BENJIE WANG

UCLA - Department of Computer Science, 404 Westwood Plaza, Los Angeles, CA 90024
https://web.cs.ucla.edu/~benjiewang ◇ benjiewang@ucla.edu ◇ +1 (341) 766-9058

RESEARCH INTERESTS

My research interests are in the foundations of artificial intelligence and machine learning. In particular, my work focuses on the theory and design of *tractable* architectures for probabilistic generative AI that enable computationally efficient and exact probabilistic inference at scale. My research extends to applications where tractable and accurate inference is vital, including structured modeling, causality, constrained language generation, and robotics.

EDUCATION

DPhil (Ph.D.) in Computer Science, University of Oxford 2019 - 2023
Thesis: Tractable Probabilistic Models for Causal Learning and Reasoning
Advisor: Marta Kwiatkowska

MSc in Statistical Science, University of Oxford 2018 - 2019
Dissertation: Statistically Robust Neural Network Classification
Advisor: Tom Rainforth
Result: Distinction, 89%; **Gutierrez-Toscano Prize** (top-ranked in cohort)

BA in Mathematics, University of Cambridge 2015 - 2018
Result: First Class Honours, 82%

RESEARCH EXPERIENCE

Postdoctoral Scholar January 2024 - Present
University of California, Los Angeles *Los Angeles, CA*

Simons-Berkeley Research Fellow August - December 2023
Simons Institute for the Theory of Computing, UC Berkeley *Berkeley, CA*

Research Intern Summer 2023
Microsoft Research *Cambridge, UK*

Research Assistant 2020 - 2021
University of Oxford *Oxford, UK*

TEACHING EXPERIENCE

Guest Lecturer, CS161 Fundamentals of Artificial Intelligence at CS Dept, UCLA
Propositional Logic Winter 2024
Reasoning under Uncertainty Fall 2024

Class Tutor, Machine Learning at CS Dept, Oxford 2022 - 2023

Class Tutor, Quantum Information at CS Dept, Oxford 2021 - 2022

Class Tutor, Probabilistic Model Checking at CS Dept, Oxford 2021 - 2022

Class Tutor, Information Theory at Mathematical Institute, Oxford 2019 - 2020

Practical Demonstrator, Machine Learning at CS Dept, Oxford 2019 - 2020

PUBLICATIONS

Peer-Reviewed Journal Publications (* = equal contribution)

1. *PREMAP: A Unifying PREiMage APproximation Framework for Neural Networks.* Xiyue Zhang*, Benjie Wang*, Marta Kwiatkowska, Huan Zhang. Journal of Machine Learning Research (JMLR), 2025.

Peer-Reviewed Conference Publications (* = equal contribution)

14. *TRACE Back from the Future: A Probabilistic Reasoning Approach to Controllable Language Generation.* Gwen Yidou-Weng*, Benjie Wang*, Guy Van den Broeck. 42nd International Conference on Machine Learning (ICML), 2025.

13. *Scaling Probabilistic Circuits via Monarch Matrices.* Honghua Zhang*, Meihua Dang*, Benjie Wang*, Stefano Ermon, Nanyun Peng, Guy Van den Broeck. 42nd International Conference on Machine Learning (ICML), 2025.

12. *Restructuring Tractable Probabilistic Circuits.* Honghua Zhang*, Benjie Wang*, Marcelo Arenas, Guy Van den Broeck. 28th International Conference on Artificial Intelligence and Statistics (AISTATS), 2025. **Oral (2% acceptance).**

11. *On the Relationship Between Monotone and Squared Probabilistic Circuits.* Benjie Wang, Guy Van den Broeck. 39th AAAI Conference on Artificial Intelligence (AAAI), 2025.

10. *A Compositional Atlas for Algebraic Circuits.* Benjie Wang, Denis Deratani Mauá, Guy Van den Broeck, YooJung Choi. 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.

9. *Where is the signal in tokenization space?* Renato Lui Geh, Honghua Zhang, Kareem Ahmed, Benjie Wang, Guy Van den Broeck. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2024, **Oral (3.2% acceptance)**.

8. *Neural Structure Learning with Stochastic Differential Equations.* Benjie Wang*, Joel Jennings, Wenbo Gong*. Twelfth International Conference on Learning Representations (ICLR), 2024.

7. *Provable Preimage Under-approximation for Neural Networks.* Xiyue Zhang, Benjie Wang, and Marta Kwiatkowska. 30th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), 2024.

6. *Compositional Probabilistic and Causal Inference using Tractable Circuit Models.* Benjie Wang and Marta Kwiatkowska. 26th International Conference on Artificial Intelligence and Statistics (AISTATS), 2023.

5. *Tractable Uncertainty for Structure Learning.* Benjie Wang, Matthew Wicker, Marta Kwiatkowska. 39th International Conference on Machine Learning (ICML), 2022, **Oral (2% acceptance)**.

4. *Robustness Guarantees for Credal Bayesian Networks via Constraint Relaxation over Probabilistic Circuits.* Hjalmar Wijk, Benjie Wang, Marta Kwiatkowska. 31st International Joint Conference on Artificial Intelligence (IJCAI), 2022.

3. *Statistically Robust Neural Network Classification.* Benjie Wang, Stefan Webb, Tom Rainforth. 36th Conference on Uncertainty in Artificial Intelligence (UAI), 2021.

2. *Provable Guarantees on the Robustness of Decision Rules to Causal Interventions.* Benjie Wang*, Clare Lyle*, Marta Kwiatkowska. 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021.

1. *Assessing Robustness of Text Classification through Maximal Safe Radius Computation.* Emanuele La Malfa, Min Wu, Luca Laurenti, Benjie Wang, Anthony Hartshorn, Marta Kwiatkowska. Findings of the Association for Computational Linguistics: EMNLP, 2020.

Peer-Reviewed Workshop Publications and Working Papers

How to Marginalize in Causal Structure Learning? William Zhao, Guy Van den Broeck, Benjie Wang. Workshop on Graphs and more Complex Structures For Learning and Reasoning (GCLR) @ AAAI, 2026.

Conditional Causal Discovery. Cixuan Zhang, Benjie Wang. Workshop on Causal Abstractions and Representations (CAR) @ UAI, 2025.

Large Language Models as Tools to Improve Bayesian Causal Discovery. Bruna Bazaluk, Benjie Wang, Denis Mauá, Flavio S Correa da Silva. Workshop on Causal Abstractions and Representations (CAR) @ UAI, 2025.

Bitblasting for Constrained Decorrelation in Tractable Image Modeling. Poorva Garg, Benjie Wang, Oliver Broadrick, Guy Van den Broeck, Todd Millstein. 8th Workshop on Tractable Probabilistic Modeling (TPM) @ UAI, 2025.

A Circus of Circuits: Connections Between Decision Diagrams, Circuits, and Automata. Antoine Amarilli, Marcelo Arenas, YooJung Choi, Mikaël Monet, Guy Van den Broeck, Benjie Wang. Tech Report.

Probabilistic Circuits for Cumulative Distribution Functions. Oliver Broadrick, William Cao, Benjie Wang, Martin Trapp and Guy Van den Broeck. 7th Workshop on Tractable Probabilistic Modeling (TPM) @ UAI, 2024.

Bayesian network models of causal interventions in healthcare decision making: literature review and software evaluation. Artem Velikzhanin, Benjie Wang, Marta Kwiatkowska. Tech Report.

Symbolic Causal Inference via Operations on Probabilistic Circuits. Benjie Wang and Marta Kwiatkowska. Workshop on Neuro Causal and Symbolic AI (nCSI) @ NeurIPS, 2022.

Tractable Uncertainty for Structure Learning. Benjie Wang, Matthew Wicker, Marta Kwiatkowska. 5th Workshop on Tractable Probabilistic Modeling (TPM) @ UAI, 2022, **Best Paper Award**.

AWARDS

Simons-Berkeley Research Fellowship 2023
Research Fellowship associated with program on "Logic and Algorithms in Database Theory and AI".

Best Paper Award, 5th Workshop on Tractable Probabilistic Modeling @ UAI 2022
Awarded for "a novel tractable distribution over orderings with important applications to scaling causal discovery".

Gutierrez-Toscano Prize, Department of Statistics, University of Oxford 2019
Awarded for best performance in MSc Statistical Science.

Durham Prize, Keble College, University of Oxford 2019
Awarded for outstanding performance in Masters degree.

Honorary Bachelor Scholarship, Emmanuel College, University of Cambridge 2018
Awarded for outstanding performance in Bachelor's degree.

Braithwaite Batty Senior Scholarship, Emmanuel College, University of Cambridge 2017
Awarded for outstanding performance in Mathematics.

National Finalist, British Informatics Olympiad 2015
Among top 15 nationally.

IBM Prize, National Cipher Challenge 2014
Awarded to 2nd ranked team nationally.

ORGANIZATION AND ACADEMIC SERVICE

Co-organizer of the **7th Workshop on Tractable Probabilistic Modeling (TPM)** at UAI 2024

Journal and Conference Refereeing JMLR, IJAR, ICML (2022, 2023, 2024, 2025 - **Top Reviewer**), AISTATS (2021, 2023, 2024, 2025), AAAI 2025, NeurIPS (2021 - **Outstanding Reviewer Award**, 2022, 2024, 2025), ICLR 2024.

Volunteering Student Volunteer at UAI 2022; Volunteer Teaching Assistant in local schools (STIMULUS program - Cambridge, UK)

PRESENTATIONS

Probabilistic Reasoning in Large-Scale Models
UCSD Huang Lab — July 2025

Restructuring Tractable Probabilistic Circuits
International Conference on Artificial Intelligence and Statistics — May 2025
Simons Institute for the Theory of Computing, UC Berkeley — Jan 2025

Scaling Up Marginalization in Probabilistic Modeling
NYU Wilson Lab — Feb 2025

Tractable Deep Generative Models
National AI Research Resource (NAIRR) Annual Meeting — Feb 2025

Marginal Determinism in Structured Decomposable Circuits
Simons Institute for the Theory of Computing, UC Berkeley — Oct 2023

Compositional Probabilistic and Causal Inference
Causality Discussion Group — May 2023

Tractable Uncertainty for Structure Learning
Aalto University/ELLIS Seminar on Advances in Probabilistic Machine Learning — Feb 2023
International Conference on Machine Learning — Jul 2022

SUPERVISION AND MENTORING

Cixuan Zhang	Undergrad at UCLA → M.S. at Yale University
Ian Li	Undergrad at Harvey Mudd College→ Ph.D. at UCSD
William Zhao	Undergrad at UCLA
Sophie Q Li	Undergrad at UCLA
Bruna Bazaluk	MSc at University of São Paulo
Renato Lui Geh	Ph.D. at UCLA (mentoring)
Gwen Yidou-Weng	Ph.D. at UCLA (mentoring)
Hjalmar Wijk	DPhil at Oxford (mentoring) → Researcher, METR