Zachary Horvitz

zacharyhorvitz | zachary-horvitz | zacharyhorvitz.com
zacharyhorvitz@gmail.com | tel:+1.425.999.9273

Summary

Third-year PhD Student at Columbia University, formerly a Machine Learning Engineer at Rad AI. I am excited about **inference-time scaling**, **text diffusion models**, building robust **verifiers** for text generation tasks, and improving language models with **synthetic data**!

Experience

PhD Student - McKeown, Yu Labs (Columbia University) Winter 2023 - Present

Currently researching controllable text generation, along with exploring properties and applications of LLMs, synthetic data, and text diffusion models.

Senior MLE, Research Lead - Rad AI Fall 2020 - Winter 2023

- Researched, engineered, and deployed personalized summarization models to radiologists.
- Developed language modeling libraries and fielded models that dramatically improved medical accuracy.
- Reduced training times by 80 percent and enabled parameter-efficient scaling to 30+ practices.
- Supervised model training and mentored three machine learning engineers.
- Rad AI now serves $>$ 10% of American radiologists, generates millions of summaries each month.

Graduate Researcher - IRL (Brown University) Winter 2019 - Fall 2020

Incorporated linguistic and graph-based priors to improve multitask learning for DQNs at Intelligent Robot Lab, advised by Prof. George Konidaris.

Graduate Researcher - RLAB (Brown University) Fall 2019 - Summer 2020

Built pipeline for generating topical satire given news articles, advised by Prof. Michael Littman. One of our system's generated headlines was accepted and published in a satirical newspaper.

Research Intern - AI2 Spring/Summer 2015, Summer 2018

Applied entity-augmented language models for common-sense story generation under the guidance of Yejin Choi, Maarten Sap, and Antoine Bosselut at AI2 and UW. Previously, I worked on automated pipelines for parsing data in academic figures.

Education

2023 - present	PhD (Computer Science) at **Columbia University**	(GPA: 4.1/4.0)
2019 - 2020	ScM (Computer Science) **Brown University**	(GPA: 4.0/4.0)
2015 - 2019	AB (Computer Science, Anthropology) **Brown University**	(GPA: 3.96/4.0)

Publications

Singhal*, Raghav, **Zachary Horvitz***, Ryan Teehan*, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath (Jan. 2025). *A General Framework for Inference-time Scaling and Steering of Diffusion Models*. **Accepted to ICML 2025**. arXiv: 2501.06848 [cs.LG]. URL: https://arxiv.org/abs/2501.06848.

Patel, Ajay, Jiacheng Zhu, Justin Qiu, **Zachary Horvitz**, Marianna Apidianaki, Kathleen McKeown, and Christopher Callison-Burch (Oct. 2024). "StyleDistance: Stronger Content-Independent Style Em-

beddings with Synthetic Parallel Examples". In: Accepted to NAACL 2025. URL: https://api.semanticscholar.org/CorpusID:273375301.

Zachary Horvitz, Ajay Patel, Kanishk Singh, Chris Callison-Burch, Kathleen McKeown, and Zhou Yu (Nov. 2024). "TinyStyler: Efficient Few-Shot Text Style Transfer with Authorship Embeddings". In: *Findings of the Association for Computational Linguistics: EMNLP 2024*. Ed. by Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen. Miami, Florida, USA: Association for Computational Linguistics, pp. 13376–13390. URL: https://aclanthology.org/2024.findings-emnlp.781.

Zachary Horvitz*, Jingru Chen*, Rahul Aditya, Harshvardhan Srivastava, Robert West, Zhou Yu, and Kathleen McKeown (Aug. 2024). "Getting Serious about Humor: Crafting Humor Datasets with Unfunny Large Language Models". In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Ed. by Lun-Wei Ku, Andre Martins, and Vivek Srikumar. **Outstanding Paper Award**. Bangkok, Thailand: Association for Computational Linguistics, pp. 855–869. DOI: 10.18653/v1/2024.acl-short.76. URL: https://aclanthology.org/2024.acl-short.76.

Zachary Horvitz, Ajay Patel, Chris Callison-Burch, Zhou Yu, and Kathleen McKeown (Mar. 2024). "ParaGuide: Guided Diffusion Paraphrasers for Plug-and-Play Textual Style Transfer". In: *Proceedings of the AAAI Conference on Artificial Intelligence* 38, pp. 18216–18224. DOI: 10.1609/aaai.v38i16.29780. URL: https://ojs.aaai.org/index.php/AAAI/article/view/29780.

Parikh*, Neev, **Zachary Horvitz***, Naveen Srinivasan*, Aansh Shah, and George Dimitri Konidaris (Dec. 2020). "Graph Embedding Priors for Multi-task Deep Reinforcement Learning". In: *NeurIPS 2020. 4th KR2ML Workshop*. URL: https://api.semanticscholar.org/CorpusID:227528201.

Zachary Horvitz, Nam Do, and Michael L. Littman (July 2020). "Context-Driven Satirical News Generation". In: *FIGLANG ACL Workshop*. URL: https://api.semanticscholar.org/CorpusID:220330989.

Siegel, Noah, **Zachary Horvitz**, Roie Levin, Santosh Kumar Divvala, and Ali Farhadi (2016). "FigureSeer: Parsing Result-Figures in Research Papers". In: *European Conference on Computer Vision*. URL: https://api.semanticscholar.org/CorpusID:7857660.

Teaching

Contributor - Deep Learning with Tensorflow (Codecademy) 2020

– Developed content for Codecademy's first deep learning course.
– Materials included assignments, articles, and videos for over 14,000 students.

Head Teaching Assistant - Deep Learning (Brown University) 2019
Teaching Assistant - NLP (Columbia University) 2024
Teaching Assistant - Computational Linguistics (Brown University) 2018 - 2019

Awards / Selected Activities

CAIT Fellowship Recipient 2024

Columbia Center of Artificial Intelligence Technology

Outstanding Paper Award (ACL 2024) 2024

For *Getting Serious about Humor: Crafting Humor Datasets with Unfunny Large Language Models*

CS Department Rep, Eng. Council 2023 - 2024
Columbia University

Best in General Anthropology 2019
Brown University

CS2951-O Transport Logistics, 1st Place 2020
Brown University

Senior Staff Writer/Staff Writer 2016 - 2020
The Brown Noser

Associate Director/Associate, Data Board 2017 - 2019
Brown Political Review

Skills

Programming Languages	Python, Matlab, Scala, Java, C/C++, LaTeX, SQL, Javascript
Tools	PyTorch, TensorFlow, Caffe, Scikit-Learn, NLTK, SpaCy, Horovod, Docker, Huggingface, AWS, Flask, Git, ...