

Ari Holtzman

419 NE 71st St.,
Seattle, WA 98115
+1 (650) 228 6145

aholtzman@uchicago.edu
<http://ariholtzman.com>

ACADEMIC POSITIONS

Acting Primary Investigator, University of Chicago 2023 - 2024
Department of Computer Science & The Data Science Institute
Research Area: Machine Communication

Assistant Profesor, University of Chicago starting 2024
Department of Computer Science & The Data Science Institute
Toyota Technological Institute at Chicago (courtesy appointment)
Research Area: Machine Communication

EDUCATION

Doctor of Philosophy, University of Washington, Seattle, WA 2016 - 2023
Paul G. Allen School of Computer Science & Engineering
Research Areas: Natural Language Processing, Machine Learning
Advisor: Luke Zettlemoyer
Thesis: Interpretation Errors: Extracting Functionality from Generative Models of Language by Understanding Them Better

Bachelor of Arts, New York University 2012 - 2015
Individualized Major Information Theory: Bridging the Sciences and the Humanities
magna cum laude
Advisor: Peter Rajsingh

PUBLICATIONS & PREPRINTS

QLoRA: Efficient Finetuning of Quantized LLMs
Tim Dettmers, Artidoro Pagnoni, **Ari Holtzman**, Luke Zettlemoyer
In Submission. 2023.

Generative Models as a Complex Systems Science: How can we make sense of large language model behavior?
Ari Holtzman, Peter West, Luke Zettlemoyer
Self-published. 2023.

Artificial Intelligence and Aesthetic Judgment
Jessica Hullman, **Ari Holtzman**, Andrew Gelman
Self-published. 2023.

Toward Human Readable Prompt Tuning: Kubrick’s The Shining is a good movie, and a good prompt too?
Weijia Shi, Xiaochuang Han, Hila Gonen, **Ari Holtzman**, Yulia Tsvetkov, Luke Zettlemoyer
In Submission. 2023.

Contrastive Decoding: Open-ended Text Generation as Optimization
Xiang Lisa Li, **Ari Holtzman**, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, Mike Lewis
Association for Computational Linguistics (ACL). 2023.

What Do NLP Researchers Believe? Results of the NLP Community Metasurvey
Julian Michael, **Ari Holtzman**, Alicia Parrish, Aaron Mueller, Alex Wang, Angelica Chen, Divyam Madaan, Nikita Nangia, Richard Yuanzhe Pang, Jason Phang, Samuel R. Bowman
Association for Computational Linguistics (ACL). 2023.

Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?

Sewon Min, Xinxì Lyu, **Ari Holtzman**, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer
Empirical Methods in Natural Language Processing (EMNLP). 2022.

DEMix layers: Disentangling Domains for Modular Language Modeling
Suchin Gururangan, Mike Lewis, **Ari Holtzman**, Noah A Smith, Luke Zettlemoyer
Association for Computational Linguistics (ACL). 2022.

Surface Form Competition: Why the Highest Probability Answer Isn't Always Right
Ari Holtzman*, Peter West*, Vered Shwartz, Yejin Choi, Luke Zettlemoyer
Empirical Methods in Natural Language Processing (EMNLP). 2021.

Clipscore: A Reference-free Evaluation Metric for Image Captioning
Jack Hessel, **Ari Holtzman**, Maxwell Forbes, Ronan Le Bras, Yejin Choi
Empirical Methods in Natural Language Processing (EMNLP). 2021.

PIGLeT: Language Grounding Through Neuro-Symbolic Interaction in a 3D World
Rowan Zellers, **Ari Holtzman**, Matthew Peters, Roozbeh Mottaghi, Aniruddha Kembhavi, Ali Farhadi, Yejin Choi
Association for Computational Linguistics (ACL). 2021.

TuringAdvice: A Generative and Dynamic Evaluation of Language Use
Rowan Zellers, **Ari Holtzman**, Elizabeth Clark, Lianhui Qin, Ali Farhadi, Yejin Choi.
North American Chapter of the Association for Computational Linguistics (NAACL). 2021.

Multitalk: A Highly-branching Dialog Testbed for Diverse Conversations
Yao Dou, Maxwell Forbes, **Ari Holtzman**, Yejin Choi
Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). 2021.

Reflective Decoding: Beyond Unidirectional Generation with Off-the-Shelf Language Models
Peter West, Ximing Lu, **Ari Holtzman**, Chandra Bhagavatula, Jena Hwang, Yejin Choi
Association for Computational Linguistics (ACL). 2021.

The Curious Case of Neural Text Degeneration
Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi.
International Conference on Learning Representations (ICLR). 2020.

Experience Grounds Language
Yonatan Bisk*, **Ari Holtzman***, Jesse Thomason*, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, Joseph Turian.
Empirical Methods in Natural Language Processing (EMNLP). 2020.

Abductive Commonsense Reasoning
Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, **Ari Holtzman**, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, and Yejin Choi.
International Conference on Learning Representations (ICLR). 2020.

Defending Against Neural Fake News
Rowan Zellers, **Ari Holtzman**, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi.
Neural Information Processing Systems (NeurIPS). 2019.

BottleSum: Unsupervised and Self-supervised Sentence Summarization Using the Information Bottleneck Principle

Peter West, **Ari Holtzman**, Jan Buys, and Yejin Choi.
Empirical Methods in Natural Language Processing (EMNLP). 2019.

Counterfactual Story Reasoning and Generation
Lianhui Qin, Antoine Bosselut, **Ari Holtzman**, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi.
Empirical Methods in Natural Language Processing (EMNLP). 2019.

Do Neural Language Models Learn Physical Commonsense?
Maxwell Forbes, **Ari Holtzman**, and Yejin Choi.
Cognitive Science. 2019.

HellaSwag: Can a Machine Really Finish Your Sentence?
Rowan Zellers, **Ari Holtzman**, Yonatan Bisk, and Yejin Choi.
Association for Computational Linguistics (ACL). 2019.

Tactical Rewind: Self-correction via Backtracking in Vision-and-Language Navigation
Liyiming Ke, Xiujun Li, Yonatan Bisk, **Ari Holtzman**, Zhe Gan, Jingjing Liu, Jianfeng Gao, Yejin Choi, and Siddhartha Srinivasa.
Computer Vision and Pattern-Recognition (CVPR). 2019.

Sounding Board: A User-centric and Content-driven Social Chatbot
Hao Fang, Hao Cheng, Maarten Sap, Elizabeth Clark, **Ari Holtzman**, Yejin Choi, Noah A Smith, and Mari Ostendorf.
North American Chapter of the Association for Computational Linguistics (NAACL): Demonstrations. 2018.

Learning to Write with Cooperative Discriminators
Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi.
Association for Computational Linguistics (ACL). 2018.

Simulating Action Dynamics with Neural Process Networks
Antoine Bosselut, Omer Levy, **Ari Holtzman**, Corin Ennis, Dieter Fox, and Yejin Choi.
International Conference on Learning Representations (ICLR). 2018.

Connotation Frames of Power and Agency in Modern Films
Maarten Sap, Marcella Cindy Prasettio, **Ari Holtzman**, Hannah Rashkin, and Yejin Choi.
Empirical Methods in Natural Language Processing (EMNLP). 2017.

PRESENTATIONS	<i>What are language models actually modeling?</i>	
	International Conference of Social Computing (Keynote)	2023
	<i>NLP as a Complex Systems Science</i>	
	University of Washington NLP Retreat	2022
	<i>Surface Form Competition: Why the Highest Probability Answer Isn't Always Right</i>	
	Empirical Methods in Natural Language Processing (EMNLP)	2021
	<i>Clipscore: A Reference-free Evaluation Metric for Image Captioning</i>	
	Empirical Methods in Natural Language Processing (EMNLP)	2021
	<i>Beyond Unidirectional Generation with Off-the-Shelf Language Models</i>	
	Association for Computational Linguistics (ACL)	2021

	<i>The Curious Case of Neural Text Degeneration</i>	
	International Conference on Learning Representations (ICLR)	2020
	DARPA Communicating with Computers Meeting	2019
	University of Washington NLP Retreat	2019
	<i>Experience Grounds Language</i>	
	<i>Empirical Methods in Natural Language Processing (EMNLP)</i>	2020
	<i>Defending Against Neural Fake News</i>	
	Neural Information Processing Systems (NeurIPS)	2019
	<i>Do Neural Language Models Learn Physical Commonsense?</i>	
	Annual Meeting of the Cognitive Science Society	2019
	<i>Tactical Rewind: Self-correction via Backtracking in Vision-and-Language Navigation</i>	
	Computer Vision and Pattern-Recognition (CVPR).	2019
	<i>Learning to Write with Cooperative Discriminators</i>	
	Association for Computational Linguistics Conference	2018
	Samsung Research Seminar	2018
	<i>Connotation Frames of Power and Agency in Modern Films</i>	
	Empirical Methods in Natural Language Processing (EMNLP)	2017
AWARDS & HONORS	Amazon Alexa Prize Winner	2017
	Phi Beta Kappa	2015
	NYU University Honors Scholar	2015
	Dean's Honor List	2012 - 2015
TEACHING	Teaching Assistant	
	Natural Language Processing (Capstone)	2018
	Natural Language Processing (lectured for recitation)	2018
INDUSTRY EXPERIENCE	Post-Doctoral Researcher	2023 - 2024
	Meta	Seattle, WA
	Research on how Large Language Models can be integrated into virtual environments such as video games, to explore capabilities and error analysis of LLM social coordination capacity.	
	Research Intern	2018 - 2019
	Allen Institute for Artificial Intelligence	Seattle, WA
	Research on integrating commonsense into natural language generation and reasoning about social and physical causality in narrative text.	
	Software Engineer	2016
	Brave	San Francisco, CA
	Prototyped system for gauging user advertising preferences.	
	Education Curriculum Intern	2015
	CK-12 Foundation	Palo Alto, CA
	Designed and consulted on systems that allow teachers to build their own interactive demos without coding knowledge.	
	Research Assistant	2015
	Courant Institute of Mathematical Sciences, New York University	New York, NY

Coded a transpiler for GRASShopper verification tool that compiles an academic research language into compilable C code.

Software Engineering Intern 2014

Redshift Technologies

New York, NY

Built new features for North American Butterfly Association site. Helped design a chemical data-sharing platform. Various other tasks including planning the merging of two websites and their infrastructures.

Software Engineer 2014

I Love Velvet, Inc.

New York, NY

Built model Android App used at trade shows and assisted with various technical challenges.

Research Assistant 2014

Digital Experiments at NYU

New York, NY

Designed systems for using programs and statistical measurements to determine a text's relationships within a community of texts. Provided technical guidance to a team of people with varying levels of technical knowledge.

**PROGRAM
COMMITTEE &
REVIEWING**

Program Committee/Reviewer (Conference):

ACL Roling Review	2021-2023
North American Association for Computational Linguistics (NAACL)	2019-2023
Association for Computational Linguistics (ACL)	2017-2023
Empirical Methods in Natural Language Processing (EMNLP)	2017-2023
International Conference on Learning Representations (ICLR)	2021-2022
Neural Information Processing Systems (NeurIPS)	2020
AAAI Conference on Artificial Intelligence (AAAI)	2018
International Conference on Computational Linguistics (COLING)	2018

Program Committee/Reviewer (Workshop)

NAACL Workshop on Optimizing and Evaluating Neural Generation (NeuralGen)	2019
NAACL Workshop on Narrative Understanding (WNU)	2019-2020
ACL Joint Workshop on Narrative Understanding, Storylines, and Events	2020
EMNLP SustaiNLP Workshop	2020