

Jacob Morrison

📞 206-504-9132 | ✉️ jacobm@allenai.org | 🏠 jacobmorrison.com

Education

University of Washington

M.S. in Computational Linguistics

- Advisor: Noah Smith

Seattle, WA

2021 - 2022

University of Washington

B.S. in Computer Science

- Minor in Mathematics

Seattle, WA

2013 - 2017

Awards & Fellowships

2024 **Theme Paper Award**, ACL 2024

2024 **Best Resource Paper Award**, ACL 2024

2023 **National Science Foundation**, Computer Science Graduate Fellowship

Experience

Allen Institute for Artificial Intelligence

Predoctoral Young Investigator

- Mentors: Pradeep Dasigi & Jesse Dodge

Seattle, WA

Jan 2023 - Present

X, the moonshot factory

AI Resident

- Confidential research applying large language models to problems in code generation & program synthesis

Mountain View, CA

Summer 2022

Twitter

Data Scientist Intern

- Investigated negative feedback loops in training and deployment of content moderation machine learning models

Seattle, WA

Summer 2021

Allen Institute for Artificial Intelligence

Research Engineer Intern

- Implemented language + vision models into the AllenNLP platform

Seattle, WA

Spring 2021

Google

Software Engineer

Kirkland, WA

Jan 2019 - Feb 2021

Tableau

Software Engineer

Seattle, WA

Sep 2017 - Dec 2018

Publications

OLMo 2: Pushing Frontiers in Open Language Model Pretraining

Forthcoming, 2024

Tulu 3: Pushing Frontiers in Open Language Model Post-Training

Nathan Lambert, **Jacob Morrison**, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, Hannaneh Hajishirzi

arXiv preprint, 2024

Holistically Evaluating the Environmental Impact of Creating Language Models

Jacob Morrison, Clara Na, Jared Fernandez, Tim Dettmers, Emma Strubell, Jesse Dodge

Under review, ICLR 2025

OLMoE: Open Mixture-of-Experts Language Models

Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, **Jacob Morrison**, Sewon Min, Weijia Shi, Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, Ali Farhadi, Noah A. Smith, Pang Wei Koh, Amanpreet Singh, Hannaneh Hajishirzi

NeurIPS 2024 Workshop on Efficient Natural Language and Speech Processing (spotlight)

Under review, ICLR 2025

Merge to Learn: Efficiently Adding Skills to Language Models with Model Merging

Jacob Morrison, Noah A. Smith, Hannaneh Hajishirzi, Pang Wei Koh, Jesse Dodge, Pradeep Dasigi

Findings of EMNLP 2024

SciRIF: A Resource to Enhance LM Instruction-Following over Scientific Literature

David Wadden*, Kejian Shi*, **Jacob Morrison**, Aakanksha Naik, Shruti Singh, Nitzan Barzilay, Kyle Lo, Tom Hope, Luca Soldaini, Shannon Shen, Doug Downey, Hanna Hajishirzi, Arman Cohan

NeurIPS 2024 Workshop on Foundation Models for Science

Under review, NAACL 2025

RewardBench: Evaluating Reward Models for Language Modeling

Nathan Lambert, Valentina Piatkin, **Jacob Morrison**, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Tom Zick, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi

Under review, NAACL 2025

OLMo: Accelerating the Science of Language Models

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, **Jacob Morrison**, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, Hannaneh Hajishirzi

*ACL 2024 **Theme Paper Award***

Dolma: An Open Corpus of Three Trillion Tokens for Language Model Pretraining Research

Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxu Lyu, Nathan Lambert, Ian Magnusson, **Jacob Morrison**, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Evan Pete Walsh, Hannaneh Hajishirzi, Noah A. Smith, Luke Zettlemoyer, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, Kyle Lo

*ACL 2024 **Best Resource Paper Award***

Intentionally Unintentional Speech: Why Generative AI Models Are Not Protected by the First Amendment

David Atkinson, Jena D. Hwang, **Jacob Morrison**

First Amendment Law Review (University of North Carolina), 2024

Unsettled Law in the Age of Generative AI: Time to Generate New Approaches?

David Atkinson, **Jacob Morrison**

Journal of Law and Technology at Texas, 2024

A Legal Risk Taxonomy for Generative Artificial Intelligence

David Atkinson, **Jacob Morrison**

arXiv report, 2024

Bidimensional Leaderboards: Generate and Evaluate Language Hand in Hand

Jungo Kasai, Keisuke Sakaguchi, Ronan Le Bras, Lavinia Dunagan, **Jacob Morrison**, Alexander Fabbri, Yejin Choi, Noah A. Smith

NAACL 2022

Transparent Human Evaluation for Image Captioning

Jungo Kasai, Keisuke Sakaguchi, Lavinia Dunagan, **Jacob Morrison**, Ronan Le Bras, Yejin Choi, Noah A. Smith

NAACL 2022

* indicates equal contribution

Invited Talks & Presentations

Washington State House Consumer Protection & Business Committee

July 2024

Technical and Policy Considerations for Generative AI

Morgan Stanley

July 2024

Legal Risks for Generative AI

Bill & Melinda Gates Foundation

May 2024

OLMo: A Truly Open Language Model

State of Washington AI Community of Practice

Nov 2023

Responsible, Real World AI Development and Deployment

Press

City of Seattle Releases Generative Artificial Intelligence Policy for City Employees

Nov 2023

Office of the Mayor of Seattle

Tech policy leaders explore new frontiers for supporting (and regulating) AI at forum

Aug 2023

GeekWire

• Organized the Future of AI Forum with U.S. Senator Maria Cantwell's office

Teaching Experience

University of Washington, Computer Science & Engineering

Seattle, WA

Head Teaching Assistant

Fall 2022

- Lead a team of five TAs for a natural language processing course with over fifty undergraduate and graduate students
- Coordinated and managed all grading, office hours and guest lectures

University of Washington, Mathematics Department

Seattle, WA

Grading Assistant

Mar 2014 - Jun 2016

- Graded weekly differential equations, linear analysis and multivariable calculus assignments
- Calculated, recorded and ensured the accuracy of grades for 100-200 students per quarter

University of Washington, Computer Science & Engineering

Seattle, WA

Teaching Assistant

Fall 2015

- Taught weekly sections, covering introductory programming concepts and departmental best practices
- Created lesson plans and grading criteria, and set points of emphasis with other TAs and instructors
- Received an above average rating, including in teaching effectiveness and feedback quality

Service

Reviewer

CoLM (2024), ICML GenLaw Workshop (2024), ACL Rolling Review (Jun, Oct 2024), ICLR (2025)

State of Washington

Olympia, WA

Artificial Intelligence Task Force Education & Workforce Development Subcommittee Member

Aug 2024 - Present

City of Seattle

Seattle, WA

Generative AI Policy Advisory Group Member

Jun 2023 - Aug 2023

- Advised the City of Seattle on the creation of its Generative AI Policies