

# Raha Moraffah

CAUSAL RESPONSIBLE MACHINE LEARNING, DATA SCIENCE & CAUSAL LEARNING

+1-480-652-3040 | ✉ [rmoraffah@wpi.edu](mailto:rmoraffah@wpi.edu) | 🏠 <https://rmoraffa.github.io> | 🎓 Google Scholar

## Employment

### Worcester Polytechnic Institute (WPI)

ASSISTANT PROFESSOR OF COMPUTER SCIENCE

Assistant Professor

Fall 2024- Present

## Education

### Arizona State University

DOCTOR OF PHILOSOPHY, COMPUTER SCIENCE

Advisor: Professor Huan Liu

Jan. 2017 - Present

### Sharif University of Technology

BACHELOR OF SCIENCE, COMPUTER SCIENCE

Advisor: Professor Hossein Sameti

Sep. 2011 - Aug. 2015

## Professional Experience

### Data Mining and Machine Learning Lab, ASU

ADVISOR: PROFESSOR HUAN LIU

Graduate Research Assistant

Mar. 2017 - Present

### Army Research Lab (ARL)

SUPERVISOR: DR. ADRIENNE RAGLIN

Graduate Researcher

May 2018 - Present

### USC Information Sciences Institute, USC

SUPERVISOR: DR. KRISTINA LERMAN

Research Intern

May 2018 - Sep. 2018

### Speech Processing Laboratory (SPL), Sharif University of Technology

SUPERVISOR: DR. HOSSEIN SAMETI

Undergraduate Researcher

Aug. 2013 - Aug. 2015

## Research Interests

### Responsible Machine Learning (esp. security and robustness)

- Generative AI and Foundation models
- Discriminative Deep Models

### Causal Inference (causality for responsible Machine Learning)

- Causality for security and robustness
- Causal Machine Learning

### AI for Social Good

- AI-generated text identification, hate speech detection, and disaster relief
- Cybersecurity
- Health care

## Publications

Raha Moraffah has co-authored over 40 papers in top-tier data mining, machine learning, and Natural Language Processing venues. As of 2/2024, her work has been cited over 550 times, with h-index = 10. Below are her selected publications: † denotes the author is her mentee; \* indicates an equal contribution.

## Books

- [1] **Raha Moraffah**, Saketh Vishnubhatla, and Huan Liu. "Causal Robust Machine Learning". *Spring Briefs* (2023).
- [2] **Raha Moraffah**, Shu Wan, and Huan Liu. "Causality: Learning to Discover with Data". *Springer Nature* (2023).

## Conferences , Journals , and Submissions

\* indicates equal contributions

† indicates the author is her mentee

- [3] Amrita Bhattacharjee<sup>†</sup>, **Raha Moraffah**, Joshua Garland, and Huan Liu. “EAGLE: A Domain Generalization Framework for AI-generated Text Detection”. *Submitted to European Conference on Machine Learning and Principles and Practice of Knowledge Discovery (ECML PKDD)* (2024).
- [4] Amrita Bhattacharjee<sup>\*†</sup>, **Raha Moraffah**, Joshua Garland, and Huan Liu. “Zero-shot LLM-guided Counterfactual Generation for Text”. *Submitted to European Conference on Artificial Intelligence (ECAI)* (2024).
- [5] Tharindu Kumarage<sup>\*†</sup>, Garima Agrawal, Paras Sheth<sup>\*†</sup>, **Raha Moraffah**, Aman Chadha, Joshua Garland, and Huan Liu. “A survey of ai-generated text forensic systems: Detection, attribution, and characterization”. *Submitted to SIGKDD Explorations* (2024).
- [6] **Raha Moraffah**, Shubh Khandelwal<sup>†</sup>, Amrita Bhattacharjee<sup>†</sup>, and Huan Liu. “Adversarial Text Purification: A Large Language Model Approach for Defense”. *PAKDD* (2024).
- [7] **Raha Moraffah** and Huan Liu. “A Generative Approach to Surrogate-based Black-box Attacks”. *Submitted to IJCAI* (2024).
- [8] **Raha Moraffah** and Huan Liu. “Exploiting Class Probabilities for Black-box Sentence-level Attacks”. *EACL* (2024).
- [9] **Raha Moraffah**, Paras Sheth<sup>\*†</sup>, Saketh Vishnubhatla<sup>\*†</sup>, and Huan Liu. “Causal Feature Selection for Responsible Machine Learning”. *Submitted to SIGKDD Explorations* (2024).
- [10] **Raha Moraffah**, Suraj Jyothi Unni, and Huan Liu. “Domain Generalization for VQA: A Cross-prompt Approach”. *Submitted to ECML* (2024).
- [11] **Raha Moraffah**, Isabel Valera, and Huan Liu. “Adversarial Transferability Through the Lens of Causality”. *Submitted to ICML* (2024).
- [12] **Raha Moraffah**, Chaowei Xiao, and Huan Liu. “What Should They Look Like? Reinforcing Surrogate-based Black-box Attacks with Distribution Feedback”. *Submitted to CVPR* (2024).
- [13] Paras Sheth<sup>†</sup>, Tharindu Kumarage, **Raha Moraffah**, Joshua Garland, and Huan Liu. “Cross-Platform Hate Speech Detection with Weakly Supervised Causal Disentanglement”. *Submitted to European Conference on Machine Learning and Principles and Practice of Knowledge Discovery (ECML PKDD)* (2024).
- [14] Paras Sheth<sup>†</sup>, **Raha Moraffah**, Tharindu Kumarage<sup>†</sup>, Aman Chadha, and Huan Liu. “Causality Guided Disentanglement for Cross-Platform Hate Speech Detection”. *ACM International Conference on Web Search and Data Mining (WSDM)* (2024).
- [15] Zhen Tan, Chengshuai Zhao, **Raha Moraffah**, Yifan L, Yu Kong, Tianlong Chen, and Huan Liu. “The Wolf Within: Covert Injection of Malice into MLLM Societies via an MLLM Operative”. *IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR) - ReGenAI* (2024).
- [16] Suraj Jyothi Unni<sup>†</sup>, **Raha Moraffah**, and Huan Liu. “VQA-GEN: A Visual Question Answering Benchmark for Domain Generalization”. *Submitted to ECML-PKDD* (2024).
- [17] Amrita Bhattacharjee<sup>†</sup>, Tharindu Kumarage<sup>†</sup>, **Raha Moraffah**, and Huan Liu. “ConDA: Contrastive Domain Adaptation for AI-generated Text Detection”. *IJCNLP-AAACL* (2023), **Outstanding Paper Award**.
- [18] Amrita Bhattacharjee<sup>†</sup>, **Raha Moraffah**, Joshua Garland, and Huan Liu. “Towards LLM-guided Causal Explainability for Black-box Text Classifiers”. *AAAI-ReLM* (2023).
- [19] Amrita Bhattacharjee<sup>†</sup>, **Raha Moraffah**, Joshua Garland, and Huan Liu. “Zero-shot Counterfactual Explanation Generation with LLMs for Text Classifiers”. *Submitted to ECML-PKDD* (2023).
- [20] Tharindu Kumarage<sup>†</sup>, Paras Sheth<sup>†</sup>, **Raha Moraffah**, Joshua Garland, and Huan Liu. “How Reliable Are AI-Generated-Text Detectors? An Assessment Framework Using Evasive Soft Prompts”. *EMNLP* (2023).
- [21] **Raha Moraffah**, Amir-Hossein Karimi, Adrienne Raglin, and Huan Liu. “Socially Responsible Machine Learning: A Causal Perspective”. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)* (2023).
- [22] Paras Sheth<sup>\*†</sup>, Tharindu Kumarage<sup>\*†</sup>, **Raha Moraffah**, Aman Chadha, and Huan Liu. “Peace: Cross-platform hate speech detection-a causality-guided framework”. *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD)* (2023).

- [23] Lu Cheng, Ruocheng Guo, **Raha Moraffah**, Paras Sheth, K Selçuk Candan, and Huan Liu. “Evaluation methods and measures for causal learning algorithms”. *IEEE Transactions on Artificial Intelligence* (2022).
- [24] **Raha Moraffah** and Huan Liu. “Query-Efficient Target-Agnostic Black-Box Attack”. *IEEE International Conference on Data Mining (ICDM)* (2022).
- [25] **Raha Moraffah**, Paras Sheth<sup>†</sup>, and Huan Liu. “Exploring the Target Distribution for Surrogate-Based Black-Box Attacks”. *IEEE International Conference on Big Data (Big Data)* (2022).
- [26] **Raha Moraffah**, Suraj Jyothi Unni<sup>†</sup>, Adrienne Raglin, and Huan Liu. “Causal Data Fusion for Multimodal Disaster Classification in Social Media”. *SBP-BRiMS* (2022).
- [27] Paras Sheth<sup>†</sup>, **Raha Moraffah**, K Selçuk Candan, Adrienne Raglin, and Huan Liu. “Domain Generalization—A Causal Perspective”. *arXiv preprint arXiv:2209.15177* (2022).
- [28] **Raha Moraffah**, Paras Sheth<sup>†</sup>, Mansooreh Karami<sup>†</sup>, Anchit Bhattacharya, Qianru Wang<sup>†</sup>, Anique Tahir, Adrienne Raglin, and Huan Liu. “Causal inference for time series analysis: Problems, methods and evaluation”. *Knowledge and Information Systems (KAIST)* (2021).
- [29] Bahman Moraffah, Christ Richmond, **Raha Moraffah**, and Antonia Papandreou-Suppappola. “Metric-bayes: Measurements estimation for tracking in high clutter using bayesian nonparametrics”. *Asilomar Conference on Signals, Systems, and Computers* (2020).
- [30] Bahman Moraffah, Christ Richmond, **Raha Moraffah**, and Antonia Papandreou-Suppappola. “Use of bayesian nonparametric methods for estimating the measurements in high clutter”. *arXiv preprint arXiv:2012.09785* (2020).
- [31] **Raha Moraffah**, Mansooreh Karami<sup>†</sup>, Ruocheng Guo, Adrienne Raglin, and Huan Liu. “Causal interpretability for machine learning-problems, methods and evaluation”. *ACM SIGKDD Explorations* (2020).
- [32] **Raha Moraffah**, Bahman Moraffah, Mansooreh Karami, Adrienne Raglin, and Huan Liu. “CAN: A Causal Adversarial Network for Learning Observational and Interventional Distributions”. <https://arxiv.org/abs/2008.11376> (2020).
- [33] Adrienne Raglin, **Raha Moraffah**, and Huan Liu. “Causality and Uncertainty of Information for Content Understanding”. *IEEE International Conference on Cognitive Machine Intelligence (CogMI)* (2020).
- [34] Lu Cheng, Ruocheng Guo\*, **Raha Moraffah\***, K Selçuk Candan, Adrienne Raglin, and Huan Liu. “A practical data repository for causal learning with big data”. *Benchmarking, Measuring, and Optimizing: Second BenchCouncil International Symposium, Bench* (2019).
- [35] **Raha Moraffah**, Kai Shu, Adrienne Raglin, and Huan Liu. “Deep causal representation learning for unsupervised domain adaptation”. *arXiv preprint arXiv:1910.12417* (2019).
- [36] Vineeth Rakesh\*, Ruocheng Guo\*, **Raha Moraffah**, Nitin Agarwal, and Huan Liu. “Linked causal variational autoencoder for inferring paired spillover effects”. *ACM International Conference on Information and Knowledge Management (CIKM)* (2018).
- [37] Mohamed Sarwat, **Raha Moraffah**, Mohamed F Mokbel, and James L Avery. “Database system support for personalized recommendation applications”. *International Conference on Data Engineering (ICDE)* (2017).
- [38] Jia Yu, **Raha Moraffah**, and Mohamed Sarwat. “Hippo in action: Scalable indexing of a billion new york city taxi trips and beyond”. *IEEE 33rd International Conference on Data Engineering (ICDE)* (2017).

## Peer-reviewed Posters (with poster papers and abstracts)

- [39] **Raha Moraffah**, Suraj Jyothi Unni, Adrienne Raglin, and Huan Liu. “CAUSEMMD: A multi-modal classification platform for disaster relief”. *SBP-BRiMS* (2022).
- [40] **Raha Moraffah**, Lu Cheng, Ruocheng Guo, and Huan Liu. “CAUSE: A data repository for causal inference from observational data”. *SBP-BRiMS* (2019).

## Proposal Writing

---

### DISCOVER: Decision Intelligence System for Causal Exploration and Verification

PI: DR. HUAN LIU

- Designed and wrote two out of three tasks: “Causal Decision Exploration before Action” and “Causal Action Analysis after Action”
- Submitted to AMC in 2023

## Elements: CausalBench: A Cyberinfrastructure for Causal-Learning Benchmarking for Efficacy, Reproducibility, and Scientific Collaboration

PI: DR. SELCUK CANDAN, Co-PI: DR. HUAN LIU

- Designed and wrote one task out of 4 tasks “Open Sourcing and Validation”
- NSF (#2311716)

## High-Performance Computing Cluster for NLP and Deep Learning Approaches to Understanding Social, Psychological, and Rhetorical Dimensions of Influence in Cross-Cultural Contexts

PI: DR. SCOTT RUSTON, Co-PI: DR. HUAN LIU

- Design and wrote one task out of 3 for “Benchmarking Causal Discovery”
- DURIP ONR

## Characterizing and Understanding of Dynamic Time-Series Information for Decision Making: A Causal Learning Approach

PI: DR. HUAN LIU

- Designed and wrote three tasks out of five tasks “Causal discovery for time series”, “Causal inference for time series”, and “Causal Interpretable models for time series prediction”
- AMC (ARL W911NF2020124)

## Knowing the Unknown: Exploring the Space of Adversarial Attacks via Causal Learning

PI: DR. HUAN LIU

- Designed and wrote 2 out of 4 tasks “Causal Learning for Exploring the Space of Adversarial Examples” and “Coverage of Causal Learning and Transferability of Adversarial Examples”
- AMC (ARO W911NF2110030 )

## Presentation

### Invited Talks at Universities and Summer Schools

- Causal feature selection in the era of big data, Causal Inference for Time Series Workshop, UAI July. 2024
- Causal feature selection in the era of big data, SIAM Annual Meeting July. 2024
- Causal Adversarial Attacks: A New Frontier, Max Plank Institute, Oct. 2023
- Pragmatic Responsible ML: From Foundation to Application, Sharif University of Technology Aug. 2023
- A mini-lecture series on Introduction to Causal Inference, ASU Aug. 2023
- Causal Machine Learning: Issues, Purposes, and Challenges, Sharif University of Technology Jul. 2022
- Lecture series on Causal Machine Learning for Adversarial Robustness, ASU Nov. 2023
- A mini-lecture series on Causal Learning from Observational Data, ASU Oct. 2023
- Lectures on Causal Inference, Army Research Lab Mar. 2018- Present

### Keynote/Invited Talks at Conferences and Workshops

- Causal Feature Selection in the Era of Big Data, Minisymposium on Learning from Data with Information Theory and Causal Inference, SIAM Annual Meeting July. 2024
- Responsible for Responsible Generative AI and Causality Discussion Sessions, Workshop on Generative AI and Large Language Models: Opportunities and Challenges for Healthcare and Medicine, ASU-Mayo Clinic Nov. 2023
- Tutorial on Socially Responsible Machine Learning: A Causal Perspective, KDD 2023 Aug. 2023
- PhD Doctrol Forumn, SDM 2023 Nov. 2023
- Tutorial on Machine Learning for Causal Inference, SBP-BRIMS 2022 Sep. 2022

## Students and Mentees

### Brian Vincent, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Aug. 2023- Present

- Research: Causal Feature Selection with Application in E-commerce

### Shu Wan, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

July 2023- Present

- Research: Large language models for causal discovery: A data-driven approach

### Paras Sheth, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Sep. 2022- Present

- Co-authored 4 papers, one accepted to IEEE Bigdata 2022, one accepted to ECML-PKDD 2023, one accepted to WSDM 2024, and one submitted to JAIR.

### Kumarage Tharindu Kumarage, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Sep. 2022- Present

- Co-authored “How Reliable Are AI-Generated-Text Detectors? An Assessment Framework Using Evasive Soft Prompts”, Accepted to EMNLP 2023

### Amrita Bhattacharjee, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Sep. 2022- Present

- Co-authored 4 papers, one accepted and won the outstanding paper award at IJCNLP-AACL 2023, one accepted to AAAI-ReLM, one submitted to Lrec-Coling 2024, and one submitted to PAKDD
- Won best paper award

### Ali Mohammad Beigi, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Sep. 2022- Present

- Causal feature selection

### Mansoor Karami, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Jan 2019-Nov.2023

- Co-authored 3 papers, one accepted at SIGKDD Explorations, one accepted to KAIST 2021
- Joined Microsoft Research as a Data Scientist

### Qianru Wang, Visiting Ph.D. student

SCHOOL OF COMPUTER SCIENCE, NORTHWESTERN POLYTECHNICAL UNIVERSITY

Sep. 2019-Sep. 2020

- Co-authored one paper accepted to KAIST 2021

### Anchit Bhattacharya, Ph.D. student

DATA MINING AND MACHINE LEARNING LAB, ASU

Jun. 2019-Sep. 2020

- Co-authored one paper accepted to KAIST 2021
- Joined University of Kentucky as a Ph.D. Student

### Saketh Vishnubhatla, Master's student

DATA MINING AND MACHINE LEARNING LAB, ASU

Sep. 2023- Present

- Research: Causal feature selection for responsible ML
- Coauthored one paper to be submitted to IJCAI 2024

### Somayeh Khazaei, Visiting Master's student

UNIVERSITY OF LUXEMBOURG

Mar. 2023- Present

- Research: Causal adversarial attacks

### Suraj Jyoti Unni, Master's student

DATA MINING AND MACHINE LEARNING LAB, ASU

Feb. 2019- Present

- Co-authored two papers, one submitted to SIGIR 2024, and one submitted to ECCV 2024
- Successfully defended his master's and cofounded a company

### Ishita Mathur, Undergraduate student

CALIFORNIA INSTITUTE OF TECHNOLOGY (CALTECH)

Jan. 2023- Present

- Research: Textual adversarial attacks

### Shubh Khandelwal, Undergraduate student

ARIZONA STATE UNIVERSITY (ASU)

Mar. 2023- Present

- Research: Textual adversarial defense
- Coauthored a paper submitted to PAKDD 2024

### Ryan Schwartz, Undergraduate student

ARIZONA STATE UNIVERSITY (ASU) - NSF REU

Jan. 2019- Sep. 2020

- Research: Benchmarking adversarial attack

### Qiang Fu, Undergraduate student

ARIZONA STATE UNIVERSITY (ASU) - FURI

Jan. 2019- Sep. 2019

- Received the ASU Fulton School of Engineering Impact Award for his dissertation on "Benchmarking adversarial attack"
- Joined Google as a Software Engineer

## Teaching Experience

### CSE 472: Social Media Mining

Arizona State University

GUEST LECTURER - TAUGHT 7 LECTURES ON CAUSAL INFERENCE IN ML

Fall 2020- Fall 2023

### Socially Responsible Machine Learning: A Causal Perspective

KDD 2023

CONFERENCE TUTOR

Fall 2023

### Machine Learning for Causal Inference

SBP-BRIMS 2022

CONFERENCE TUTOR

Fall 2023

### CSE 472: Social Media Mining

Arizona State University

TEACHING ASSISTANT

Fall 2019- Fall 2023

## Editorial Boards and Reviewing Activities

### Senior Program Committee

Association for the Advancement of Artificial Intelligence (AAAI, Main Track) (2024), ACM International Conference on Web Search and Data Mining (WSDM) (2023)

### Program Committee Member

ICLR(2020-2024), NeurIPS(2020-2024), ICML(2021-2024), CVPR (2022-2023), AAAI (2020 - 2023), ACM SIGKDD (2020-2023), ICDM (2020-2023), IEEE Big data (2020-2023), (ECML-PKDD) (2020-2023), (SBP-BRIMS) (2019-2023), IEEE CogMI (2021-2023), ACL ARR (2022 - 2023), IJCAI (2021 - 2024), ACM MM (2021 - 2022), (CIKM) (2019-2020), WSDM (2022-2023), ICRA (2022), ECCV (2022), ICME (2022), EMNLP (2022), ICCV (2021), ACM SIGIR (2020), etc.

### Journal Review

JMLR, TPAMI, TIP, TCSVT, TNNLS, TMLR, TC, IPM, MACH, TKDE, ISIT, IEEE Access, Computers & Security (COSE), IEEE Transactions on Computational Social Systems (TCSS), IEEE Transactions on Artificial Intelligence (TAI), ACM Computing Survey (CSUR), ACM Transactions on Knowledge Discovery from Data (TKDD), Data & Policy, etc.

## Awards & Honors

2024	ASU's nominee for the "CGS/ProQuest Distinguished Dissertation Awards 2024" in the category of Mathematics, Physical Science & Engineering	
2024	Dean's Dissertation Award in Fulton Schools of Engineering	ASU
2024	SCAI Doctoral Fellowship for Excellent Research (Highest Distinction)	ASU
2023	Best Paper Award , ASU SCAI Symposiom	ASU
2023	Outstanding Paper Award , IJCNLP-AAACL	Bali, Indonesia
2022-2023	Doctoral Fellowship award , ASU SCAI	ASU
2022-2023	Travel award , ASU SCAI	ASU
2023	Travel Grant Award , KDD	Long Beach, CA
2023	Doctoral Forum Poster First Runner Up , SDM	Minneapolis, MN
2023	Travel Grant Award , SDM	Minneapolis, MN
2022	Travel Grant Award , ICDM	Orlando, FL
2022-2023	GPSA Individual Travel Funding	ASU
2022	Travel Grant Award , SBP-BRIMS	Pittsburgh, PA
2020	Graduate Outstanding Research Award	ASU
2019	Travel Grant Award , SBP-BRIMS	Washington, D.C.
2018	Grace Hopper Celebration Scholarship	ASU
2017	Travel Grant Award , ICDE	San Diego, CA
2014-2015	Outstanding Undergraduate Award	Tehran, IR
2011-2015	Undergraduate Scholarship	Tehran, IR
2010	Ranked 12 <sup>th</sup> out of 600,000 in Mathematics Nationwide Entrance Qualification Exam (Konkour)	Tehran, IR
2009	Silver Medalist of Iranian National Mathematics Olympiad	Tehran, IR
2008-Present	Honorary Member of the Council of Young Mathematician Society	Tehran, IR
2008	Bronze Medal in Fryer, Galois, and Hypatia Mathematics Contests , University of Waterloo	Waterloo, CA
2007-Present	Member of Young Mathematician Society	Tehran, IR