

William B. Held

MACHINE LEARNING (ML) & NATURAL LANGUAGE PROCESSING (NLP)

✉ (+1) 646-961-2627 | ✉ wheld3@gatech.edu | 🏷 www.williamheld.com | 📡 Helw150 | 🌐 williambarrheld

Education

Georgia Institute of Technology

3.84 GPA Ph.D. in Machine Learning

Atlanta, GA, United States

Aug. 2021 - PRESENT

- Advised by Prof. Diyi Yang; Research on inclusive & robust Large Language Models
- NSF Graduate Research Fellowship Program Honorable Mention Recipient.

New York University Abu Dhabi

3.73 GPA B.S. in Computer Science with a concentration in Economics

Abu Dhabi, United Arab Emirates

Aug. 2015 - May 2019

- Research focused on semantic relation extraction and Arabic computational linguistics under Prof. Nizar Habash

Highlighted Work

Distilling an End-to-End Voice Assistant from Speech Recognition Data

William Held, Ella Li, Michael Ryan, Weiyan Shi, Yanzhe Zhang, Diyi Yang

63rd Annual Meeting of the Association for Computational Linguistics, Oral (Top 5% of Papers)

2025

Optimizing Pretraining Data Mixtures with LLM-Estimated Utility

William Held, Bhargavi Paranjape, Punit Singh Koura, Mike Lewis, Frank Zhang, Todor Mihaylov

Preprint w/ Meta GenAI, Part of Llama 4 Research

2024

Multi-VALUE: A Framework for Cross-Dialectal English NLP

William Held[🔥], Caleb Ziems[🔥], Jingfeng Yang, Jwala Dhamala, Rahul Gupta, Diyi Yang

Equal Contribution

61st Annual Meeting of the Association for Computational Linguistics (20.8% Acceptance Rate)

2023

Can Large Language Models Transform Computational Social Science?

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, Diyi Yang

Computational Linguistics, Volume 50

2023

Professional Experience

Meta AI

Research Scientist Intern - GenAI Llama Data Team

Menlo Park, CA, United States

Summer 2023 & Summer 2024

- Worked on methods for data mix optimization for Large Language Model Pretraining.
- Contributed to both Llama 3 and Llama 4 LLMs

Google Research

Research Intern - Neural Semantic Parsing

New York, NY, United States

Summer 2022

• Worked on methods to help extend Google Assistant's Multilingual capabilities. Work published in our paper: "DAMP".

Sunshine/Lumi Labs

Software Engineer - Machine Learning

Palo Alto, CA, United States

June 2019 - June 2021

- Worked with Prof. Dan Jurafsky on coreference resolution. Work published in our paper: "Focus on What Matters".
- 9th engineer hired to build out AI-powered applications focused on contact management

Foursquare

Machine Learning Engineer Intern - Pilgrim Data Team

New York, NY, United States

Summer 2018

Knotch

Data Scientist - Knowledge Search Engine

New York, NY, United States

May 2017 - February 2018

Quorum Analytics

Software Engineering Intern - User Accounts and Permissions

Washington, D.C., United States

Spring 2017

Sandia National Laboratories

Research Intern - Computer Science Research Institute

Albuquerque, NM, United States

Summer 2015 & Summer 2016

Service

Reviewer ACL 2022-, NeurIPS 2022- ICLR 2024-, ICML 2023-

Organizer 2021 Georgia Tech Graduate Application Support (GT-GAS)

Volunteer Coordinator The 2021 Conference on Empirical Methods in Natural Language Processing

Languages

Programming Scala(Spark), Python(PyTorch, Jax), Javascript(React) **Natural** English, Arabic

All Publications

Marin: An Open Lab for Building Foundation Models

David Hall, Ahmed Ahmed, Christopher Chou, Abhinav Garg, Rohith Kuditipudi, **Will Held**, ... Percy Liang

Open-Source Software

2025

Distilling an End-to-End Voice Assistant from Speech Recognition Data

William Held, Ella Li, Michael Ryan, Weiyan Shi, Yanzhe Zhang, Diyi Yang

63rd Annual Meeting of the Association for Computational Linguistics, Oral (Top 5% of Papers)

2025

Mind the gap! static and interactive evaluations of large audio models

William Held[🔥], Ella Li[🔥], Michael J Ryan, Kunat Pipatanakul, Potsawee Manakul, Hao Zhu, Diyi Yang

63rd Annual Meeting of the Association for Computational Linguistics, Oral (Top 5% of Papers)

2025

SynthesizeMe! Inducing Persona-Guided Prompts for Personalized Reward Models in LLMs

Michael J Ryan, Omar Shaikh, Aditri Bhagirath, Daniel Frees, **William Held**, Diyi Yang

63rd Annual Meeting of the Association for Computational Linguistics

2025

Optimizing Pretraining Data Mixtures with LLM-Estimated Utility

William Held, Bhargavi Paranjape, Punit Singh Koura, Mike Lewis, Frank Zhang, Todor Mihaylov

Preprint

2024

Perceptions of Language Technology Failures from South Asian English Speakers

William Held[🔥], Faye Holt[🔥], Diyi Yang

Findings of the Association for Computational Linguistics ACL 2024

2024

Measuring and Addressing Indexical Bias in Information Retrieval

Caleb Ziems, **William Held**, Jane Dwivedi-Yu, Diyi Yang

Findings of the Association for Computational Linguistics ACL 2024

2024

Unintended Impacts of LLM Alignment on Global Representation

Michael Ryan, **William Held**, Diyi Yang

Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics

2024

- Main, Long Paper Acceptance Rate: 23.5%

A Material Lens on Coloniality in NLP

William Held, Camille Harris, Michael Best, Diyi Yang

Pre-Print

2023

Task-Agnostic Low-Rank Adapters for Unseen English Dialects

Zedian Xiao, **William Held**, Yanchen Liu, Diyi Yang

The 2023 Conference on Empirical Methods in Natural Language Processing

2023

DADA: Dialect Adaptation via Dynamic Aggregation of Linguistic Rules

Yanchen Liu, **William Held**, Diyi Yang

The 2023 Conference on Empirical Methods in Natural Language Processing

2023

TADA : Task Agnostic Dialect Adapters for English

William Held, Caleb Ziems, Diyi Yang

61st Annual Meeting of the Association for Computational Linguistics

2023

- Findings, Short Paper Acceptance Rate: 35.58%

Can Large Language Models Transform Computational Social Science?

Caleb Ziems, **William Held**, Omar Shaikh, Jiaao Chen, Zhehao Zhang, Diyi Yang

Journal of Computational Linguistics

2023

Modeling Cross-Cultural Pragmatic Inference with Codenames Duet

Omar Shaikh, Caleb Ziems, **William Held**, Aryan Pariani, Fred Morstatter, Diyi Yang

61st Annual Meeting of the Association for Computational Linguistics

2023

- Findings, Long Paper Acceptance Rate: 41.89%

Multi-VALUE: A Framework for Cross-Dialectal English NLP

William Held[🔥], Caleb Ziems[🔥], Jingfeng Yang, Jwala Dhamala, Rahul Gupta, Diyi Yang

 Equal Contribution

2023

61st Annual Meeting of the Association for Computational Linguistics

- Main Conference, Long Paper Acceptance Rate: 20.75%

DAMP: Doubly Aligned Multilingual Parser for Task-Oriented Parsing

William Held, Chris Hidey, Fei Liu, Eric Zhu, Rahul Goel, Diyi Yang, Rushin Shah

61st Annual Meeting of the Association for Computational Linguistics

2023

- Main Conference, Long Paper Acceptance Rate: 20.75%

On Second Thought, Let's Not Think Step by Step! Bias and Toxicity in Zero-Shot Reasoning

Omar Shaikh, Hongxin Zhang, **William Held**, Michael Bernstein, Diyi Yang

61st Annual Meeting of the Association for Computational Linguistics

2023

- Main Conference, Long Paper Acceptance Rate: 20.75%

Shapley Head Pruning: Identifying and Removing Interference in Multilingual Transformers

William Held, Diyi Yang

17th Conference of the European Chapter of the Association for Computational Linguistics

2022

- Main Conference, Long Paper Acceptance Rate: 24.7%

Focus on what matters: Applying Discourse Coherence Theory to Cross Document Coreference

William Held, Dan Iter, Dan Jurafsky

Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing

2021

- Selected for Oral Presentation, Long Paper Acceptance Rate: 24.6%

The Effectiveness of Simple Hybrid Systems for Hypernym Discovery

William Held, Nizar Habash

Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics

2019

- Short Paper Acceptance Rate: 18.2%

Transitioning Green-Gauss Gradients to the Kokkos Framework

William Held, Andrew Bradley

Sandia National Labs Center for Computing Research Summer Proceedings

2016

A Testing Framework for a Hybrid Triangular Solver

William Held, Andrew Bradley

Sandia National Labs Center for Computing Research Summer Proceedings

2015