Mohit Iyyer

miyyer@umd.edu
http://cs.umd.edu/~miyyer

Academic Appointments

University of Maryland, College Park
Associate Professor, Computer Science — 2025 — *present*
Research Interests: *natural language processing, large language models*

University of Massachusetts Amherst
Adjunct Associate Professor, Computer Science — 2025 — *present*
Associate Professor, Computer Science — 2023 — 2025
Assistant Professor, Computer Science — 2018 — 2023

Education

University of Maryland, College Park — 2012 — 2017
Ph.D. in Computer Science
Master of Science, Computer Science, 2012 - 2014

Washington University in St. Louis — 2008 — 2012
BS in Computer Science
Minors: Bioinformatics & Creative Writing

Research & Professional Experience

Postdoctoral Researcher, Allen Institute for AI — 2017 — 2018
Supervisors: Luke Zettlemoyer, Scott Yih

Research Assistant, University of Maryland, College Park — 2013 — 2017
Advisors: Jordan Boyd-Graber, Hal Daumé III

Research Intern, NLP Group, Microsoft Research, Redmond — Summer 2016
Supervisors: Scott Yih, Ming-Wei Chang

Research Intern, MetaMind — Spring 2015
Supervisor: Richard Socher

Awards

Distinguished Paper Award, ACM CCS 2023

Outstanding Paper Award, EACL 2023

Samsung AI Researcher of the Year Award, 2022

NSF CAREER Award, 2021

Best Long Paper Award, NAACL 2018

Larry S. Davis Doctoral Dissertation Award, 2017

Best Long Paper Award, NAACL 2016

Best Demonstration Award, NeurIPS 2015

Conference Publications

As of April 2026, my papers have been cited over 34K times (h-index 51) according to my Google Scholar profile. ACL, NAACL, EMNLP, ICML, NeurIPS, CVPR, and ICLR are peer reviewed conferences with acceptance rates typically around 25-30%.

1. Katherine Thai, Bradley Emi, Elyas Masrour, **Mohit Iyyer**. EditLens: Quantifying the Extent of AI Editing in Text. *ICLR 2026.*
2. Yapei Chang, Yekyung Kim, Michael Krumdick, Amir Zadeh, Chuan Li, Chris Tanner, **Mohit Iyyer**. BLEUBERI: BLEU is a surprisingly effective reward for instruction following. *NeurIPS 2025.*
3. Alisha Srivastava*, Emir Korukluoglu*, Minh Nhat Le*, Duyen Tran, Chau Minh Pham, Marzena Karpinska, **Mohit Iyyer**. OWL: Probing Cross-Lingual Recall of Memorized Texts via World Literature. *EMNLP 2025.*
4. Anmol Mekala*, Anirudh Atmakuru*, Yixiao Song, Marzena Karpinska, **Mohit Iyyer**. Does quantization affect models' performance on long-context tasks? *EMNLP 2025.*
5. Rishanth Rajendhran, Amir Zadeh, Matthew Sarte, Chuan Li, **Mohit Iyyer**. VeriFastScore: Speeding up long-form factuality evaluation. *Findings of EMNLP 2025.*
6. Nischal Ashok Kumar, Chau Minh Pham, **Mohit Iyyer**, Andrew Lan. Whose story is it? Personalizing story generation by inferring author styles. *AACL 2025.*
7. Yixiao Song, Katherine Thai, Chau Minh Pham, Yapei Chang, Mazin Nadaf, **Mohit Iyyer**. BEARCUBS: A benchmark for computer-using web agents. *COLM 2025.*
8. Yekyung Kim, Jenna Russell, Marzena Karpinska, **Mohit Iyyer**. One ruler to measure them all: Benchmarking multilingual long-context language models. *COLM 2025.*
9. Chau Minh Pham, Yapei Chang, **Mohit Iyyer**. CLIPPER: Compression enables long-context synthetic data generation. *COLM 2025.*
10. Jenna Russell, Marzena Karpinska, **Mohit Iyyer**. People who frequently use ChatGPT for writing tasks are accurate and robust detectors of AI-generated text. *ACL 2025.*
11. Shane Arora*, Marzena Karpinska*, Hung-Ting Chen, Ipsita Bhattacharjee, **Mohit Iyyer**, Eunsol Choi. CaLMQA: Exploring culturally specific long-form question answering across 23 languages. *ACL 2025.*
12. Katherine Thai, **Mohit Iyyer**. Literary Evidence Retrieval via Long-Context Language Models. *ACL 2025 (short).*
13. Rachneet Sachdeva, Yixiao Song, **Mohit Iyyer**, Iryna Gurevych. Localizing and Mitigating Errors in Long-form Question Answering. *Findings of ACL 2025.*

14. Chaitanya Malaviya, Joseph Chee Chang, Dan Roth, **Mohit Iyyer**, Mark Yatskar, Kyle Lo. Contextualized Evaluations: Judging Language Model Responses to Underspecified Queries. *TACL 2025.*

15. Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, **Mohit Iyyer**. One Thousand and One Pairs: A "novel" challenge for long-context language models. *EMNLP 2024.*

16. Yapei Chang, Kalpesh Krishna, Amir Houmansadr, John Wieting, **Mohit Iyyer**. PostMark: A Robust Blackbox Watermark for Large Language Models. *EMNLP 2024.*

17. Yixiao Song, Yekyung Kim, **Mohit Iyyer**. VERISCORE: Evaluating the factuality of verifiable claims in long-form text generation. *Findings of EMNLP 2024.*

18. Chau Minh Pham, Simeng Sun, **Mohit Iyyer**. Suri: Multi-constraint Instruction Following for Long-form Text Generation. *Findings of EMNLP 2024.*

19. Yekyung Kim, Yapei Chang, Marzena Karpinska, Aparna Garimella, Varun Manjunatha, Kyle Lo, Tanya Goyal, **Mohit Iyyer**. FABLES: Evaluating faithfulness and content selection in book-length summarization. *COLM 2024.*

20. Ali Naseh, Katherine Thai, **Mohit Iyyer**, Amir Houmansadr. Iteratively Prompting Multimodal LLMs to Reproduce Natural and AI-Generated Images. *COLM 2024.*

21. Jiachen Zhao, Wenlong Zhao, Andrew Drozdov, Benjamin Rozonoyer, Md Arafat Sultan, Jay-Yoon Lee, **Mohit Iyyer**, Andrew McCallum. Multistage Collaborative Knowledge Distillation from a Large Language Model for Semi-Supervised Sequence Generation. *ACL 2024.*

22. Tu Vu, **Mohit Iyyer**, Xuezhi Wang, Noah Constant, Jerry Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny Zhou, Quoc Le, Thang Luong. FreshLLMs: Refreshing Large Language Models with Search Engine Augmentation. *Findings of ACL 2024.*

23. Chau Pham, Alexander Hoyle, Simeng Sun, **Mohit Iyyer**. TopicGPT: A Prompt-based Topic Modeling Framework. *NAACL 2024.*

24. Yixiao Song, Kalpesh Krishna, Rajesh Bhatt, Kevin Gimpel, **Mohit Iyyer**. GEE! Grammar Error Explanation with Large Language Models. *Findings of NAACL 2024.*

25. Prasanna Lakkur Subramanyam, **Mohit Iyyer**, Brian Levine. Triage of Messages and Conversations in a Large-Scale Child Victimization Corpus . *ACM The Web Conference 2024 (Web4Good Track).*

26. Yapei Chang, Kyle Lo, Tanya Goyal, **Mohit Iyyer**. BooookScore: A systematic exploration of book-length summarization in the era of LLMs. *ICLR 2024 (oral).*

27. Simeng Sun, Yang Liu, Shuohang Wang, Chenguang Zhu, **Mohit Iyyer**. PEARL: Prompting Large Language Models to Plan and Execute Actions Over Long Documents. *EACL 2024.*

28. Simeng Sun, Yang Liu, Dan Iter, Chenguang Zhu, **Mohit Iyyer**. How Does In-Context Learning Help Prompt Tuning? *Findings of EACL 2024.*

29. Shufan Wang, Yixiao Song, Andrew Drozdov, Aparna Garimella, Varun Manjunatha, **Mohit Iyyer**. kNN-LM Does Not Improve Open-ended Text Generation. *EMNLP 2023.*

30. Sewon Min, Kalpesh Krishna, Xinxi 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. *EMNLP 2023*.

31. Nader Akoury, Qian Yang, **Mohit Iyyer**. Disco Elysium: Exploring Player Perceptions of LLM-Generated Dialogue within a Commercial Video Game. *Findings of EMNLP 2023*.

32. Andrew Drozdov, Honglei Zhuang, Zhuyun Dai, Zhen Qin, Razieh Rahimi, Xuanhui Wang, Dana Alon, **Mohit Iyyer**, Andrew McCallum, Donald Metzler, Kai Hui. PaRaDe: Passage Ranking using Demonstrations with LLMs. *Findings of EMNLP 2023 (short)*.

33. Marzena Karpinska, **Mohit Iyyer**. Large language models effectively leverage document-level context for literary translation, but critical errors persist. *WMT 2023*.

34. Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, **Mohit Iyyer**. Paraphrasing evades detectors of AI-generated text, but retrieval is an effective defense. *NeurIPS 2023*.

35. Fangyuan Xu*, Yixiao Song*, **Mohit Iyyer**, Eunsol Choi. A Critical Evaluation of Evaluations for Long-form Question Answering. *ACL* 2023.

36. Ali Naseh, Kalpesh Krishna, **Mohit Iyyer**, Amir Houmansadr. Stealing the Decoding Algorithms of Language Models. *CCS* 2023. *Distinguished paper award.*

37. Kalpesh Krishna, Erin Bransom, Bailey Kuehl, **Mohit Iyyer**, Pradeep Dasigi, Arman Cohan, Kyle Lo. LongEval: Guidelines for Human Evaluation of Faithfulness in Long-form Summarization. *EACL* 2023. *Outstanding paper award.*

38. Ankita Gupta, Marzena Karpinska, Wenlong Zhao, Kalpesh Krishna, Jack Merullo, Luke Yeh, **Mohit Iyyer**, Brendan O'Connor. ezCoref: Towards Unifying Annotation Guidelines for Coreference Resolution. *Findings of EACL* 2023.

39. Kalpesh Krishna, Yapei Chang, John Wieting, **Mohit Iyyer**. RankGen: Improving Text Generation with Large Ranking Models. *EMNLP* 2022

40. Tu Vu, Aditya Barua, Brian Lester, Daniel Cer, **Mohit Iyyer**, Noah Constant. Overcoming Catastrophic Forgetting in Zero-Shot Cross-Lingual Generation. *EMNLP* 2022

41. Yixiao Song, Kalpesh Krishna, Rajesh Bhatt, **Mohit Iyyer**. SLING: Sino Linguistic Evaluation of Large Language Models. *EMNLP* 2022

42. Katherine Thai*, Marzena Karpinska*, Kalpesh Krishna, William Ray, Moira Inghilleri, John Wieting, **Mohit Iyyer**. Exploring Document-Level Literary Machine Translation with Parallel Paragraphs from World Literature. *EMNLP* 2022

43. Marzena Karpinska, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, **Mohit Iyyer**. DEMETR: Diagnosing Evaluation Metrics for Translation. *EMNLP* 2022

44. Andrew Drozdov, Shufan Wang, Razieh Rahimi, Andrew McCallum, Hamed Zamani, **Mohit Iyyer**. You can't pick your neighbors, or can you? When and How to Rely on Retrieval in the KNN-LM. *Findings of EMNLP* 2022

45. Shufan Wang, Fangyuan Xu, Laure Thompson, Eunsol Choi, **Mohit Iyyer**. Modeling Exemplification in Long-form Question Answering via Retrieval. *NAACL* 2022.

46. Simeng Sun, Katherine Thai, **Mohit Iyyer**. ChapterBreak: A Challenge Dataset for Long-Range Language Models. *NAACL* 2022 (short).

47. Katherine Thai, Yapei Chang, Kalpesh Krishna, **Mohit Iyyer**. RELIC: Retrieving Evidence for Literary Claims. *ACL* 2022.

48. Simeng Sun, Kalpesh Krishna, Andrew Mattarella-Micke, and **Mohit Iyyer**. Do Long-Range Language Models Actually Use Long-Range Context? *EMNLP* 2021.

49. Tu Vu, Minh-Thang Luong, Quoc Le, Grady Simon, and **Mohit Iyyer**. STraTA: Self-Training with Task Augmentation for Better Few-shot Learning. *EMNLP* 2021.

50. Marzena Karpinska, Nader Akoury, and **Mohit Iyyer**. The Perils of Using Mechanical Turk to Evaluate Open-Ended Text Generation. *EMNLP* 2021.

51. Shufan Wang, Laure Thompson, and **Mohit Iyyer**. Phrase-BERT: Improved Phrase Embeddings from BERT with an Application to Corpus Exploration. *EMNLP* 2021.

52. Zhiyang Xu, Andrew Drozdov, Jay Yoon Lee, Tim O'Gorman, Subendhu Rongali, Dylan Finkbeiner, Shilpa Suresh, **Mohit Iyyer**, and Andrew McCallum. Improved Latent Tree Induction with Distant Supervision via Span Constraints. *EMNLP* 2021.

53. Simeng Sun, Wenlong Zhao, Varun Manjunatha, Rajiv Jain, Vlad Morariu, Franck Dernoncourt, Balaji Vasan Srinivasan, and **Mohit Iyyer**. IGA: An Intent-Guided Authoring Assistant. *EMNLP* 2021.

54. Sumanta Bhattacharyya, Amirmohammad Rooshenas, Subhajit Naskar, Simeng Sun, **Mohit Iyyer**, Andrew McCallum. Energy-Based Reranking: Improving Neural Machine Translation Using Energy-Based Models. *ACL* 2021.

55. Iman Deznabi, **Mohit Iyyer**, Madalina Fiterau. Predicting In-Hospital Mortality by Combining Clinical Notes with Time-Series Data. *Findings of ACL* 2021.

56. Amee Trivedi, Kate Silverstein, Emma Strubell, **Mohit Iyyer**, Prashant Shenoy. WiFiMod: Transformer-based Indoor Human Mobility Modeling using Passive Sensing. *ACM COMPASS* 2021.

57. Kalpesh Krishna, Aurko Roy, **Mohit Iyyer**. Hurdles to Progress in Long-form Question Answering. *NAACL* 2021.

58. Hiroshi Iida, June Thai, Varun Manjunatha, **Mohit Iyyer**. TABBIE: Pretrained Representations of Tabular Data. *NAACL* 2021.

59. Simeng Sun, **Mohit Iyyer**. Revisiting Simple Neural Probabilistic Language Models. *NAACL* 2021 (short).

60. Haw-Shiuan Chang, Jiaming Yuan, **Mohit Iyyer**, Andrew McCallum. Changing the Mind of Transformers for Topically-Controllable Language Generation. *EACL* 2021.

61. Chen Qu, Liu Yang, Cen Chen, W. Bruce Croft, Kalpesh Krishna, **Mohit Iyyer**. Weakly-Supervised Open-Retrieval Conversational Question Answering. *ECIR* 2021.

62. Tu Vu, Tong Wang, Tsendsuren Munkhdalai, Alessandro Sordoni, Adam Trischler, Andrew Mattarella-Micke, Subhransu Maji, **Mohit Iyyer**. Exploring and Predicting Transferability across NLP Tasks. *EMNLP* 2020.

63. Kalpesh Krishna, John Wieting, **Mohit Iyyer**. Reformulating Unsupervised Style Transfer as Paraphrase Generation. *EMNLP* 2020.

64. Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, **Mohit Iyyer**. STORIUM: A Dataset and Platform for Machine-in-the-Loop Story Generation. *EMNLP* 2020.

65. Andrew Drozdov, Subendhu Rongali, Yi-Pei Chen, Tim O'Gorman, **Mohit Iyyer**, Andrew McCallum. Unsupervised Parsing with S-DIORA: Single Tree Encoding for Deep Inside-Outside Recursive Autoencoders. *EMNLP* 2020.

66. Weiqiu You*, Simeng Sun*, **Mohit Iyyer**. Hard-Coded Gaussian Attention for Neural Machine Translation. *ACL* 2020.

67. Fenfei Guo, Jordan Boyd-Graber, **Mohit Iyyer**, Leah Findlater. Which Evaluations Uncover Sense Representations that Actually Make Sense? *LREC* 2020.

68. Kalpesh Krishna, Gaurav Singh Tomar, Ankur Parikh, Nicolas Papernot, **Mohit Iyyer**. Thieves on Sesame Street! Model Extraction of BERT-based APIs. *ICLR* 2020.

69. Jack Merullo*, Luke Yeh*, Abram Handler, Alvin Grissom II, Brendan O'Connor, **Mohit Iyyer**. Investigating Sports Commentator Bias within a Large Corpus of American Football Broadcasts. *EMNLP* 2019 (short).

70. Andrew Drozdov, Patrick Verga, Yi-Pei Chen, **Mohit Iyyer**, Andrew McCallum. Unsupervised Labeled Parsing with Deep Inside-Outside Recursive Autoencoders. *EMNLP* 2019 (short).

71. Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W. Bruce Croft, **Mohit Iyyer**. Attentive History Selection for Conversational Question Answering. *CIKM* 2019.

72. Nader Akoury, Kalpesh Krishna, **Mohit Iyyer**. Syntactically Supervised Transformers for Faster Neural Machine Translation. *ACL* 2019.

73. Kalpesh Krishna, **Mohit Iyyer**. Generating Question-Answer Hierarchies. *ACL* 2019.

74. Tu Vu, **Mohit Iyyer**. Encouraging Paragraph Embeddings to Remember Sentence Identity Improves Classification. *ACL* 2019 (short).

75. Chen Qu, Liu Yang, Minghui Qiu, W. Bruce Croft, Yongfeng Zhang, **Mohit Iyyer**. BERT with History Modeling for Conversational Question Answering. *SIGIR* 2019 (short).

76. Andrew Drozdov, Patrick Verga, Mohit Yadav, **Mohit Iyyer**, Andrew McCallum. Unsupervised Latent Tree Induction with Deep Inside-Outside Recursive Auto-Encoders. *NAACL* 2019.

77. Shufan Wang, **Mohit Iyyer**. Casting Light on Invisible Cities: Computationally Engaging with Literary Criticism. *NAACL* 2019 (short).

78. Eunsol Choi*, He He*, **Mohit Iyyer***, Mark Yatskar*, Wen-tau Yih, Yejin Choi, Percy Liang, Luke Zettlemoyer. QuAC: Question Answering in Context. *EMNLP* 2018.

79. Shi Feng, Eric Wallace, Alvin Grissom II, **Mohit Iyyer**, Pedro Rodriguez, Jordan Boyd-Graber. Pathologies of Neural Models Make Interpretation Difficult. *EMNLP* 2018.

80. Kalpesh Krishna, Preethi Jyothi, **Mohit Iyyer**. Revisiting the Importance of Encoding Logic Rules in Sentiment Classification. *EMNLP* 2018 (short).

81. **Mohit Iyyer***, John Wieting*, Kevin Gimpel, Luke Zettlemoyer. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. *NAACL* 2018.

82. Matthew E. Peters, Mark Neumann, **Mohit Iyyer**, Matt Gardner, Christopher Clark, Kenton Lee, Luke Zettlemoyer. Deep contextualized word representations. *NAACL* 2018. *Best long paper award*

83. Varun Manjunatha*, **Mohit Iyyer***, Jordan Boyd-Graber, Larry Davis. Learning to Color from Language. *NAACL* 2018 (short); also appeared at *NIPS 2017 ViGIL Workshop*.

84. **Mohit Iyyer**, Wen-tau Yih, Ming-Wei Chang. Search-based Neural Structured Learning for Sequential Question Answering. *ACL* 2017.

85. **Mohit Iyyer***, Varun Manjunatha*, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, Hal Daumé III, Larry Davis. The Amazing Mysteries of the Gutter: Drawing Inferences Between Panels in Comic Book Narratives. *CVPR* 2017.

86. Snigdha Chaturvedi, **Mohit Iyyer**, Hal Daumé III. Unsupervised Learning of Evolving Relationships Between Literary Characters. *AAAI* 2017.

87. **Mohit Iyyer**, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, Hal Daumé III. Feuding Families and Former Friends: Unsupervised Learning for Dynamic Fictional Relationships. *NAACL* 2016. *Best long paper award.*

88. Ankit Kumar, Ozan Irsoy, Peter Ondruska, **Mohit Iyyer**, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. Ask Me Anything: Dynamic Memory Networks for Natural Language Processing. *ICML* 2016.

89. Jordan Boyd-Graber, **Mohit Iyyer**, He He, Hal Daumé III. Interactive Incremental Question Answering. *NIPS Demonstration Track*, 2015. *Outstanding Demonstration Award.*

90. **Mohit Iyyer**, Varun Manjunatha, Jordan Boyd-Graber, Hal Daumé III. Deep Unordered Composition Rivals Syntactic Methods for Text Classification. *ACL* 2015.

91. Anupam Guha, **Mohit Iyyer**, Danny Bouman, Jordan Boyd-Graber. Removing the Training Wheels: A Coreference Dataset that Entertains Humans and Challenges Computers. *NAACL* 2015.

92. **Mohit Iyyer**, Jordan Boyd-Graber, Richard Socher, Hal Daumé III. A Neural Network for Factoid Question Answering over Paragraphs. *EMNLP* 2014.

93. **Mohit Iyyer**, Peter Enns, Jordan Boyd-Graber, Philip Resnik. Political Ideology Detection Using Recursive Neural Networks. *ACL* 2014.

Workshop Publications

Nader Akoury, Ronan Salz, **Mohit Iyyer**. Towards Grounded Dialogue Generation in Video Game Environments *AAAI Workshop on Creative AI Across Modalities*, 2023.

Simeng Sun, Brian Dillon, **Mohit Iyyer**. How Much Do Modifications to Transformer Language Models Affect Their Ability to Learn Linguistic Knowledge? *ACL Workshop on Insights from Negative Results in NLP*, 2022.

Dhruvil Gala, Mohammad Omar Khursheed, Hannah Lerner, Brendan O'Connor, **Mohit Iyyer**. Analyzing Gender Bias within Narrative Tropes. *EMNLP Workshop on NLP and CSS*, 2020.

Anupam Guha, **Mohit Iyyer**, Jordan Boyd-Graber. A Distorted Skull Lies in the Bottom Center... Identifying Paintings from Text Descriptions. *NAACL Human-Computer QA Workshop*, 2016.

Mohit Iyyer, Jordan Boyd-Graber, Hal Daumé III. Generating Sentences from Semantic Vector Space Representations. *NIPS Workshop on Learning Semantics*, 2014.

Preprints

Yapei Chang, Kyle Lo, **Mohit Iyyer**, Luca Soldaini. How2Everything: Mining the Web for How-To Procedures to Evaluate and Improve LLMs. *arXiv 2026*.

Prithviraj Tarale, Kiet Chu, Abhishek Varghese, Kai-Chun Liu, Maxwell A. Xu, **Mohit Iyyer**, Sunghoon I. Lee. Bio-Inspired Self-Supervised Learning for Wrist-worn IMU Signals. *arXiv 2026*.

Jenna Russell, Marzena Karpinska, Destiny Akinode, Katherine Thai, Bradley Emi, Max Spero, **Mohit Iyyer**. AI use in American newspapers is widespread, uneven, and rarely disclosed. *arXiv 2025*.

Chau Minh Pham, Jenna Russell, Dzung Pham, **Mohit Iyyer**. Frankentext: Stitching random text fragments into long-form narratives. *arXiv 2025*.

Abhinav Kumar, Jaechul Roh, Ali Naseh, Marzena Karpinska, **Mohit Iyyer**, Amir Houmansadr, Eugene Bagdasarian. OverThink: Slowdown Attacks on Reasoning LLMs. *arXiv 2025*.

Garima Dhanania, Sheshera Mysore, Chau Minh Pham, **Mohit Iyyer**, Hamed Zamani, Andrew McCallum. Interactive Topic Models with Optimal Transport. *arXiv 2024*.

Simeng Sun, Dhawal Gupta, **Mohit Iyyer**. Exploring the impact of low-rank adaptation on the performance, efficiency, and regularization of RLHF. *arXiv 2023*.

Pedro Rodriguez, Shi Feng, **Mohit Iyyer**, He He, Jordan Boyd-Graber. Quizbowl: The Case for Incremental Question Answering. *arXiv 2019*.

Funding

In total, I have raised over $5M in research funding from the NSF, industry labs, and philanthropic organizations.

- Google Gemini Academic Research Credit Award, 2025
  PI: Mohit Iyyer
  $100,000

- Google Research Gift, 2025
  "Leveraging Structural Differences to Improve and Identify AI-Generated Stories"
  PI: Mohit Iyyer
  $30,000

- NSF RI MEDIUM, 2023-2027
  "Multilingual Long-form QA with Retrieval-Augmented Language Models"
  PI: Eunsol Choi, co-PI: Mohit Iyyer
  UMass share: $554,000

- Google Research Award, 2023
  "RankGen: Improving Text and Code Generation with Large Ranking Models"
  PI: Mohit Iyyer
  $30,000

- Kensho Research Award, 2023
  "Improving factuality via LLM alignment"
  PI: Mohit Iyyer
  $74,000

- NSF RETTL, 2022-2025
  "STEM Learning Embedded in a Machine-in-the-Loop Collaborative Story Writing Game"
  PI: Mohit Iyyer, co-PIs: Andrew Lan, Stephen Hood, Danielle McNamara
  $621,352

- Open Philanthropy, 2022-2024
  "Improving AI translation of novels into English"
  PI: Mohit Iyyer
  $822,365

- Alexa Prize TaskBot Challenge, 2021-2022
  PI: Hamed Zamani, co-PIs: Mohit Iyyer, Bruce Croft
  $250,000

- NSF CAREER, 2021-2026
  "Building Creative Writing Assistants for Machine-in-the-Loop Storytelling"
  PI: Mohit Iyyer
  $532,000

- NSF RI Medium, 2020-2024
  "Tree-Structured Self-Supervised Modeling for Natural Language"
  PI: Mohit Iyyer, co-PI: Andrew McCallum
  $1.2 million

- Subcontract with US Navy, 2020-2023
  "Natural language processing of fleet data"
  PI: Mohit Iyyer
  $136,295

- Adobe Faculty Research Gift, 2019-2026
  PI: Mohit Iyyer
  $270,000 cumulative

- IBM Faculty Award, 2020
  PI: Mohit Iyyer
  $20,000

- NSF CCRI Planning Grant, 2019-2020
  "Planning for the Development of a Platform to Support Multilingual and Multi-Domain Coreference Annotation for Natural Language Processing Research"
  PI: Brendan O'Connor, co-PI: Mohit Iyyer
  $100,000

- Genpact Faculty Gift, 2019
  PI: Mohit Iyyer
  $50,000

- Sony Research Award, 2019
  PI: Mohit Iyyer
  $60,000

- Intuit AI Award, 2019
  PI: Mohit Iyyer
  $133,750

- Facebook Robust Deep Learning for NLP Award, 2019
  PIs: Brendan O'Connor & Mohit Iyyer
  $76,905

- Allen Institute Young Investigator Award, 2018
  PI: Mohit Iyyer
  $125,000

Advising

PhD advisees:

Shufan Wang (2019-), CS

Yapei Chang (2022-), CS

Chau Pham (2022-), CS

Yekyung Kim (2023-), CS

Jenna Russell (2024-), CS

Rishanth Rajendhran (2024-), CS

Alumni:

- Katherine Thai (2020-2025), CS, *now founding research scientist at Pangram*
- Yixiao Song (2022-2025), Linguistics & CS, *co-advised with Rajesh Bhatt, now research scientist at Google*
- Marzena Karpinska (2020-2024), Postdoc, CS, *now Asst. Prof at Simon Fraser*
- Nader Akoury (2018-2024), CS, *postdoc at Cornell*
- Simeng Sun (2019-2024), CS, *research scientist at Nvidia*
- Andrew Drozdov (2018-2024, *co-advised with Andrew McCallum*), CS, *research scientist at Databricks*
- Kalpesh Krishna (2018-2023), CS, *Google PhD Fellow, now research scientist at Google Gemini*
- Tu Vu (2018-2023), CS, *Asst. Prof at Virginia Tech & research scientist at Google DeepMind*

On committee:

Emma Strubell (UMass CS PhD, 2019)

Patrick Verga (UMass CS PhD, 2019)

Carolyn Anderson (UMass Linguistics PhD, 2020)

Luke Vilnis (UMass CS PhD, 2020)

Su Lin Blodgett (UMass CS PhD, 2020)

Brandon Prickett (UMass Linguistics PhD, 2020)

Abe Handler (UMass CS PhD, 2021)

Katie Keith (UMass CS PhD, 2021)

Chen Qu (UMass CS PhD, 2021)

Chenyun Wu (UMass CS PhD, 2021)

Trapit Bansal (UMass CS PhD, 2021)

Rajarshi Das (UMass CS PhD, 2022)

Xiang Lorraine Li (UMass CS PhD, 2022)

Haw-Shiuan Chang (UMass CS PhD, 2022)

Subendhu Rongali (UMass CS PhD, 2022)

Difan Liu (UMass CS PhD, 2022)

Virat Shejwalkar (UMass CS PhD, 2022)

Jyoti Iyer (UMass Linguistics PhD, 2022)

Pinelopi Papalampidi (University of Edinburgh Informatics PhD, 2022)

Mahmood Jasim (UMass CS PhD, 2023)

Shahrzad Naseri (UMass CS PhD, 2023)

Pengshan Cai (UMass CS PhD, 2023)

Dung Thai (UMass CS PhD, 2023)

Pengshan Cai (UMass CS PhD, 2023)

Youngwoo Kim (UMass CS PhD, 2024)

Sheshera Mysore (UMass CS PhD, 2024)

Zhiqi Huang (UMass CS PhD, 2024)

Taisiya Glushkova (Instituto Superior Técnico CS PhD, 2024)

Rohit Saxena (University of Edinburgh CS PhD, 2025)

Mollie Schichman (UMD CS PhD, 2025)

Erica Cai (UMass CS PhD, 2025)

Josh Davis (UMD CS PhD, 2026)

Mentoring:

Mazin Nadaf (2025-), UMD CS UG

James Zhou (2025-), UMD CS UG

Jiarui Liu (2023-2024), UMass CS MS

Anmol Mekala (2024), UMass CS MS

Anirudh Atmakuru (2024), UMass CS MS

Alisha Srivastava (2023-2024), UMass CS UG

Emir Korukluoglu (2023-2024), UMass CS UG

Minh Nhat Le (2023-2024), UMass CS UG

Duyen Tran (2023-2024), UMass CS UG

Ipsita Bhattacharjee (2023-2024), UMass CS UG

Ronan Salz (2022-2023), UMass CS UG

Naveen Nizar (2022), UMass CS MS

Nishant Raj (2022), UMass CS MS

Xiaoyu Song (2021-2022), UMass CS MS

George Wei (2021-2022), UMass CS UG

Yapei Chang (2020-), Smith CS UG, *now PhD student in my group*

Shriya Atmakuri (2021), UMass CS MS

Evan Moore (2020-2021), UMass CS MS

Kavya Jeganathan (2020-2021), UMass CS UG

Luke Yeh (2018-2020), UMass CS UG

Wenlong Zhao (2020-2021), UMass CS MS, *now CS PhD student at UMass*

Sangeetha Balasubramanian (2020-2021), UMass CS MS, *now at ML engineer at Amazon*

Akshita Bhagia (2019-2020), UMass CS MS *now research engineer at AI2*

Dhruvil Gala (2019-2020), UMass CS UG, *now at Microsoft*

Varun Sharma (2019-2020), UMass CS MS, *now at Bloomberg*

Weiqiu You (2018-2020), UMass CS MS, *now PhD student at UPenn*

Jack Merullo (2018-2020), UMass CS UG *now PhD student at Brown*

Davis Yoshida (2016), CU Boulder undergraduate, *now PhD student at TTIC*

Danny Bouman (2015), UMD College Park undergraduate

Teaching

Fall 2025: Natural Language Processing (CMSC 723), *72 students*

Spring 2025: Seminar on Long-Context LLMs (CMSC 848O), *45 students*

Fall 2024: NLP Seminar (CS 692L)

Spring 2024: Advanced Natural Language Processing (CS685), *204 students*

Spring 2023: Advanced Natural Language Processing (CS685), *168 students*

Spring 2022: Advanced Natural Language Processing (CS685), *110 students*

Fall 2021: Advanced Natural Language Processing (CS685), *120 students*

Fall 2020: Advanced Natural Language Processing (CS685), *122 students*

Spring 2019: Deep Learning for Natural Language Processing (CS690D), *106 students, new course*

Fall 2019: Introduction to Natural Language Processing (CS585), *200 students*

Fall 2018: Introduction to Natural Language Processing (CS585), *107 students*

Invited talks

Apr 2026: talk at MASC-SLL 2026

Nov 2025: talk at Ohio State TDAI seminar

Apr 2025: talk at University of Edinburgh ILCC seminar

Aug 2024: keynote at KDD CUP 2024 RAG workshop

Jul 2024: keynote at Workshop on Long Context Foundation Models (ICML 2024)

Jun 2024: talk at Tel Aviv University NLP seminar

Feb 2024: talk at University of Chicago / TTIC

Dec 2023: talk at University of Tokyo

Nov 2023: talk at MIT Embodied Intelligence seminar

Nov 2023: talk at UPenn CLunch

Oct 2023: talk at UMD CLIP Colloquium

Jul 2023: keynote at SIGIR Workshop on Retrieval-Enhanced Machine Learning

May 2023: talk at EACL Workshop on Insights from Negative Results in NLP

Apr 2023: talk at NYU Text-as-Data series

Dec 2022: talk at IndoML 2022

Jul 2022: lecture at 6th International Gran Canaria School on Deep Learning

May 2022: talk at Baidu Research *(remote)*

Apr 2022: talk at UNC Chapel Hill *(remote)*

Oct 2021: talk at Cornell AI seminar *(remote)*

Jun 2021: talk at Yandex NLP seminar *(remote)*

Jun 2021: talk at Cambridge NLP seminar *(remote)*

Mar 2021: talk at Georgia Tech NLP seminar *(remote)*

Sep 2020: talk at UMass CICS Computing and Social Justice series *(remote)*

Sep 2020: talk at UPenn CLunch *(remote)*

Aug 2020: talk at Data Science fwdays'20, Kyiv *(remote)*

Jul 2020: talk at USC/ISI Boston office *(remote)*

Nov 2019-Jan 2020: *Towards story generation*: NYU, Google AI Language, CMU LTI

Nov 2019: *Racial bias in sports commentary*: Brown University

Sep 2019: *Neural Question Answering and Generation*: IBM QA and Semantic Parsing workshop

Mar 2019: *Neural Language Modeling*: UMass Cognitive Brown Bag

Nov 2018: *Deep Learning for Question Answering*: University of Antwerp

Oct 2018-May 2019: *Towards Understanding Narratives with Artificial Intelligence*: University of Vermont (Science of Stories Symposium), WPI, UMass Data Science for the Humanities Series, Amazon Cambridge

July 2018: *Generating QA Dialogs from Documents*: TTIC Language Generation workshop

Mar 2018: *Generating Adversarial Examples with Syntactically Controlled Paraphrase Networks*: USC/ISI NL Seminar

Feb 2018: *Towards Understanding Creative Language with Artificial Intelligence*: Ursinus College

Jan 2018: *Generating Syntactic Adversaries with Controlled Paraphrasing*: Indian Institute of Science, Bengaluru

Feb-April 2017, *Deep Learning for Creative Language Understanding*: UC Davis, TTIC, CMU, Dartmouth, Northeastern, UC Santa Cruz, UMass Amherst, WUSTL, Microsoft Research, Allen Institute for AI

Jan 2017: Talk at CU Boulder Stats, Optimization, and Machine Learning seminar

Nov 2016: Talk at UMass Amherst Machine Learning & Friends Lunch

External Service & Outreach

Senior Area Chair for ACL 2025, ACL 2024, ACL 2023, ACL 2022, COLING 2022, EMNLP 2021

Ethics Chair for COLM 2024

Area Chair for EMNLP 2024, ICLR 2024, EMNLP 2023, EMNLP 2022, ACL 2021, NAACL 2021, EMNLP 2020, ACL 2020, CoNLL 2019, EMNLP 2019

Program Committee/Reviewer for TACL, ACL, EMNLP, NAACL, ICLR, NIPS, ICML, CVPR

Sponsorship co-chair of NAACL 2024

Co-organizer of 5th Workshop on Narrative Understanding (WNU) at ACL 2023

Co-organizer of 4th Workshop on Narrative Understanding (WNU) at NAACL 2022

Co-organizer of 1st Workshop on Document-grounded Dialogue and Conversational QA (DialDoc) at ACL 2021

Co-organizer of 3rd Workshop on Narrative Understanding (WNU) at NAACL 2021

Co-organizer of 1st Joint Workshop on Narrative Understanding, Storylines, and Events (NUSE) at ACL 2020

Guest Editor of PLOS ONE Science of Stories Collection

Co-organizer of Narrative Understanding Workshop at NAACL 2019

Publicity Chair of EMNLP 2018

Co-organizer of Human-Computer Question Answering Competition at NIPS 2017

Co-organizer of Human-Computer Question Answering Workshop at NAACL 2016

Co-organizer of Mid-Atlantic Student Colloquium on Speech, Language and Learning (MASC-SLL) 2016

Talk at DC-NLP Meetup on "Deep Learning for Question Answering"

Introductory talk to high schoolers about the technology behind our quiz bowl system

Institutional Service

Internal awards committee (2021-)

Graduate program committee (2020)

Data science faculty recruiting committee (2019)

Annual faculty review committee (2019)

Graduate program committee (2018)

PhD admissions committee (2018)

Media

2024 TechCrunch article on our NoCha benchmark

2023 UMass Magazine feature on our literary machine translation work

2023 UMass press release covering our award-winning work on LongEval

Spoke about large language models and ChatGPT on local news stations in 2023:
GBH, WWLP (and again here), Western Mass News (and again here)

EMNLP 2019 football racial bias press: The Undefeated

CVPR 2017 comic books press: MIT Tech Review, Digital Trends

NAACL 2016 relationships press: Aeon

EMNLP 2014 quiz bowl press: UMIACS, Terp, Diamondback, Colorado CS

QANTA defeats Ken Jennings, 300-160

QANTA quiz bowl system exhibition match against a team of four former Jeopardy champions in May 2015 (result: 200-200 tie)

Patents

Richard Socher, Ankit Kumar, Ozan Irsoy, Mohit Iyyer, Caiming Xiong, Stephen Merity, Romain Paulus. *Dynamic memory network*. United States Patent #11113598. September 7, 2021.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, Luke Zettlemoyer. *System and methods for performing NLP related tasks using contextualized word representations*. United States Patent #11030414. June 8, 2021.