

Todd Morrill

PhD Student, Computer Science · Columbia University

 [toddmorrill.github.io](https://github.com/toddmorrill)
 todd@cs.columbia.edu
 New York, NY

EDUCATION

Columbia University <i>PhD, Computer Science</i> Advisor: Richard Zemel	2024 – present (expected May 2028) New York, NY
Columbia University <i>MS, Computer Science</i> Advisors: Richard Zemel, Kathleen McKeown	Dec 2023 New York, NY
Harvard Extension School <i>Computer Science coursework</i>	May 2021 Online
George Washington University <i>BBA, Finance; minors in Economics and Chinese</i>	May 2014 Washington, DC
Middlebury College <i>Chinese Summer Language Program</i>	Summer 2011 Middlebury, VT

RESEARCH EXPERIENCE

Columbia University <i>PhD Student, Prof. Richard Zemel, Computer Science</i>	Sep 2024 – present New York, NY
• Derived a test-time sequence memory model for predictive coding (under review). • Founded and lead the Biological Learning Working Group of ARNI, the NSF AI Institute for Artificial and Natural Intelligence, focused on credit assignment (how systems adjust their parameters) in biological and artificial systems; direct 6 PhD students, set the group's research agenda, and report progress via presentations and written updates to ARNI leadership and the NSF. • Developed mixture-of-experts survival models with Columbia University Irving Medical Center clinicians that cluster similar patients while cutting calibration error by up to 70% and improving concordance [C2, W1].	
Cold Spring Harbor Laboratory <i>NeuroAI Research Intern & Visiting Scholar, Prof. Anthony Zador, Neuroscience</i>	May 2024 – present Cold Spring Harbor, NY
• Published <i>Bullet Trains</i>, a first-author ICML 2026 paper, and open-sourced our JAX system that trains event-based SNNs up to 44x faster than sequential baselines while preserving exact neuron dynamics [C1].	
Columbia University <i>Graduate Research Assistant (MS), Profs. Richard Zemel, Kathleen McKeown</i>	Sep 2022 – May 2024 New York, NY
• Co-developed <i>Prompt Risk Control</i>, a framework for bounding the deployment risk of large language models, in Prof. Richard Zemel's lab; published at ICLR 2024 [C4, W2]. • Developed a circumplex-theory approach to dialogue analysis for the DARPA Computational Cultural Understanding program in Prof. Kathleen McKeown's lab; oral presentation at LREC-COLING 2024 [C3].	

TEACHING EXPERIENCE

Columbia University <i>Teaching Assistant, Continual Learning & Memory Models (Prof. Richard Zemel)</i>	Spring 2026 New York, NY
• Lectured on neural memory models to a class of 50, moderated weekly in-class paper discussions, and supervised 6 student research projects.	

RESEARCH INTERESTS

Machine learning and deep learning; biologically plausible learning and credit assignment; memory models; spiking neural networks; unsupervised and multimodal representation learning; neuroscience- and cognitive-science-inspired machine learning.

PROFESSIONAL EXPERIENCE

PricewaterhouseCoopers (PwC)

Aug 2014 – Mar 2024

Data Scientist, Senior Manager

New York, NY

Large Language Model Development

- Built PwC's in-house LLM pipeline in 2023, from FSDP training of Llama 2 (7B–13B) on a 16-GPU A100 cluster through inference benchmarking, load testing, and a hosting cost model for Falcon 40B-class models; first author of the resulting technical report [R1].

Automated Knowledge Graph Construction

- Implemented and compared 7 methods for extracting entities and relations from arbitrary text; accepted to the AKBC 2021 Workshop on Unstructured and Structured Knowledge Bases [W3]. Co-inventor on 5 patents (2 granted) [P3–P7].

Information Retrieval and Classification for a Bank

- Managed a team of 6 and collaborated with CMU (Eduard Hovy) to build models that classify regulatory text for compliance purposes.
- Lead inventor on 2 patents (1 granted) [P1, P2]; reduced time spent by compliance experts by up to 50% [bit.ly/32f2lx3].

Enterprise Search

- Reduced the number of clicks required for 50K+ users to find information in PwC's firm-wide search portal; presented at Google Cloud Next '18 [youtu.be/63EANKPzuJY].

Selected Projects Led

- Domain-adapting large language models; an EEG-based brain–computer interface [youtu.be/f81T0KcprpM]; deep learning on edge devices [youtu.be/iqVKjye-J68]; structural causal modeling with CMU students; dynamic pricing for an online ticketing platform; and forecasting, classification, and information-extraction models for technology, insurance, automotive, food service, and financial services clients.

PEER-REVIEWED CONFERENCE PAPERS

- [C1] **Todd Morrill**, Christian Pehle, Anthony Zador. Bullet Trains: Parallelizing Training of Temporally Precise Spiking Neural Networks. *International Conference on Machine Learning (ICML)*, 2026.
- [C2] **Todd Morrill**, Aahlad Manas Puli, Murad Megjhani, Soojin Park, Richard Zemel. Let the Experts Speak: Improving Survival Prediction & Calibration via Mixture-of-Experts Heads. *Machine Learning for Health (ML4H) Symposium*, 2025.
- [C3] **Todd Morrill**, Zhaoyuan Deng, Yanda Chen, Amith Ananthram, Colin Wayne Leach, Kathleen McKeown. Social Orientation: A New Feature for Dialogue Analysis. *LREC-COLING*, 2024. Oral presentation.
- [C4] Thomas Zollo, **Todd Morrill**, Zhun Deng, Jake Snell, Toniann Pitassi, Richard Zemel. Prompt Risk Control: A Rigorous Framework for Responsible Deployment of Large Language Models. *International Conference on Learning Representations (ICLR)*, 2024.

PEER-REVIEWED WORKSHOP PAPERS

- [W1] **Todd Morrill**, Aahlad Manas Puli, Murad Megjhani, Soojin Park, Richard Zemel. Let the Experts Speak: Improving Survival Prediction & Calibration via Mixture-of-Experts Heads. *Learning from Time Series for Health Workshop, NeurIPS*, 2025.

- [W2] Thomas Zollo, **Todd Morrill**, Zhun Deng, Jake Snell, Toniann Pitassi, Richard Zemel. Prompt Risk Control: A Rigorous Framework for Responsible Deployment of Large Language Models. *Socially Responsible Language Modelling Research (SoLaR) Workshop, NeurIPS*, 2023.
- [W3] **Todd Morrill**. Cluster Based Named Entity Recognition. *Workshop on Unstructured and Structured Knowledge Bases, AKBC*, 2021.

TECHNICAL REPORTS

- [R1] **Todd Morrill**, Michael Kirchner, Vasudeva Sankaranarayanan, Joseph Voyles. Technical Report: Large Language Models. 2023.
- [R2] **Todd Morrill**, Satyam Sharma. Comparing Hyperdimensional Computing to Deep Learning for Natural Language Processing Tasks. 2023.

PATENTS

- [P1] **Todd Morrill**, Eric Roma, Nicolas Kuzak, Neelam Sharma, Andrew Runge, Jayvardhan Rathi, Waqar Sarguroh, Wenting Zhao. Regulatory Tree Parser. US Patent 12,072,861, granted 2024 (filed 2022).
- [P2] **Todd Morrill**, Eric Roma, Neelam Sharma, Alistair Moore, Andrew Runge. Regulatory Obligation Identifier. US Patent App. 2022/0374914, pending (filed 2022).
- [P3] Paul Sheward, Chung-Sheng Li, Scott Likens, and 9 others incl. **Todd Morrill**. Method and Apparatus for an AI-Assisted Virtual Consultant. US Patent 12,282,501, granted 2025.
- [P4] Paul Sheward, Chung-Sheng Li, Scott Likens, and 9 others incl. **Todd Morrill**. Method and Apparatus for Knowledge Representation and Reasoning in Accounting. US Patent 12,229,195, granted 2025.
- [P5] Paul Sheward, Chung-Sheng Li, Scott Likens, and 9 others incl. **Todd Morrill**. Systems and Methods for Constructing Topic-Specific Knowledge Graphs. US Patent App. 2024/0394563, pending.
- [P6] Paul Sheward, Chung-Sheng Li, Scott Likens, and 9 others incl. **Todd Morrill**. Method and Apparatus for an AI-Assisted Virtual Assistant. US Patent App. 2024/0403302, pending.
- [P7] Paul Sheward, Chung-Sheng Li, Scott Likens, and 9 others incl. **Todd Morrill**. Method and Apparatus for AI-Assisted Virtual Assistant for SME Agent. US Patent App. 2025/0028742, pending.

TALKS AND PRESENTATIONS

1. Memory Models in Artificial and Biological Systems. Entrypoint Capital, 2026.
2. Local Circuits: Memory, Neuroscience and AI. Zuckerman Institute, Columbia University, 2026.
3. Machine Learning Enabled Chatbots. Google Cloud Next, 2018.
4. Machine Learning Enabled Chatbots. AI4 Finance, 2018.

TECHNICAL SKILLS

Programming languages	Python, C/C++ (+CUDA), JavaScript, SQL, Bash, LaTeX, Markdown
ML libraries	JAX, PyTorch, TensorFlow/Keras, scikit-learn, pandas, NumPy, SciPy
Tools	SLURM, Docker, Kubernetes, Git/GitHub, VS Code
Cloud	Google Cloud, Amazon Web Services, Microsoft Azure
Spoken languages	Professional proficiency in Mandarin Chinese

LEADERSHIP AND SERVICE

- Midnight Run NYC volunteer, 2024 – present
- Harlem Grown volunteer, 2014 – 2020
- GWU Alternative Breaks volunteer, rebuilding after Hurricane Katrina, 2011 – 2013
- Eagle Scout, Boy Scouts of America, 2010