Main Contributor, Undergraduate Researcher mentored by Prof. Joseph Gonzalez 05/2025 - 10/2025

- Integrated computer use agent and its environment into SkyRL-Agent and trained multi-step agents with VeRL or SkyRL.
- Optimized the library `desktop_env` for computer use agents, allowing supports for async workflow and evaluation and resulting in a speedup from 1 hr. to 25 min for a single training step.
- Implemented a distributed setup for online RL training with virtual desktop environments, allowing an efficient use of both CPU and GPU nodes and scaling up the number of environments from 8 to 128.

Train SQL Data Agents with Multi-turn RL Sky Computing Lab, UC Berkeley
Undergraduate Researcher mentored by Prof. Ion Stoica and Prof. Joseph Gonzalez 05/2025 - 09/2025

- Trained Qwen2.5-Coder-7B on multi-turn SQL generation with GRPO, improving the performance on the Spider benchmark for 7% with only 600 data. The work is accepted by NeurIPS 2025 MTI-LLM workshop.
- Tested our model with no schema information and proved its ability to explore and probe the database by itself after training (a 13% improvement over the baseline).
- Scaled the experiments on the number of turns and passes and analyzed the query plans generated by the LLMs, arguing for the inefficiency in modern database design.
- Proposed a new paradigm for databases to adapt to workloads posed by agents and LLMs. The vision paper is accepted by CIDR 2026.

Teaching & Academic Service

TA for EECS 183: Natural Language Processing 08/2026 - Present
TA for COMPSCI 170: Efficient Algorithms and Intractable Problems 08/2024 - 05/2026
Reviewer for NeurIPS 2025 MTI-LLM Workshop 09/2025

Work Experience

Linear Algebra Game Development Department of Mathematics, UC Berkeley
Software Developer supervised by Prof. David Nadler 08/2024 - 12/2024

- Designed and visualized mini-games for a college level linear algebra course that is taken by more than 300 students (MATH 56, Linear Algebra).
- Implemented the web games (vector addition, vector projection, etc.) with typescript. See [Demo].
- Worked with professors from math department together on better linear algebra learning experience for students.

Skills

- **Programming Languages:** Python, Java, C, SQL, MATLAB, R, RISC-V, HTML, LaTeX
- **Related Libraries:** PyTorch, NumPy, Scikit-learn, vLLM, Ray, VeRL, etc.
- **Other skills:** Statistical Analysis, Algorithm Design, Software Development, Teaching