

# Sewon Min

Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley  
2121 Berkeley Way, Berkeley, CA 94704

Home: [sewonmin.com](http://sewonmin.com)  
Email: [sewonm@berkeley.edu](mailto:sewonm@berkeley.edu)  
Github: [/shmsw25](https://github.com/shmsw25)

## Current Employment

---

07/2025– Assistant Professor, UC Berkeley EECS  
07/2024– Research scientist, Allen Institute for AI

## Education

---

09/2018–06/2024 Ph.D., Computer Science & Engineering, University of Washington  
Supervisors: Hannaneh Hajishirzi, Luke Zettlemoyer  
Thesis: *Rethinking Data Use in Large Language Models*  
Committee: Noah Smith, Shane Steinert-Threlkeld, Graham Neubig

03/2020 M.Sc., Computer Science & Engineering, University of Washington

03/2014–08/2018 B.Sc., Computer Science & Engineering, Seoul National University  
Summa Cum Laude, Rank: 1st in the department

## Past Employment

---

09/2018–06/2024 Research assistant, University of Washington  
Supervisors: Hannaneh Hajishirzi, Luke Zettlemoyer

06/2023–12/2023, 04/2021–03/2023, 10/2019–08/2020 Visiting researcher/research intern, Meta AI (Facebook AI Research)  
Host: Luke Zettlemoyer

09/2020–03/2021 Research intern/student researcher, Google AI Research  
Hosts: Kenton Lee, Kristina Toutanova / Collaborator: Ming-Wei Chang

11/2017–02/2018 Research intern, Salesforce Research (formerly Metamind)  
Host: Caiming Xiong

10/2016–02/2017 Undergraduate research assistant, University of Washington  
Supervisor: Hannaneh Hajishirzi

07/2016–08/2016 Undergraduate intern, Vision & Learning Lab, Seoul National University  
Supervisor: Gunhee Kim

## Honors & Awards

---

2023 Microsoft Accelerate Foundation Models Research Grant

2023 Outstanding Paper, Regulatable ML Workshop @ NeurIPS 2023

2023 Best Paper Honorable Mentions, ACL 2023

2022 EECS Rising Stars

2022–2023 JP Morgan Ph.D. Fellowship

2018–2023 Doctoral Study Fellowship granted by Korea Foundation for Advanced Studies (KFAS)

2019, 2021 Outstanding Reviewer (EMNLP 2019, ACL 2021, EMNLP 2021)

2018–2019 Wissner-Slivka Endowed Graduate Fellowship granted by Paul G. Allen School

- 2018 Best Undergraduate Thesis Award granted by CSE, Seoul National University
- 2014–2018 Full Scholarship for all semesters attended at Seoul National University (merit-based)
- 2017 Internship Abroad Support Fund granted by Seoul National University
- 2017 Samsung Convergence Software Course Mentor Scholarship

## Publications

---

GScholar <https://scholar.google.com/citations?user=jU4IZs4AAAAAJ>

Semantic Scholar <https://www.semanticscholar.org/author/Sewon-Min/48872685>

\* denotes equal contributions

### Peer-Reviewed Conference/Journal Papers

- [C34] Michael Duan, Anshuman Suri, Niloofar Mireshghallah, **Sewon Min**, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, Hannaneh Hajishirzi. “Do Membership Inference Attacks Work on Large Language Models?.” In: *Proceedings of COLM*. 2024. [pdf]
- [C33] Jiacheng Liu, **Sewon Min**, Luke Zettlemoyer, Yejin Choi, Hannaneh Hajishirzi. “Infini-gram: Scaling Unbounded n-gram Language Models to a Trillion Tokens.” In: *Proceedings of COLM*. 2024. [pdf]
- [C32] Weijia Shi, **Sewon Min**, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, Wen-tau Yih. “REPLUG: Retrieval-Augmented Black-Box Language Models.” In: *Proceedings of NAACL*. 2024. [pdf]
- [C31] Weijia Shi, **Sewon Min**, Maria Lomeli, Chunting Zhou, Margaret Li, Xi Victoria Lin, Noah A. Smith, Luke Zettlemoyer, Wen-tau Yih, Mike Lewis. “In-Context Pretraining: Language Modeling Beyond Document Boundaries.” In: *Proceedings of ICLR*. 2024 (**Spotlight**). [pdf]
- [C30] Qingqing Cao, **Sewon Min**, Yizhong Wang, Hannaneh Hajishirzi. “BTR: Binary Token Representations for Efficient Retrieval Augmented Language Models.” In: *Proceedings of ICLR*. 2024 (**Spotlight**). [pdf]
- [C29] **Sewon Min**\*, Suchin Gururangan\*, Eric Wallace, Weijia Shi, Hannaneh Hajishirzi, Noah A. Smith, Luke Zettlemoyer. “SILO Language Models: Isolating Legal Risk in a Nonparametric Datastore.” In: *Proceedings of ICLR*. 2024 (**Spotlight**). [pdf]  
**Oral presentation at the Regulatable ML Workshop at NeurIPS 2023 (🏆 Outstanding paper).**  
**Spotlight presentation at the Workshop on Distribution Shifts (DistShift) at NeurIPS 2023.**
- [C28] **Sewon Min**\*, Kalpesh Krishna\*, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, Hannaneh Hajishirzi. “FACTScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation.” In: *Proceedings of EMNLP (long)*. 2023. **EMNLP 2023’s Most Influential Paper #4**. [pdf]  
**Oral presentation at the Third Workshop on the Generation, Evaluation & Metrics (GEM).**
- [C27] Ofir Press, Muru Zhang, **Sewon Min**, Ludwig Schmidt, Noah A. Smith, Mike Lewis. “Measuring and Narrowing the Compositionality Gap in Language Models.” In: *Proceedings of Findings of EMNLP (long)*. 2023. [pdf]
- [C26] Boshi Wang, **Sewon Min**, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, Huan Sun. “Towards Understanding Chain-of-Thought Prompting: An Empirical Study of What Matters.” In: *Proceedings of ACL (long)*. 2023. 🏆 **Best Paper Honorable Mentions. ACL 2023’s Most Influential Paper #13**. [pdf]
- [C25] Xinxu Lyu, **Sewon Min**, Iz Beltagy, Luke Zettlemoyer, Hannaneh Hajishirzi. “Z-ICL: Zero-Shot In-Context Learning with Pseudo-Demonstrations.” In: *Proceedings of ACL (long)*. 2023. [pdf]
- [C24] **Sewon Min**, Weijia Shi, Mike Lewis, Xilun Chen, Wen-tau Yih, Hannaneh Hajishirzi, Luke Zettlemoyer. “Nonparametric Masked Language Modeling.” In: *Proceedings of Findings of ACL (long)*. 2023. [pdf]

- [C23] Xinyan Velocity Yu, **Sewon Min**, Luke Zettlemoyer, Hannaneh Hajishirzi. “CREPE: Open-Domain Question Answering with False Presuppositions.” In: *Proceedings of ACL (long)*. 2023. [pdf]
- [C22] Zeqiu Wu, Ryu Parish, Hao Cheng, **Sewon Min**, Prithviraj Ammanabrolu, Mari Ostendorf, Hannaneh Hajishirzi. “INSCIT: Information-Seeking Conversations with Mixed-Initiative Interactions .” In: *TACL*. 2023. [pdf]
- [C21] Chenglei Si, Chen Zhao, **Sewon Min**, Jordan Boyd-Graber. “Re-Examining Calibration: The Case of Question Answering.” In: *Proceedings of Findings of EMNLP (long)*. 2022. [pdf]
- [C20] Anas Awadalla, Mitchell Wortsman, Gabriel Ilharco, **Sewon Min**, Hannaneh Hajishirzi, Ludwig Schmidt. “Exploring The Landscape of Distributional Robustness for Question Answering Models.” In: *Proceedings of Findings of EMNLP (long)*. 2022. [pdf]
- [C19] **Sewon Min**, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer. “Rethinking the Role of Demonstrations: What makes In-context Learning Work?” In: *Proceedings of EMNLP (long)*. 2022. *EMNLP 2022’s Most Influential Paper #1*. [pdf]
- [C18] Daniel Khashabi, Xinxi Lyu, **Sewon Min**, Lianhui Qin, Kyle Richardson, Sameer Singh, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, Yejin Choi. “Prompt Waywardness: The Curious Case of Discretized Interpretation of Continuous Prompts.” In: *Proceedings of NAACL-HLT (long)*. 2022. [pdf]
- [C17] **Sewon Min**, Mike Lewis, Luke Zettlemoyer, Hannaneh Hajishirzi. “MetaICL: Learning to Learn In Context.” In: *Proceedings of NAACL-HLT (long)*. 2022. *NAACL 2022’s Most Influential Paper #3*. [pdf]
- [C16] **Sewon Min**, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer. “Noisy Channel Language Model Prompting for Few-Shot Text Classification.” In: *Proceedings of ACL (long)*. 2022. *ACL 2022’s Most Influential Paper #12*. [pdf]
- [C15] Jungsoo Park\*, **Sewon Min**\*, Jaewoo Kang, Luke Zettlemoyer, Hannaneh Hajishirzi. “FAVIQ: FAct Verification from Information-seeking Questions.” In: *Proceedings of ACL (long)*. 2022. [pdf]
- [C14] **Sewon Min**, Kenton Lee, Ming-Wei Chang, Kristina Toutanova, Hannaneh Hajishirzi. “Joint Passage Ranking for Diverse Multi-Answer Retrieval.” In: *Proceedings of EMNLP (long)*. 2021. [pdf]
- [C13] **Sewon Min**, Jordan Boyd-Graber, Chris Alberti, Danqi Chen, Eunsol Choi, Michael Collins, Kelvin Guu, Hannaneh Hajishirzi, Kenton Lee, Jennimaria Palomaki, Colin Raffel, Adam Roberts, Tom Kwiatkowski (with EfficientQA participants). “NeurIPS 2020 EfficientQA Competition: Systems, Analyses and Lessons Learned.” In: *Proceedings of Maching Learning Research (PMLR)*. 2021. [pdf]
- [C12] Srinivasan Iyer, **Sewon Min**, Yashar Mehdad, Wen-tau Yih. “ReConsider: Re-Ranking using Span-Focused Cross-Attention for Open Domain Question Answering.” In: *Proceedings of NAACL-HLT (short)*. 2021. [pdf]
- [C11] Belinda Z. Li, **Sewon Min**, Srinivasan Iyer, Yashar Mehdad and Wen-tau Yih. “Efficient One-Pass End-to-End Entity Linking for Questions.” In: *Proceedings of EMNLP (short)*. 2020. [pdf]
- [C10] **Sewon Min**, Julian Michael, Hannaneh Hajishirzi, Luke Zettlemoyer. “AmbigQA: Answering Ambiguous Open-domain Questions.” In: *Proceedings of EMNLP (long)*. 2020. [pdf]
- [C9] Vladimir Karpukhin\*, Barlas Oguz\*, **Sewon Min**, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, Wen-tau Yih. “Dense Passage Retrieval for Open-domain Question Answering.” In: *Proceedings of EMNLP (long)*. 2020. *EMNLP 2020’s Most Influential Paper #1*. [pdf]
- [C8] Daniel Khashabi, **Sewon Min**, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark and Hannaneh Hajishirzi. “UnifiedQA: Crossing Format Boundaries With a Single QA System.” In: *Proceedings of Findings of EMNLP (long)*. 2020. *EMNLP 2020’s Most Influential Paper #6*. [pdf]
- [C7] **Sewon Min**, Danqi Chen, Hannaneh Hajishirzi, Luke Zettlemoyer. “A Discrete Hard EM Approach for Weakly Supervised Question Answering.” In: *Proceedings of EMNLP (long)*. 2019. [pdf]

- [C6] **Sewon Min**, Victor Zhong, Luke Zettlemoyer, Hannaneh Hajishirzi. “Multi-hop Reading Comprehension through Question Decomposition and Rescoring.” In: *Proceedings of ACL (long)*. 2019. [\[pdf\]](#)
- [C5] **Sewon Min\***, Eric Wallace\*, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, Luke Zettlemoyer. “Compositional Questions Do Not Necessitate Multi-hop Reasoning.” In: *Proceedings of ACL (short)*. 2019. [\[pdf\]](#)
- [C4] **Sewon Min**, Victor Zhong, Richard Socher, Caiming Xiong. “Efficient and Robust Question Answering from Minimal Context over Documents.” In: *Proceedings of ACL (long)*. 2018. [\[pdf\]](#)
- [C3] Minjoon Seo\*, **Sewon Min\***, Ali Farhadi, Hannaneh Hajishirzi. “Neural Speed Reading via Skim-RNN.” In: *Proceedings of ICLR*. 2018. [\[pdf\]](#)
- [C2] **Sewon Min**, Minjoon Seo, Hannaneh Hajishirzi. “Question Answering through Transfer Learning from Large Fine-grained Supervision Data.” In: *Proceedings of ACL (short)*. 2017. [\[pdf\]](#)
- [C1] Minjoon Seo, **Sewon Min**, Ali Farhadi, Hannaneh Hajishirzi. “Query-Reduction Networks for Question Answering.” In: *Proceedings of ICLR*. 2017. [\[pdf\]](#)

## Tutorials

- [T3] Akari Asai, **Sewon Min**, Zexuan Zhong, Danqi Chen. “Retrieval-based Language Models and Applications.” In: *Proceedings of ACL (Tutorial)*. 2023. [\[pdf\]](#)
- [T2] Iz Beltagy, Arman Cohan, Robert L. Logan IV, **Sewon Min**, Sameer Singh. “Zero- and Few-Shot NLP with Pretrained Language Models.” In: *Proceedings of ACL (Tutorial)*. 2022. [\[pdf\]](#)
- [T1] Iz Beltagy, Arman Cohan, Hannaneh Hajishirzi, **Sewon Min**, Matthew E. Peters. “Beyond Paragraphs: NLP for Long Sequences.” In: *Proceedings of NAACL-HLT (Tutorial)*. 2021. [\[pdf\]](#)

## Peer-Reviewed Workshop Papers

- [W2] Rulin Shao, **Sewon Min**, Luke Zettlemoyer, Pang Wei Koh. “Retrieval-based Language Models Using a Multi-domain Datastore.” In: *Proceedings of Workshop on Distribution Shifts (DistShift) @ NeurIP (Non-archival)*. 2023.
- [W1] Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, **Sewon Min**. “On Making Reading Comprehension More Comprehensive.” In: *Proceedings of Workshop on Machine Reading for Question Answering (MRQA) @ EMNLP*. 2019. [\[pdf\]](#)

## Preprints

- [P4] Tong Chen, Akari Asai, Niloofar Miresghallah, **Sewon Min**, James Grimmermann, Yejin Choi, Hannaneh Hajishirzi, Luke Zettlemoyer, Pang Wei Koh. “CopyBench: Measuring Literal and Non-Literal Reproduction of Copyright-Protected Text in Language Model Generation.” *arXiv preprint*. 2024. [\[pdf\]](#)  
**Submitted to ARR 2024.**
- [P3] Rulin Shao, Jacqueline He, Akari Asai, Weijia Shi, Tim Dettmers, **Sewon Min**, Luke Zettlemoyer, Pang Wei Koh. “Scaling Retrieval-Based Language Models with a Trillion-Token Datastore.” *arXiv preprint*. 2024. [\[pdf\]](#)  
**Submitted to NeurIPS 2024.**
- [P2] **Sewon Min**, Danqi Chen, Luke Zettlemoyer, Hannaneh Hajishirzi. “Knowledge Guided Text Retrieval and Reading for Open Domain Question Answering.” *arXiv preprint*. 2019. [\[pdf\]](#)
- [P1] Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, **Sewon Min**. “Question Answering is a Format; When is it Useful?.” *arXiv preprint*. 2019. [\[pdf\]](#)

## Invited Talks

---

- 09/2024 (Expected) Simons Workshop on [Transformers as a Computational Model](#)
- 09/2024 (Expected) Simons Workshop on [Emerging Settings in Generalization](#)
- 07/2024 National University of Singapore
- 07/2024 DSO National Laboratories
- 07/2024 [LLM Safety Workshop](#)
- 04/2024 Cornell Tech
- 04/2024 Harvard University
- 04/2024 University of Toronto
- 03/2024 University of Texas, Austin
- 03/2024 MIT Media Lab
- 03/2024 MIT EECS
- 03/2024 Carnegie Mellon University - Machine Learning Department
- 03/2024 Carnegie Mellon University - Language Technologies Institute
- 03/2024 Princeton University
- 02/2024 Cornell University
- 02/2024 New York University
- 02/2024 UC Berkeley
- 02/2024 Columbia University
- 02/2024 University of Pennsylvania
- 02/2024 Caltech
- 01/2024 Northwestern University
- 01/2024 Allen School Colloquium (a short talk, representing NLP Research Lab)
- 
- 12/2023 EMNLP 2023 [Big Picture Workshop](#)
- 09/2023 UMass Amherst NLP Group (Host: Mohit Iyyer)
- 09/2023 Privacy in ML Seminar, Google Deepmind (Host: Sewoong Oh)
- 09/2023 IBS Data Science Group, KAIST (Host: Meeyoung Cha)
- 06/2023 Microsoft Semantic Machines (Host: Hao Fang)
- 03/2023 Gradient Podcast Interview (Host: Daniel Bashir)
- 02/2023 Conversation & Reasoning Group, Meta (Host: Zhe Zhang)
- 
- 12/2022 Data Science & Business Analytics Lab, Korea University (Host: Pilsung Kang)
- 10/2022 ISI NL Seminar, University of Southern California (Host: Jonathan May)
- 07/2022 MLOPT Research Group, University of Wisconsin–Madison (Host: Kangwook Lee)
- 04/2022 Language & Intelligence Lab, MIT (Host: Jacob Andreas)
- 
- 12/2021 Hyundai AIRS (Host: Cheoneum Park)
- 10/2021 Facebook AI (Host: Zornitsa Kozareva)
- 08/2021 AI Summer School, Seoul National University (Host: Hyun Oh Song)
- 08/2021 Summer AI Seminar Series, Postech (Host: Gary Geunbae Lee)
- 04/2021 Colloquium Talk, Graduate School of AI at KAIST (Host: Juho Lee)
- 02/2021 School of Electronic Engineering at KAIST (Host: Sung-Ju Lee)
- 
- 09/2020 Apple (Hosts: Chris DuBois and Shayne Longpre)
- 08/2020 Facebook NLP Summit, Question Answering Workshop (Host: Yashar Mehdad)
- 
- 08/2018 Kakao Brain (Host: Benjamin Lee)
- 08/2018 Samsung Advanced Institute of Technology (Host: Jehun Jeon)
- 
- 08/2017 NAVER (Host: Sung Kim)

## Teaching

---

### Tutorial Co-instructor

ACL 2023	Retrieval-based Language Models and Applications with Akari Asai, Zexuan Zhong and Danqi Chen
ACL 2022	Zero- and Few-Shot NLP with Pretrained Language Models with Iz Beltagy, Arman Cohan, Sameer Singh and Robert Logan
NAACL-HLT 2021	Beyond Paragraphs: NLP for Long Sequences with Iz Beltagy, Arman Cohan, Hannaneh Hajishirzi and Matthew E. Peters

### Teaching Assistant

Spring 2022	CSE 573 Artificial Intelligence, University of Washington Instructor: Hannaneh Hajishirzi
Winter 2021	CSE 447+517 Natural Language Processing, University of Washington Instructor: Noah Smith

### Guest Lecturer

Fall 2023	CPSC 488/588 AI Foundation Models, Yale University Instructor: Arman Cohan Topic: Retrieval-based Language Modeling
Winter 2023	CSE 447 Natural Language Processing, University of Washington Instructor: Sofia Serrano Topic: Zero- and Few-shot NLP with Pretrained Language Models
Winter 2023	CSE 573 Artificial Intelligence, University of Washington Instructor: Hannaneh Hajishirzi
Spring 2022	CSE 573 Artificial Intelligence, University of Washington Instructor: Hannaneh Hajishirzi

### Predocctoral Student Mentoring

01/2023–12/2023	Michael Duan, BS student at UW
09/2021–12/2023	Xinxi (Shane) Lyu, BS→MS student at UW→predocctoral young investigator at AI2
12/2021–12/2022	Xinyan (Velocity) Yu, BS→MS student at UW; now a Ph.D. student at USC
06/2021–12/2021	Skyler Hallinan, BS→MS student at UW
01/2021–07/2021	Jungsoo Park, MS student at Korea University; now at NAVER
02/2020–08/2020	Belinda Li, software engineer at Facebook; now a PhD student at MIT

## Professional Services

---

### Workshop Co-organizer

- ICLR 2024 The 2nd Workshop on Mathematical and Empirical Understanding of Foundation Models (ME-FoMo)
- ACL 2022 7th Workshop on Representation Learning for NLP (Repl4NLP)
- ACL 2022 Semiparametric Methods in NLP: Decoupling Logic from Knowledge (Spa-NLP)
- EMNLP 2021 3rd Workshop on Machine Reading for Question Answering (MRQA)
- AKBC 2021 2nd Workshop on Unstructured and Structured KBs (USKB)
- NeurIPS 2020 Competition on Efficient Open-Domain Question Answering (EfficientQA)
- AKBC 2020 Workshop on Unstructured and Structured KBs (USKB)

**Area Chair** EMNLP 2024, COLM 2024, ACL 2023, EMNLP 2023, COLING 2022

### Reviewer/Program Committee

Journals:

- Transactions of the Association for Computational Linguistics (TACL; 2022–)
- Transactions on Machine Learning Research (TMLR; 2022–)
- Journal of Artificial Intelligence Research (JAIR; 2023)

Conferences: (<sup>†</sup> denotes outstanding reviewer)

- ACL (2019, 2020, 2021<sup>†</sup>)
- EMNLP (2019<sup>†</sup>, 2020, 2021<sup>†</sup>, 2022)
- NAACL (2021)
- AACL (2020)
- ACL Rolling Review (2021, 2022, 2023)
- NeurIPS (2018, 2020, 2021, 2022)
- ICLR (2019, 2020, 2021, 2022, 2023)
- ICML (2021, 2023)
- AAAI (2020, 2021)
- AKBC (2019, 2020, 2021, 2022)

Workshops:

- Workshop on Long-Context Foundation Models: ICML 2024
- Workshop on Distribution Shifts: NeurIPS 2023
- Workshop on Retrieval-Enhanced Machine Learning: SIGIR 2023
- Workshop on Mathematical and Empirical Understanding of Foundation Models: ICLR 2023
- Workshop on Knowledge Retrieval and Language Models: ICML 2022
- Student Research Workshop: ACL 2019, ACL 2020, AACL 2020, EACL 2021, NAACL 2021, ACL 2021
- Workshop on NLP for Positive Impact: ACL 2021, EMNLP 2022
- The 2nd Workshop on Machine Reading for Question Answering: EMNLP 2019

Competition Track, NeurIPS 2021

### Departmental Service

- 2021–2023 UW CSE: Admission Committee Member for the Ph.D. Program
- 2021 UW NLP: Student Co-organizer of the Invited Talk Series
- 2021 UW NLP: Student Co-organizer of the NLP Retreat
- 2019 UW CSE: Prospective Student Committee Co-Chair for the Ph.D. Program