

# Teddi Worledge

PhD candidate working on LLM citations, human-AI collaboration, and legal applications

[worledge@stanford.edu](mailto:worledge@stanford.edu) | Stanford, CA | [LinkedIn](#) | [GitHub](#) | [Personal](#)

## EDUCATION

---

**Stanford University** 2022–Present

Computer Science PhD in Artificial Intelligence • Advisor: Carlos Guestrin • Expected Graduation: 2027

Teaching: CS 281: Ethics of Artificial Intelligence

**UC Berkeley** 2018–2022

Computer Science B.A. • Research Advisor: Ben Recht

## EXPERIENCE

---

**Stanford Law School Deborah L. Rhode Center • AI Fellow** 2025–Present

- Built an LLM assistant for the Los Angeles Superior Court (LASC) to recommend judgments in debt collection cases (30k cases per year) with trustworthy citations to case filings. 2026 LASC deployment into live judicial workflows.
- Evaluated human-AI collaboration in a controlled study with 65+ law students; the LLM assistant reduced errors on legal requirements by over 40% and cut review time by 25%.
- Worked with LASC attorneys to translate CA Civil Code §§ 1788.58–60 into a technical system.
- Designed a user interface for legal review, verification, and defect tracking.

**Stanford Artificial Intelligence Lab • PhD Researcher** 2022–Present

- Conducted human-subject experiments to quantify trade-offs across 6 LLM systems (implemented 5); observed a 50% reduction in sentences fully supported by citations, +200% perceived utility, and 3x longer user verification times as outputs became more abstractive w.r.t. cited sources.
- Developed a unifying framework for LLM citations over training data attribution and NLP citation to enable the design of more reliable and verifiable LLM applications in high-stakes domains.

**Berkeley Artificial Intelligence Research • Research Assistant** 2018–2021

- Trained deep neural networks with ML fairness algorithms (e.g., group DRO) and applied scaling laws to project optimal training data composition across population subgroups at larger dataset sizes.

**Dropbox, Core Desktop Team • Software Engineering Intern** Summer 2020

- Engineered and shipped a Python backend for the user-facing sync history. Collaborated with UX designers, product managers, and frontend engineers to deliver an intuitive user experience.

## PUBLICATIONS

---

**Theodora Worledge**, Othman Bensouda Koraichi, Daniel Bernal, Aviv Caspi, Tatsunori Hashimoto, Carlos Guestrin, and David Freeman Engstrom. "AI Assistance for Court Review of Default Judgments." *Under Review*. 2026.

**Theodora Worledge**, Tatsunori Hashimoto, and Carlos Guestrin. "The Extractive-Abstractive Spectrum: Uncovering Verifiability Trade-offs in LLM Generations." *Accepted pending revision: Transactions of the Association for Computational Linguistics*. 2025.

**Theodora Worledge\***, Judy Hanwen Shen\*, Nicole Meister, Caleb Winston, and Carlos Guestrin. "Unifying Corroborative and Contributive Attributions in Large Language Models." *IEEE Conference on Secure and Trustworthy Machine Learning*. 2024. \* Co-first authors

Esther Rolf, **Theodora Worledge**, Ben Recht, and Michael Jordan. "Representation Matters: Assessing the Importance of Subgroup Allocations in Training Data." *International Conference on Machine Learning*. 2021.

## TALKS & AWARDS

---

**Talks:** NeurIPS ATTRIB, JURIX AIDA2J, ACM CS & Law, LA Superior Court Civil Bench Judicial Seminar, Stanford Legal Design AI+A2J, Wonderfest Science, Code for Humanities

**Awards & Grants:** NSF Graduate Research Fellow (2022–2027), HAI-Google Research Token (2025), UC Berkeley Regents' & Chancellor's Scholar (2018–2022)