Zhe Zeng

CONTACT INFORMATION

Phone: (424) 535-8639 Email: zhezeng@cs.ucla.edu Address: 368 ENG VI, UCLA, CA 90095 Website: http://web.cs.ucla.edu/~zhezeng

Research Interests

AI/ML with focus on NeuroSymbolic AI; Probabilistic Machine Learning; AI for Science

	Education	
09/2018 -	University of California, Los Angeles	Los Angeles, CA
06/2024 (expected)	Ph.D. in Computer Science	
	Advisor: Guy Van den Broeck	
09/2014 - 07/2018	Zhejiang University	Hangzhou, China
	B.S. in Mathematics and Applied Mathematics with Honors	
	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	g University
2015/2016	Excellent Student Award (top 5 %)	
	Publications	
	Note: * below denotes equal contribution	
	Peer-Reviewed Conference Publications	
[1]	Zhe Zeng and Guy Van den Broeck. Collapsed inference for bayesia learning. In Advances in Neural Information Processing Systems (Neu	an deep rIPS) 2023
[2]	Vinay Shukla, Zhe Zeng [*] , Kareem Ahmed [*] , and Guy Van den Broed approach to count-based weakly-supervised learning. In <i>Advances in</i> <i>Information Processing Systems (NeurIPS)</i> , 2023	ck. A unified n Neural

- [3] 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
- [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] 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 Aritifical Intelligence (UAI)*, 2021
- [6] 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%
- [7] 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
- [8] 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
- [9] 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

- [10] 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
- [11] 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%)
- [12] 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

- [13] 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
- [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 Southern California Natural Language Processing (SoCal NLP) Symposium, 2022
- [16] 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
- [17] 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
- [18] Zhe Zeng*, Fanqi Yan*, Paolo Morettin*, Antonio Vergari, and Guy Van den Broeck. Hybrid probabilistic inference with logical constraints: Tractability and message-passing. In Workshop on Knowledge Representation & Reasoning Meets Machine Learning at NeurIPS, 2019
- [19] 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

09/2018 – Present	University of California, Los Angeles	Los Angeles, CA
	Graduate Student Researcher Advisor: Professor Guy Van den Broec	k
	Topic: Artificial Intelligence and Machine Learning	
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	1
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 Samueli Undergraduate Research Program
- 2023 Women in Machine Learning (WiML) Un-Workshop at ICML
- 2021 CSPhD@UCLA Mentorship Program

TEACHING EXPERIENCE

Spring 2023	Guest Lecturer
	CS267A Probabilistic Programming and Relational Learning at UCLA
Fall 2022	Head Teaching Assistant
	CS161 Fundamentals of Artificial Intelligence at UCLA
Fall 2021	Teaching Assistant
	CS161 Fundamentals of Artificial Intelligence at UCLA
Fall 2020	Teaching Assistant
	CS161 Fundamentals of Artificial Intelligence at UCLA
	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, now Ph.D. student at Princeton
	University)
2019	Fanqi Yan (master student at Chinese Academy of Sciences, now Ph.D. student at
	University of Texas at Austin)