

Hangfeng He

Address: 250 Hutchison Rd, Rochester, NY 14620, USA
Tel: (585) 275-8848 Email: hangfeng.he@rochester.edu

Research Interests

My research interests include natural language processing and machine learning, with a focus on deep learning interpretability, reasoning in natural language, and evaluation of large language models.

Academic Positions

University of Rochester, Rochester, NY, USA 2022-present
Assistant Professor in Computer Science and Data Science

Education

University of Pennsylvania, Philadelphia, PA, USA 2017-2023
Ph.D. in Computer and Information Science
Advisors: Dan Roth and Weijie Su

Peking University, Beijing, China 2013-2017
B.S. in Computer Science, *Summa Cum Laude*

Publications

1. Xinyi Liu, Pinxin Liu, and **Hangfeng He**. 2024. An Empirical Analysis on Large Language Models in Debate Evaluation. In *The 62nd Annual Meeting of the Association for Computational Linguistics (ACL). Short papers*.

2. **Hangfeng He**, Hongming Zhang, and Dan Roth. 2024. SocREval: Large Language Models with the Socratic Method for Reference-Free Reasoning Evaluation. In *the 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL). Findings*.
3. Sindhu Kishore and **Hangfeng He**. 2024. Unveiling Divergent Inductive Biases of LLMs on Temporal Data. In *the 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL). Short papers*.
4. Matteo Sordello, Niccolo Dalmaso, **Hangfeng He**, and Weijie Su. 2024. Robust Learning Rate Selection for Stochastic Optimization via Splitting Diagnostic. In *Transactions on Machine Learning Research (TMLR)*.
5. **Hangfeng He** and Weijie Su. 2023. A Law of Data Separation in Deep Learning. In *Proceedings of the National Academy of Sciences (PNAS). Direct submission*.
6. Kaifu Wang, **Hangfeng He**, Tin Nguyen, Piyush Kumar, and Dan Roth. 2023. On Regularization and Inference with Label Constraints. In *International Conference on Machine Learning (ICML)*.
7. Mohammad Rostami, **Hangfeng He**, Muhao Chen, and Dan Roth. 2022. Transfer Learning via Representation Learning. In *Federated and Transfer Learning. Book Chapter*.
8. Shuxiao Chen, Koby Crammer, **Hangfeng He**, Dan Roth, and Weijie Su (**alphabetical order**). 2022. Weighted Training for Cross-Task Learning. In *International Conference on Learning Representations (ICLR). Oral presentation*.
9. Cong Fang, **Hangfeng He**, Qi Long, and Weijie Su (**alphabetical order**). 2021. Exploring Deep Neural Networks via Layer-Peeled Model: Minority Collapse in Imbalanced Training. In *Proceedings of the National Academy of Sciences (PNAS). Direct submission*.
10. **Hangfeng He**, Mingyuan Zhang, Qiang Ning, and Dan Roth. 2021. Foreseeing the Benefits of Incidental Supervision. In *Proceedings of the 2021 Conference on*

Empirical Methods in Natural Language Processing (EMNLP).

11. Zhun Deng, **Hangfeng He**, and Weijie Su. 2021. Toward Better Generalization Bounds with Locally Elastic Stability. In *International Conference on Machine Learning (ICML)*.
12. Ayal Klein, Jonathan Mamou, Valentina Pyatkin, Daniela Brook Weiss, **Hangfeng He**, Dan Roth, Luke Zettlemoyer, and Ido Dagan. 2020. QANom: Question-Answer driven SRL for Nominalizations. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING)*.
13. Shuxiao Chen, **Hangfeng He**, and Weijie Su (**alphabetical order**). 2020. Label-Aware Neural Tangent Kernel: Toward Better Generalization and Local Elasticity. In *Advances in Neural Information Processing Systems (NeurIPS)*.
14. Zhun Deng, **Hangfeng He**, Jiaoyang Huang, and Weijie Su. 2020. Towards Understanding the Dynamics of the First-Order Adversaries. In *International Conference on Machine Learning (ICML)*.
15. **Hangfeng He**, Qiang Ning, and Dan Roth. 2020. QuASE: Question-Answer Driven Sentence Encoding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
16. Soham Dan, **Hangfeng He**, and Dan Roth. 2020. Understanding Spatial Relations through Multiple Modalities. In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC). Short papers*.
17. **Hangfeng He** and Weijie Su. 2020. The Local Elasticity of Neural Networks. In *International Conference on Learning Representations (ICLR)*.
18. Qiang Ning, **Hangfeng He**, Chuchu Fan, and Dan Roth. 2019. Partial or Complete, That’s The Question. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
19. Jingjing Xu, **Hangfeng He**, Xu Sun, Xuancheng Ren, and Sujian Li. 2018. Cross-domain and semisupervised named entity recognition in chinese social

media: A unified model. In *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*.

20. Federico Fancellu, Adam Lopez, Bonnie Webber, and **Hangfeng He**. 2017. Detecting negation scope is easy, except when it isn't. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL). Short papers*.
21. **Hangfeng He** and Xu Sun. 2017. F-Score Driven Max Margin Neural Network for Named Entity Recognition in Chinese Social Media. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL). Short papers*.
22. **Hangfeng He** and Xu Sun. 2017. A Unified Model for Cross-Domain and Semi-Supervised Named Entity Recognition in Chinese Social Media. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
23. **Hangfeng He**, Federico Fancellu, and Bonnie Webber. 2017. Neural Networks for Negation Cue Detection in Chinese. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles (SemBEaR)*.

Preprints

1. **Hangfeng He**, and Weijie Su. 2024. A Law of Next-Token Prediction in Large Language Models. In *arXiv preprint*.
2. Hang Hua, Yunlong Tang, Ziyun Zeng, Liangliang Cao, Zhengyuan Yang, **Hangfeng He**, Chenliang Xu, and Jiebo Luo. MMCOMPOSITION: Revisiting the Compositionality of Pre-trained Vision-Language Models. In *arXiv preprint*.
3. **Hangfeng He**, Hongming Zhang, and Dan Roth. 2023. Rethinking with Retrieval: Faithful Large Language Model Inference. In *arXiv preprint*.

Teaching

CSC 511-01: Large Language Models

Fall 2024

CSC 247/447: Natural Language Processing

Spring 2023, Fall 2023

CSC 442: Artificial Intelligence

Fall 2022

Professional Service

Area Chair

IJCNLP-AAACL (2023)

Senior Program Committee Member

AAAI (2023)

Program Committee Member

ACL (2020-2021), COLING (2025), EMNLP (2019), NAACL (2019)

Conference Reviewer

ARR (2021), ICLR (2024), ICML (2021, 2023), NeurIPS (2020)

Journal Reviewer

IEEE TNNLS (2018-2020), TACL (2023)

University Service

Data Science Working Group Co-Chair	2023-present
Data Science M.S. Admission Committee	2022-present
Computer Science M.S. Admission Committee	2024-2025
Computer Science Graduate Education Committee	2024-2025
Computer Science Undergraduate Education Committee	2023-2024
Annual URCS Newsletter Committee	2023-2024
Data Science Faculty Search Committee	2022-2023
Computer Science Ph.D. Admission Committee	2022-2023

Invited Talks

A Law of Data Separation in Deep Learning

Invited Session Talk at JSM, August, 2023

Invited Session Talk at ICASA, June, 2023

Moving Beyond Scale-Driven Learning

Invited Talk at UR GIDS, October, 2022

Local Elasticity: A Phenomenological Approach Toward Understanding Deep Learning

Invited Talk at uOttawa TML seminar, November, 2021

Incidental Supervision for Natural Language Understanding

Invited Talk at USC/ISI AI Seminar, October, 2021