

William B. Held

MACHINE LEARNING & NATURAL LANGUAGE PROCESSING

☎ (+1) 646-961-2627 | ✉ wheld3@gatech.edu | 🏠 www.williamheld.com | 📱 Helw150 | 📧 williambarrheld

Education

Georgia Institute of Technology

Atlanta, GA, United States

3.8 GPA Ph.D. in Machine Learning

Aug. 2021 - PRESENT

- Research focused on natural language processing for low-resource languages and dialects advised by Prof. Diyi Yang
- NSF Graduate Research Fellowship Program Honorable Mention Recipient.

New York University Abu Dhabi

Abu Dhabi, United Arab Emirates

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

Aug. 2015 - May 2019

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

Recent Work

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

Caleb Ziems 🔥, **William Held** 🔥, Jingfeng Yang, Diyi Yang

🔥 Equal Contribution

Pre-Print & In-Submission

2022

DAMP: Doubly Aligned Multilingual Parser for Task-Oriented Parsing

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

Pre-Print & In-Submission

2022

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

Pre-Print & In-Submission

2022

Shapley Head Pruning: Identifying and Removing Interference in Multilingual Transformers

William Held, Diyi Yang

Pre-Print & In-Submission

2022

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%

Professional Experience

Google Research

New York, NY, United States

Research Intern - Neural Semantic Parsing

May - August 2022

- Developed adversarial learning techniques for Multilingual and Codeswitched semantic parsing
- Studied explicit alignment pretraining objectives for Massively Multilingual Transformer Language Models

Sunshine/Lumi Labs

Palo Alto, CA, United States

Software Engineer - Machine Learning

June 2019 - June 2021

- Engineer assigned to research partnership with Prof. Dan Jurafsky co-reference resolution
- 9th engineer hired to build out AI-powered applications focused on contact management

Foursquare

New York, NY, United States

Machine Learning Engineer Intern - Pilgrim Data Team

May 2018 - Aug 2018

Knotch

New York, NY, United States

Data Scientist - Knowledge Search Engine

May 2017 - February 2018

Quorum Analytics

Washington, D.C., United States

Software Engineering Intern - User Accounts and Permissions

January - May 2017

Sandia National Laboratories

Albuquerque, NM, United States

Research Intern - Computer Science Research Institute

April - August 2015 & 2016

Service

NeurIPS **Reviewer**, Thirty-sixth Conference on Neural Information Processing Systems

2022

EACL **Reviewer**, The 17th Conference of the European Association for Computational Linguistics

2022

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

2021

GT **Organizer**, Georgia Tech Graduate Application Support (GT-GAS)

2021

Languages

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

All Publications

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

Caleb Ziems , William Held , Jingfeng Yang, Diyi Yang

 Equal Contribution

Pre-Print & In-Submission

2022

DAMP: Doubly Aligned Multilingual Parser for Task-Oriented Parsing

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

Pre-Print & In-Submission

2022

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

Pre-Print & In-Submission

2022

Shapley Head Pruning: Identifying and Removing Interference in Multilingual Transformers

William Held, Diyi Yang

Pre-Print & In-Submission

2022

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