

Zhe Zeng

CONTACT INFORMATION

Phone: +1 (424) 535-8639

Address: 50 West 4th Street, New York, NY 10012

Email: zhezengzz@gmail.com

Website: <https://zzeng.me>

ACADEMIC APPOINTMENTS

Starting 07/2025	University of Virginia <i>Tenure-track Assistant Professor, Department of Computer Science</i>	Charlottesville, VA
09/2024 – 05/2025	New York University <i>Faculty Fellow/Assistant Professor, Computer Science Department</i>	New York, NY

EDUCATION

09/2018 – 06/2024	University of California, Los Angeles <i>Ph.D. in Computer Science</i> <i>Advisor: Guy Van den Broeck</i>	Los Angeles, CA
09/2014 – 07/2018	Zhejiang University <i>B.S. in Mathematics and Applied Mathematics with Honors</i>	Hangzhou, China

HONORS AND AWARDS

2023	Rising Star in EECS (Georgia Institute of Technology)
2022	Amazon Doctoral Student Fellowship
2021	NEC Student Research Fellowship
2018	ICML Student Travel Award
2018	Outstanding Graduate, Zhejiang University
2016	First-Class Scholarship for Elite Students in Basic Sciences, Zhejiang University
2015/2016	Excellent Student Award (top 5%), Zhejiang University

PUBLICATIONS

Note: * below denotes equal contribution

Peer-Reviewed Conference Publications

- [1] Chendi Qian, Andrei Manolache, Kareem Ahmed, **Zhe Zeng**, Guy Van den Broeck, Mathias Niepert, and Christopher Morris. Probabilistically rewired message-passing neural networks. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2024
- [2] **Zhe Zeng** and Guy Van den Broeck. Collapsed inference for bayesian deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023

- [3] Vinay Shukla, **Zhe Zeng**^{*}, Kareem Ahmed^{*}, and Guy Van den Broeck. A unified approach to count-based weakly-supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023
- [4] Kareem Ahmed^{*}, **Zhe Zeng**^{*}, Mathias Niepert, and Guy Van den Broeck. SIMPLE: A gradient estimator for k-subset sampling. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2023
- [5] Yizhuo Chen, Kaizhao Liang, **Zhe Zeng**, Shuochao Yao, and Huajie Shao. A unified knowledge distillation framework for deep directed graphical models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023
- [6] Wenzhe Li^{*}, **Zhe Zeng**^{*}, Antonio Vergari, and Guy Van den Broeck. Tractable computation of expected kernels. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021
- [7] **Zhe Zeng**^{*}, Paolo Morettin^{*}, Fanqi Yan^{*}, Antonio Vergari, and Guy Van den Broeck. Probabilistic inference with algebraic constraints: Theoretical limits and practical approximations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020 **Spotlight presentation, acceptance rate 280/9454 = 2.96%**
- [8] **Zhe Zeng**^{*}, Paolo Morettin^{*}, Fanqi Yan^{*}, Antonio Vergari, and Guy Van den Broeck. Scaling up hybrid probabilistic inference with logical and arithmetic constraints via message passing. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020
- [9] **Zhe Zeng** and Guy Van den Broeck. Efficient search-based weighted model integration. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019
- [10] Dilin Wang^{*}, **Zhe Zeng**^{*}, and Qiang Liu. Stein variational message passing for continuous graphical models. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2018

Peer-Reviewed Workshop Publications

- [11] Ruoyan Li, Dipti Ranjan Sahu, Guy Van den Broeck, and **Zhe Zeng**. Gradient estimation for exactly- k constraints. In *Proceedings of the NeurIPS Workshop on AI for Scientific Discovery: From Theory to Practice*, 2023
- [12] **Zhe Zeng** and Guy Van den Broeck. Collapsed inference for bayesian deep learning. In *Proceedings of the ICML Workshop on Structured Probabilistic Inference & Generative Modeling (SPIGM)*, 2023 **Oral Presentation (top 5%)**

- [13] Vinay Shukla, **Zhe Zeng**^{*}, Kareem Ahmed^{*}, and Guy Van den Broeck. A unified approach to count-based weakly-supervised learning. In *Proceedings of the ICML Workshop on Differentiable Almost Everything*, 2023
- [14] Chendi Qian, Andrei Manolache, Kareem Ahmed, **Zhe Zeng**, Guy Van den Broeck, Mathias Niepert, and Christopher Morris. Probabilistic task-adaptive graph rewiring. In *Proceedings of the ICML Workshop on Differentiable Almost Everything*, 2023
- [15] Kareem Ahmed^{*}, **Zhe Zeng**^{*}, Mathias Niepert, and Guy Van den Broeck. SIMPLE: A gradient estimator for k-subset sampling. In *Proceedings of the ICML Workshop on Differentiable Almost Everything*, 2023
- [16] Kareem Ahmed^{*}, **Zhe Zeng**^{*}, Mathias Niepert, and Guy Van den Broeck. SIMPLE: A gradient estimator for k-subset sampling. In *Southern California Natural Language Processing (SoCal NLP) Symposium*, 2022
- [17] **Zhe Zeng**^{*}, Paolo Morettin^{*}, Fanqi Yan^{*}, Antonio Vergari, and Guy Van den Broeck. Is parameter learning via weighted model integration tractable? In *Proceedings of the UAI Workshop on Tractable Probabilistic Modeling (TPM)*, 2021
- [18] **Zhe Zeng**^{*}, Paolo Morettin^{*}, Fanqi Yan^{*}, Antonio Vergari, and Guy Van den Broeck. Relax, compensate and then integrate. In *Proceedings of the ECML-PKDD Workshop on Deep Continuous-Discrete Machine Learning (DeCoDeML)*, 2020
- [19] **Zhe Zeng**^{*}, Fanqi Yan^{*}, Paolo Morettin^{*}, Antonio Vergari, and Guy Van den Broeck. Hybrid probabilistic inference with logical constraints: Tractability and message-passing. In *Proceedings of the NeurIPS Workshop on Knowledge Representation & Reasoning Meets Machine Learning*, 2019
- [20] **Zhe Zeng** and Guy Van den Broeck. Efficient search-based weighted model integration. In *Proceedings of the IJCAI Workshop on Declarative Learning Based Programming (DeLBP)*, 2019

RESEARCH POSITIONS

- 06/2022 – 09/2022 **IBM Almaden Research Center** San Jose, CA
Research Scientist Intern | Team: Scalable Knowledge Intelligence
Topic: Neuro-symbolic approaches for explainable sentence classification
- 06/2021 – 09/2021 **Yahoo! Research** New York, NY
Research Scientist Intern | Team: Scalable Machine Learning Group
Topic: Deep-learning based click-through rate prediction
- 07/2017 – 10/2017 **Dartmouth College** Hanover, NH
Undergraduate Researcher | Advisor: Professor Qiang Liu
Topic: Variational inference for continuous graphical models

INVITED TALKS

- 11/2023 **Georgia Institute of Technology**
EECS Rising Star Workshop
- 10/2023 **Simons Institute at University of California, Berkeley**
Probabilistic Circuits and Logic Workshop
- 10/2023 **University of California, Santa Cruz**
LINQS Statistical Relational Learning Group
- 07/2023 **Amazon Science**
- 05/2023 **Zhejiang University**
College of Computer Science and Technology
- 05/2023 **Stuttgart-RWTH-UCLA Workshop**
- 02/2023 **Amazon and UCLA**
Science Hub for Humanity and Artificial Intelligence
- 09/2022 **IBM Research – Almaden Lab**
Scalable Knowledge Intelligence Group
- 11/2021 **Microsoft Research New England Lab**
Machine Learning Seminar
- 04/2021 **Yahoo! Research**
Scalable Machine Learning Group
- 04/2021 **The Alan Turing Institute**
Statistics in Data-Centric Engineering Seminar

PROFESSIONAL ACTIVITIES AND SERVICE

Conference and Journal Reviewing

- International Journal of Approximate Reasoning (IJAR)
International Conference on Learning Representations (ICLR)
Conference on Uncertainty in Artificial Intelligence (UAI)
International Joint Conferences on Artificial Intelligence (IJCAI)
International Conference on Machine Learning (ICML)
International Conference on Artificial Intelligence and Statistics (AISTATS)
Conference on Neural Information Processing Systems (NeurIPS)

Association for the Advancement of Artificial Intelligence (AAAI)

Workshop Reviewing

- 2023 NeurIPS Workshop on Women in Machine Learning (WiML)
- 2023 ICML Workshop on Differentiable Almost Everything (DAE)
- 2023 ICML Workshop on Structured Probabilistic Inference & Generative Modeling (SPIGM)
- 2023 ICLR Workshop on Neurosymbolic Generative Models (NeSy-GeMs)
- 2022/2023 UAI Workshop on Tractable Probabilistic Modeling (TPM)
- 2022 NeurIPS Workshop on Human in the Loop Learning (HILL)

Discussant

- 2023 Conference on Uncertainty in Artificial Intelligence (UAI)

Volunteer

- 2023 Women in Machine Learning (WiML) Workshop at NeurIPS *as a super-volunteer*
- 2023 Women in Machine Learning (WiML) 2023–2024 Mentorship Program *as a mentor*
- 2023 Samuelli Undergraduate Research Program *as a panelist*
- 2023 Women in Machine Learning (WiML) Un-Workshop at ICML
- 2021 CSpH@UCLA Mentorship Program *as a mentor*

TEACHING EXPERIENCE

Instructor

- Fall 2024 CSCI-UA.0473 Fundamentals of Machine Learning at NYU

Guest Lecturer

- Spring 2023 CS267A Probabilistic Programming and Relational Learning at UCLA

Teaching Assistant at UCLA

- Spring 2024 CSC121/C221 Probabilistic Models in Computational Genomics
- Fall 2022/21/20 CS161 Fundamentals of Artificial Intelligence

Mentor/Supervisor

- 2023 Jiacheng Wang (UCLA undergraduate)
- 2023 Ruoyan Li (UCLA undergraduate)
- 2023 Dipti Ranjan Sahu (UCLA master student)
- 2021-2022 Vinay Shukla (UCLA undergraduate)
- 2020 Wenzhe Li (Tsinghua University undergraduate → Ph.D. student at Princeton University)
- 2019 Fanqi Yan (master student at Chinese Academy of Sciences → Ph.D. student at University of Texas at Austin)