

NEEL GUHA

ng3108@columbia.edu
<https://neelguha.github.io/>

ACADEMIC APPOINTMENTS

Columbia Law School

Associate Professor of Law

Courses: Torts

New York, NY
2026–Present

EDUCATION

Stanford University

Ph.D., Computer Science, expected 2026

Advisor: Christopher Ré

Thesis: *Insights from Law for Building and Evaluating AI Systems*

Fellowships: Stanford Interdisciplinary Graduate Fellowship
Stanford HAI Graduate Fellowship

Stanford, CA

Stanford Law School

J.D., 2026

Honors: Gerald Gunther Prize for Outstanding Performance in Health Law
John Hart Ely Prize for Outstanding Performance in a Policy Lab

Activities: Senior Articles Editor, Stanford Law Review (Volume 75)
Member Editor, Stanford Law Review (Volume 74)

Stanford, CA

Carnegie Mellon University

M.S., Machine Learning, 2019

Advisor: Virginia Smith

Pittsburgh, PA

Stanford University

B.S. with Honors, Computer Science, 2018

Thesis: *Aggregating Model Spaces*

Stanford, CA

WORKS-IN-PROGRESS/PREPRINTS

Law Professors Prefer AI Over Peer Answers

Alejandro Salinas, Carly Frieders, Neel Guha, Sibom Ma, Ralph Anzivino, Ian Ayres, Oren Bar-Gill, Omri Ben-Shahar, Stephen Friedman, George Geis, Sue Guan, Christoph Henkel, Stephanie Hoffer, Gregory Klass, Larasz Moody-Villarose, Sarath Sanga, Keith Sharfman, Justin Simard, Rebecca Stone, David Wishnick, and Julian Nyarko

Designing Application-Specific AI Regulation

Neel Guha

LAW REVIEW ARTICLES

The State Statutes Project

2024 WIS. L. REV. 1615

Neel Guha and Diego Zambrano

AI Regulation Has Its Own Alignment Problem: The Technical and Institutional Feasibility of Disclosure, Registration, Licensing, and Auditing

92 GEORGE WASHINGTON LAW REVIEW 1473 (2024)

Neel Guha*, Christie M. Lawrence*, Lindsey A. Gailmard, Kit T. Rodolfa, Faiz Surani, Rishi Bommasani, Inioluwa Deborah Raji, Mariano-Florentino Cuéllar, Colleen Honigsberg, Percy Liang, and Daniel E. Ho

Private Enforcement in the States

172 UNIVERSITY OF PENNSYLVANIA LAW REVIEW 61 (2023)

Diego Zambrano, Neel Guha, Austin Peters, and Jeffrey Xia

Vulnerabilities in Discovery Tech

35 HARVARD JOURNAL OF LAW & TECHNOLOGY 281 (2022)

Neel Guha, Peter Henderson, and Diego Zambrano

JOURNALS

There's No Free Benchmark: An Institutional View of Legal AI Benchmarking

Proceedings of the National Academy of Sciences (2026)

Neel Guha, Andy K. Zhang, Christine Tsang, Christopher D. Manning, Julian Nyarko, and Daniel E. Ho

Open Problems in Technical AI Governance

Transactions on Machine Learning Research (2025)

Anka Reuel*, Ben Bucknall*, Stephen Casper, Tim Fist, Lisa Soder, Onni Aarne, Lewis Hammond, Lujain Ibrahim, Alan Chan, Peter Wills, Markus Anderljung, Ben Garfinkel, Lennart Heim, Andrew Trask, Gabriel Mukobi, Rylan Schaeffer, Mauricio Baker, Sara Hooker, Irene Solaiman, Alexandra Sasha Luccioni, Nitarshan Rajkumar, Nicolas Moës, Jeffrey Ladish, Neel Guha, Jessica Newman, Yoshua Bengio, Tobin South, Alex Pentland, Sanmi Koyejo, Mykel J. Kochenderfer and Robert Trager

Understanding Liability Risk from Using Health Care Artificial Intelligence Tools

New England Journal of Medicine (2024)

Michelle M. Mello and Neel Guha

Holistic Evaluation of Language Models

Transactions on Machine Learning Research (2023)

Percy Liang*, Rishi Bommasani*, Tony Lee*, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, and Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D. Manning, Christopher Ré, Diana AcostaNavas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda

ChatGPT and Physicians' Malpractice Risk

JAMA Health Forum (2023)

Michelle M. Mello and Neel Guha

COMPUTER SCIENCE CONFERENCES

Cartridges: Lightweight and General-purpose Long Context Representations via Self-Study

International Conference on Learning Representations (2026)

Sabri Eyuboglu*, Ryan Ehrlich*, Simran Arora*, Neel Guha, Dylan Zinsley, Emily Liu, Will Tennien, Atri Rudra, James Zou, Azalia Mirhoseini and Christopher Ré

An Architecture Search Framework for Inference-Time Techniques

International Conference on Machine Learning (2025)

Jon Saad-Falcon, Adrian Gamarra Lafuente, Shlok Natarajan, Nahum Maru, Hristo Todorov, Etash Guha, E. Kelly Buchanan, Mayee Chen, Neel Guha, Christopher Ré and Azalia Mirhoseini

A Reasoning-Focused Legal Retrieval Benchmark

4th ACM Symposium on Computer Science and Law (2025)

Lucia Zheng*, Neel Guha*, Javokhir Arifov, Sarah R. Zhang, Michal Skreta, Christopher D. Manning, Peter Henderson and Daniel E. Ho

Smoothie: Label Free Language Model Routing

Conference on Neural Information Processing Systems (2024)

Neel Guha*, Mayee F. Chen*, Trevor Chow, Ishan S. Khare and Christopher Ré

Stronger Than You Think: Benchmarking Weak Supervision on Realistic Tasks

Conference on Neural Information Processing Systems, Datasets and Benchmarks Track (2024)

Also appeared in: Workshop on Data-Centric Machine Learning Research at ICML (2024)

Tianyi Zhang, Linrong Cai, Jeffrey Li, Nicholas Roberts, Neel Guha and Frederic Sala

Benchmarking and Building Long-Context Retrieval Models with LoCo and M2-BERT

International Conference on Machine Learning (2024)

Also appeared in: Workshop on Mathematical and Empirical Understanding of Foundation Models at ICLR (2024)

Jon Saad-Falcon, Daniel Y. Fu, Simran Arora, Neel Guha and Christopher Ré

Prospector Heads: Generalized Feature Attribution for Large Models & Data

International Conference on Machine Learning (2024)

Gautam Machiraju*, Alexander Derry*, Arjun Desai, Neel Guha, Amir-Hossein Karimi, James Zou, Russ Altman, Christopher Ré and Parag Mallick

LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models

Conference on Neural Information Processing Systems, Datasets and Benchmarks Track (2023)

Neel Guha*, Julian Nyarko*, Daniel E. Ho*, Christopher Ré*, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel N. Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Heg-

land, Jessica Wu, Joe Nudell, Joel Niklaus, John Nay, Jonathan H. Choi, Kevin Tobia, Margaret Hagan, Megan Ma, Michael Livermore, Nikon Rasumov Rahe, Nils Holzenberger, Noam Kolt, Peter Henderson, Sean Rehaag, Sharad Goel, Shang Gao, Spencer Williams, Sunny Gandhi, Tom Zur, Varun Iyer, and Zehua Li

Embroid: Unsupervised Prediction Smoothing Can Improve Few-Shot Classification

Conference on Neural Information Processing Systems (2023)

Neel Guha*, Mayee F. Chen*, Kush Bhatia*, Azalia Mirhoseini, Frederic Sala, and Christopher Ré

Don't Use a Cannon to Kill a Fly: An Efficient Cascading Pipeline for Long Documents

International Conference on AI and Law (2023)

Zehua Li, Neel Guha, and Julian Nyarko

Ask Me Anything: A Simple Strategy for Prompting Language Models

International Conference on Learning Representations (2023)

Simran Arora*, Avanika Narayan*, Mayee F. Chen, Laurel J. Orr, Neel Guha, Kush Bhatia, Ines Chami, Frederic Sala, and Christopher Ré

Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open Source Legal Dataset

Conference on Neural Information Processing Systems, Datasets and Benchmarks Track (2022)

Peter Henderson*, Mark S. Krass*, Lucia Zheng, Neel Guha, Christopher D. Manning, Dan Jurafsky and Daniel E. Ho

When Does Pretraining Help? Assessing Self-Supervised Learning for Law and the CaseHOLD Dataset

International Conference on AI and Law (2021)

Lucia Zheng*, Neel Guha*, Brandon R. Anderson, Peter Henderson and Daniel E. Ho

Bootleg: Chasing the Tail with Self-Supervised Named Entity Disambiguation

Conference on Innovative Data Systems Research (2021)

Laurel Orr*, Megan Leszczynski*, Simran Arora, Sen Wu, Neel Guha, Xiao Ling and Christopher Ré

Leveraging Administrative Data for Bias Audits: Assessing Disparate Coverage with Mobility Data for COVID-19 Policy

ACM Conference on Fairness, Accountability, and Transparency (2020)

Amanda Coston, Neel Guha, Lisa Lu, Derek Ouyang, Alexandra Chouldechova, and Daniel E. Ho

Machine Learning for AC Optimal Power Flow

Climate Change Workshop at the International Conference on Machine Learning (2019)

Neel Guha, Zhecheng Wang, Matt Wytock, and Arun Majumdar

One-Shot Federated Learning

2nd Workshop on Machine Learning on the Phone and other Consumer Devices at Neural Information Processing Systems (2018)

Neel Guha, Ameet Talwalkar, and Virginia Smith

WHITE PAPERS

[Proceedings of the 2025 Workshop on Law-Following AI](#)

Cullen O’Keefe, Christoph Winter, Matthijs Maas, Janna Tay, Yonathan Arbel, Katie Atkinson, Sebastian Benthall, Jack Boeglin, Claire Boine, Thomas Burri, Nicholas Caputo, Karni Chagal-Feferkorn, Alan Chan, Inyoung Cheong, Mihailis Diamantis, Talita Dias, Michael Dorff, Kevin Frazier, [Neel Guha](#), Julius Mentis Hattingh, Maarten Herbosch, Apostolos Latsonas, Dan Linna, Anat Lior, Jakob Merane, Sunayana Rane, Peter Salib, Holli Sargeant, Daniel Schwarcz, Niloufer Selvadurai, Henry A. Thompson, Paul Weitzel, and Spencer Williams

[HAI Policy Brief: Understanding Liability Risk from Healthcare AI](#) (2024)

Michelle M. Mello and [Neel Guha](#)

[HAI Policy Brief: The AI Regulatory Alignment Problem](#) (2023)

[Neel Guha](#)*, Christie M. Lawrence*, Lindsey A. Gilmard, Kit T. Rodolfa, Faiz Surani, Rishi Bommasani, Inioluwa Deborah Raji, Mariano-Florentino Cuéllar, Colleen Honigsberg, Percy Liang, and Daniel E. Ho

[Building a National AI Research Resource: A Blueprint for the National Research Cloud](#) (2021)

Contributing Author: §5 (Data Privacy Compliance), §6 (Technical Privacy and Virtual Data Safe Rooms), §8 (Managing Cybersecurity Risks).

[On the Opportunities and Risks of Foundation Models](#) (2021)

Contributing Author: §3.2 (Legal Applications) and §5.4 (Legality)

[ResX Task Force Final Report: Our Vision for Stanford’s Undergraduate Residences](#) (2018)

ResX Task Force

BOOK CHAPTERS

AI Assurance, in *Equalizing Justice: Harnessing AI for Litigants Without Lawyers* (Cambridge Univ. Press, forthcoming July 2026)

[Neel Guha](#) and Izak C. Rosenfeld

Building GenAI Benchmarks: A Case Study in Legal Applications, in *The Oxford Handbook on the Foundations and Regulation of Generative AI* (Oxford University Press, 2025)

[Neel Guha](#), Julian Nyarko, Daniel E. Ho, and Christopher Ré

Gamesmanship in Modern Discovery Tech, in *Legal Tech and the Future of Civil Justice* (Cambridge University Press, 2023)

Diego Zambrano, [Neel Guha](#), and Peter Henderson

TEACHING EXPERIENCE

CS 103: Mathematical Foundations of Computing, Stanford University (Teaching Assistant) (Fall 2024)

CS 103: Mathematical Foundations of Computing, Stanford University (Teaching Assistant) (Summer 2024)

TALKS AND PRESENTATIONS

Invited Talks and Paper Presentations

An Institutional View of Legal Benchmarking

- Max Planck Institute, virtual (May 2026)

Legal Telescopes: New Ways of Seeing the Law at Scale

- Berkman Klein Center, Cambridge, MA (February 2026)

Designing Application-Specific AI Regulation

- Accountable AI Research Conference, Wharton (February 2026)
- Law School Faculty Workshop, Harvard University (February 2026)
- Law School Faculty Workshop, Columbia University (January 2026)
- Law School Faculty Workshop, Stanford University (January 2026)
- Law School Faculty Workshop, Wayne State University (November 2025)
- Law School Faculty Workshop, Washington University in St. Louis (November 2025)
- Law School Faculty Workshop, Texas A&M University (November 2025)
- Law School Faculty Workshop, Michigan State University (November 2025)
- Law School Faculty Workshop, George Mason University (November 2025)
- Law School Faculty Workshop, University of Arizona (October 2025)
- Law School Faculty Workshop, University of Miami (October 2025)
- Law School Faculty Workshop, University of San Diego (October 2025)
- Fourth Annual Legal Scholars Roundtable on Artificial Intelligence, Emory University (April 2025)
- Emerging Scholars in Law & Economics, AALS Annual Meeting (January 2025)

The State Statutes Project

- Symposium on Legal Interpretation and Data, Georgetown University Law Center (March 2025)
- Public Law in the States Conference, University of Wisconsin (May 2024)

AI Evaluation and Governance

- Object Management Group AI Task Force, San Francisco, CA (December 2024)

Benchmarks for Legal RAG

- Max Planck Institute, virtual (November 2024)

Large Language Models for Law and Civil Justice

- Summer Institute in Computational Social Science, Adam Mickiewicz University, virtual (June 2024)
- AI & Legal Help Crossover, Stanford Law School Legal Design Lab, virtual (July 2023)

Private Enforcement in the States

- Law & Geopolitics Colloquium, University of California, Berkeley School of Law (with Diego Zambrano) (January 2024)
- Faculty Colloquium, University of Washington School of Law (with Diego Zambrano) (November 2023)
- Conference on Empirical Legal Studies (with Diego Zambrano and Austin Peters) (October 2023)
- Public Law Workshop, University of Minnesota Law School (with Diego Zambrano) (October 2023)
- Civil Procedure Workshop, Northwestern University (with Diego Zambrano, Austin Peters, and Jeffrey Xia) (May 2023)

Understanding Liability Risk from Health Care AI Tools

- Stanford Human-Centered AI Seminar (with Michelle Mello) (January 2024)
- Anesthesia Grand Rounds, Stanford University School of Medicine (with Michelle Mello) (August 2023)

LegalBench

- LAW 806Y: Justice By Design, Stanford Law School (January 2024)
- Google, Mountain View, CA (November 2023)
- LLM x Law Hackathon, New York, NY (June 2023)
- CodeX FutureLaw, Stanford Law School (April 2023)

Machine Learning for AC Optimal Power Flow

- Climate Change Workshop, International Conference on Learning Representations, virtual (update on prior work) (June 2020)
- Climate Change Workshop, International Conference on Learning Representations, Long Beach, CA (June 2019)

One-Shot Federated Learning

- Workshop on Machine Learning on the Phone and Other Consumer Devices, NeurIPS, Montreal, Canada (December 2018)

Panels and Roundtables

AI Evaluations for Law, AI4Law Workshop, ICML (July 2026)

Computer Science Perspectives on AI Legal Reasoning, Workshop on Law-Following AI, Christ's College, University of Cambridge (August 2025)

AI in Action: New and Emerging Research, Programs, and Policies, Access to Justice and AI, Arizona State University (January 2025)

Synthetic Data, Stanford Law School AI + Access to Justice Summit (October 2024)

The Risks, Rewards, and Ethics of Using Artificial Intelligence, Annual Bench Bar Conference, Bar Association of San Francisco (November 2023)

SELECTED HONORS

The National Law Review / Wickard Top 50 Legal Innovators in Academia (2026)

Finalist, Paul & Daisy Soros Fellowship for New Americans (2021)

Carole Hafner Best Paper Award, International Conference on Artificial Intelligence and Law (2021)

Finalist, Knight-Hennessy Scholars (2020)

Best Paper Honorable Mention, Climate Change Workshop at ICML (2019)

SERVICE

Referee and program committee member: NeurIPS Main Track, NeurIPS Datasets & Benchmarks Track, ICAIL, ICML, ICLR, Journal of Empirical Legal Studies.

Discussant: Conference on Empirical Legal Studies.

Committee and leadership service:

- Healthcare AI Policy Steering Committee, Stanford Institute for Human-Centered Artificial Intelligence (2025)
- ResX Provostial Task Force, Stanford University (2018)
- Committee on Residential Learning, Stanford University (2017–2018)