

Jacob Morrison

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

Education

University of Washington

Ph.D. in Computer Engineering

• Advisor: Banghua Zhu

Seattle, WA

2025 - Present

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

- 2025 **Mercor**, Graduate Fellowship Finalist
- 2025 **Massachusetts Institute of Technology**, EECS Merrill Lynch Ph.D. Fellowship (Declined)
- 2025 **University of California, Berkeley**, EECS Ph.D. Fellowship (Declined)
- 2025 **Carnegie Mellon University**, LTI Ph.D. Fellowship (Declined)
- 2025 **University of Toronto**, Departmental Ph.D. Fellowship; Vector Research Grant (Declined)
- 2024 **ACL 2024**, Theme Paper Award
- 2024 **ACL 2024**, Best Resource Paper Award
- 2024 **GeekWire**, Innovation of the Year Award
- 2023 **National Science Foundation**, Computer Science Graduate Fellowship

Experience

University of Washington

Ph.D. Student

• Advisor: Banghua Zhu

Seattle, WA

Sep 2025 - Present

Allen Institute for Artificial Intelligence

Student Researcher

• Host: Hannaneh Hajishirzi

Predoctoral Young Investigator

• Mentors: Pradeep Dasigi & Jesse Dodge

Seattle, WA

Sep 2025 - Present

Jan 2023 - Sep 2025

X, the moonshot factory

AI Resident

Mountain View, CA

Summer 2022

Twitter

Data Scientist Intern

Seattle, WA

Summer 2021

Allen Institute for Artificial Intelligence

Research Engineer Intern

Seattle, WA

Spring 2021

Google

Software Engineer

Kirkland, WA

Jan 2019 - Feb 2021

Tableau

Software Engineer

Seattle, WA

Sep 2017 - Dec 2018

Publications

FlexOLMo: Open Language Models for Flexible Data Use

Weijia Shi, Akshita Bhagia, Kevin Farhat, Niklas Muennighoff, Evan Pete Walsh, **Jacob Morrison**, Dustin Schwenk, Shayne Longpre, Jake Poznanski, Allyson Ettinger, Daogao Liu, Margaret Li, Mike Lewis, Wen-tau Yih, Dirk Groeneveld, Luca Soldaini, Kyle Lo, Noah A. Smith, Luke Zettlemoyer, Pang Wei Koh, Hannaneh Hajishirzi, Ali Farhadi, Sewon Min

Under review, NeurIPS 2025

RewardBench 2: Advancing Reward Model Evaluation

Saumya Malik, Valentina Pyatkin, Sander Land, **Jacob Morrison**, Noah A. Smith, Hannaneh Hajishirzi, Nathan Lambert

Under review, NeurIPS 2025

2 OLMo 2 Furious

OLMo Team, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, **Jacob Morrison**, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, Hannaneh Hajishirzi
COLM 2025

Tülu 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
COLM 2025

Holistically Evaluating the Environmental Impact of Creating Language Models

Jacob Morrison, Clara Na, Jared Fernandez, Tim Dettmers, Emma Strubell, Jesse Dodge
ICLR 2025 (spotlight)

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)
ICLR 2025 (oral)

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

SciRIFF: 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

RewardBench: Evaluating Reward Models for Language Modeling

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

North American Summer School for Logic, Language and Information AI for Legal Interpretation?	June 2025
Washington State House Consumer Protection & Business Committee Technical and Policy Considerations for Generative AI	July 2024
Morgan Stanley Legal Risks for Generative AI	July 2024
Bill & Melinda Gates Foundation OLMo: A Truly Open Language Model	May 2024
State of Washington AI Community of Practice Responsible, Real World AI Development and Deployment	Nov 2023

Press

A new kind of AI model gives data owners more control Wired	July 2025
Your AI models are failing in production – Here’s how to fix model selection VentureBeat	June 2025
City of Seattle Releases Generative Artificial Intelligence Policy for City Employees Office of the Mayor of Seattle	November 2023
Tech policy leaders explore new frontiers for supporting (and regulating) AI at forum GeekWire	August 2023
• Organized the Future of AI Forum with U.S. Senator Maria Cantwell’s office	

Teaching Experience

University of Washington, Computer Science & Engineering Head Teaching Assistant	Seattle, WA 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 Grading Assistant	Seattle, WA 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 Teaching Assistant	Seattle, WA 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, 2025) ACL Rolling Review (Jun 2024 - Present) ICLR (2025) ICML GenLaw Workshop (2024)	
State of Washington Artificial Intelligence Task Force Education & Workforce Development Subcommittee Member	Olympia, WA Aug 2024 - Present
City of Seattle Generative AI Policy Advisory Group Member	Seattle, WA Jun 2023 - Aug 2023
• Advised the City of Seattle on the creation of its Generative AI Policies	