

DANIEL LOWD

Department of Computer Science
University of Oregon
Eugene, OR 97403

Tel. (541) 346-4154
lowd@cs.uoregon.edu
<http://ix.cs.uoregon.edu/~lowd>

INTERESTS

Machine learning, data mining, and artificial intelligence.

EDUCATION

2003–2010: **University of Washington**, Seattle, WA

Ph.D. in Computer Science and Engineering, March 2010

M.S. in Computer Science and Engineering, June 2005

Advisor: Pedro Domingos.

Committee members: Jeff Bilmes, Marina Meila, and Mark S. Handcock.

Dissertation title: *Efficient Algorithms for Learning and Inference in Rich, Statistical Models*.

1999–2003: **Harvey Mudd College**, Claremont, CA.

B.S. in Mathematics/Computer Science with “High Honors.”

PROFESSIONAL EXPERIENCE

September 2024 – Present: Professor, Department of Computer Science, University of Oregon.

September 2017 – September 2024: Associate Professor, Department of Computer and Information Science, University of Oregon.

June 2010 – September 2017: Assistant Professor, Department of Computer and Information Science, University of Oregon.

September 2009 – June 2010: Acting Assistant Professor, Department of Computer and Information Science, University of Oregon.

September 2003 – August 2009: Research assistant for Pedro Domingos, University of Washington. Developed faster machine learning algorithms, more flexible probabilistic models, and methods for learning models with efficient inference.

Summer 2008: Intern at SmartDesktop division of Pi Corporation, Seattle, WA. Developed and evaluated automatic methods for desktop activity recognition.

Summer 2004: Intern at Microsoft Research with Christopher Meek in Redmond, WA. Developed simple yet effective attacks against linear spam filters, testing filter robustness and promoting the development of more secure spam filters.

June 2002 – September 2003: Research assistant for Jon Herlocker, Oregon State University. Conducted a third-party evaluation of prominent collaborative filtering algorithms. Funded by the NSF through the REU program.

Summer 2001: Intern at Green Hills Software in Santa Barbara, CA. Invented and implemented a binary diff algorithm, reducing code update time by 90–97% for embedded targets. Ported a linker from Solaris to vxWorks and reduced code size by 80% for

use on JPL’s Mars Expedition Rover. Ran automated compiler validations for R6K architecture and implemented appropriate fixes in library code.

Summer 2000: Intern at Adobe Systems in San Jose, CA.

Added XML support to Adobe Acrobat’s Webcapture plug-in.

Summer 1999: Intern at Spyglass in Los Gatos, CA.

Assisted in the development of a Windows CE application for physical therapists.

SCHOLARSHIPS, HONORS, AND AWARDS

Best Poster Paper Award Nominee, EDM 2025

Area Chair Favorite Paper, COLING 2018

ICML 2016 Outstanding Reviewer Award (2016)

Army Research Office Young Investigator Award (2015)

Best Paper Award, DEXA Conference (2015)

Google Faculty Research Award (2013)

Microsoft Research Fellow, sponsored by Live Labs (2007–2008)

National Science Foundation Graduate Research Fellowship (2003–2006)

BOOKS

1. P. Domingos and D. Lowd, *Markov Logic: An Interface Layer for AI*, Morgan & Claypool. 2009. 155 pages.

BOOK CHAPTERS

2. S. Natarajan, J. Davis, K. Kersting, P. Domingos, and D. Lowd, “Statistical Relational Learning,” in *An Introduction to Lifted Probabilistic Inference* (18 pages), edited by G. Van den Broeck, K. Kersting, S. Natarajan, and D. Poole, MIT Press, 2021.
3. P. Domingos, D. Lowd, S. Kok, H. Poon, M. Richardson and P. Singla, “Markov Logic: A Language and Algorithms for Link Mining,” in *Link Mining: Models, Algorithms, and Applications* (pp. 135–161), edited by P. S. Yu, J. Han, and C. Faloutsos, Springer, 2010.
4. P. Domingos, D. Lowd, S. Kok, H. Poon, M. Richardson and P. Singla, “Just Add Weights: Markov Logic for the Semantic Web,” in *Uncertainty Reasoning for the Semantic Web I* (pp. 1–25), edited by P. Costa, C. d’Amato, N. Fanizzi, K. B. Laskey, K. J. Laskey, T. Lukasiewicz, M. Nickles and M. Pool, Springer, 2008.
5. P. Domingos, S. Kok, D. Lowd, H. Poon, M. Richardson and P. Singla, “Markov Logic,” in *Probabilistic Inductive Logic Programming* (pp. 92–117), edited by L. De Raedt, P. Frasconi, K. Kersting and S. Muggleton, Springer, 2008.

JOURNAL ARTICLES

6. Z. Hammoudeh and D. Lowd, “Training data influence analysis and estimation: A survey,” in *Machine Learning Journal*, (pp. 2351–2403), vol.113, 2024.
7. J. Brophy, Z. Hammoudeh, and D. Lowd, “Adapting and Evaluating Influence-Estimation Methods for Gradient-Boosted Decision Trees,” in the *Journal of Machine Learning Research* (pp. 1–48), vol.24, 2023.

8. P. Domingos and D. Lowd, “Unifying Logical and Statistical AI with Markov Logic,” in *Communications of the ACM* (pp. 74–83), vol.62, no.7, 2019.
9. S. Jiang, D. Lowd, S. Kafle, and D. Dou, “Ontology Matching with Knowledge Rules,” in *Transactions on Large-Scale Data- and Knowledge-Centered Systems* (pp. 75–95), vol.28, 2016.
10. D. Lowd and A. Rooshenas, “The Libra Toolkit for Probabilistic Models,” in the *Journal of Machine Learning Research* (pp. 2459–2463), vol.16, 2015.
11. D. Lowd and J. Davis, “Improving Markov Network Structure Learning Using Decision Trees,” in the *Journal of Machine Learning Research* (pp. 501–532), vol.15, 2014.

REFEREED CONFERENCE PUBLICATIONS

12. W. You and D. Lowd, “The Ultimate Cookbook for Invisible Poison: Crafting Subtle Clean-Label Text Backdoors with Style Attributes,” in *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, (pp. 222–246), 2025.
13. Z. Hammoudeh and D. Lowd, “Provable Robustness Against a Union of L_0 Adversarial Attacks,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2024.
14. W. You, Z. Hammoudeh, and D. Lowd. “Large Language Models Are Better Adversaries: Exploring Generative Clean-Label Backdoor Attacks Against Text Classifiers,” in *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 12499–12527), Singapore, 2024.
15. A. Baruwa, S. Sokolowski, J. Searcy, and D. Lowd. “Machine Learning to Define Anthropometric Landmarks for Relevant Product Design 2D Blueprint Measures,” in *Proceedings of the Applied Human Factors and Ergonomics (AHFE) 2023 International Conference*, 2023.
16. Z. Hammoudeh and D. Lowd, “Reducing Certified Regression to Certified Classification for General Poisoning Attacks,” in *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2023. (Acceptance rate: 26%.)
17. J. Brophy and D. Lowd, “Instance-Based Uncertainty Estimation for Gradient-Boosted Regression Trees,” in *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022. (Acceptance rate: 26%.)
18. Z. Hammoudeh and D. Lowd, “Identifying a Training-Set Attack’s Target Using Renormalized Influence Estimation,” in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, (pp. 1367–1381), 2022. (Acceptance rate: 23%.)
19. J. Brophy and D. Lowd, “Machine unlearning for random forests,” in *Proceedings on the International Conference on Machine Learning (ICML-21)*, (pp. 1092–1104), 2021. (Acceptance rate: 22%.)
20. Z. Hammoudeh and D. Lowd, “Learning from positive and unlabeled data with arbitrary positive shift,” in *Advances in Neural Information Processing Systems 33 (NeurIPS)*, (pp. 13088–13099), 2020. (Acceptance rate: 20%.)
21. S. Jamshidi, Z. Hammoudeh, R. Durairajan, D. Lowd, R. Rejaie, and W. Willinger, “On the Practicality of Learning Models for Network Telemetry,” in *Proceedings of the 4th Network Traffic Measurement and Analysis Conference (TMA 2020)*, Berlin, Germany (Virtual), 2020. (Acceptance rate: 33%.)

22. J. Ebrahimi, D. Lowd, and D. Dou, “On Adversarial Examples for Character-Level Neural Machine Translation,” in *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*, (pp. 653–663), Santa Fe, New Mexico, USA, 2018. **Area Chair Favorite.** (Acceptance rate: 3.8% Area Chair Favorites; 37% overall.)
23. J. Ebrahimi, A. Rao, D. Lowd, and D. Dou, “HotFlip: White-Box Adversarial Examples for Text Classification,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)* (short paper), (pp. 31–36), Melbourne, Australia, 2018. (Acceptance rate: 24%.)
24. A. Pouran Ben Veyseh, J. Ebrahimi, D. Dou, and D. Lowd, “A Temporal Attentional Model for Rumor Stance Classification,” in *Proceedings of the 26th ACM International Conference on Information and Knowledge Management (CIKM)* (short paper), (pp. 2335–2338), Singapore, 2017. (Acceptance rate: 28% short papers.)
25. J. Ebrahimi, D. Dou, and D. Lowd, “A Joint Sentiment-Target-Stance Model for Stance Classification in Tweets,” in *Proceedings of the 26th International Conference on Computational Linguistics (COLING 2016)* (pp. 2656–2665), 2016. (Acceptance rate: 32%.)
26. J. Ebrahimi, D. Dou, and D. Lowd, “Weakly Supervised Tweet Stance Classification by Relational Bootstrapping,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)* (pp. 1012–1017) (short paper), 2016. (Acceptance rate: 25%.)
27. H. Wang, D. Dou, and D. Lowd, “Ontology-based Deep Restricted Boltzmann Machine,” in *Proceedings of the 27th International Conference on Database and Expert Systems Applications (DEXA 2016)* (pp. 431–445), Porto, Portugal, 2016.
28. A. Rooshenas and D. Lowd, “Discriminative Structure Learning of Arithmetic Circuits,” in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)* (pp. 1506–1514), Cadiz, Spain, 2016. (Acceptance rate: 31%.)
29. S. Jiang, D. Lowd, and D. Dou, “A Probabilistic Approach to Knowledge Translation,” in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)* (pp. 1716–1722), Phoenix, AZ, USA, 2016. (Acceptance rate: 26%.)
30. S. Jiang, D. Lowd, and D. Dou, “Ontology Matching with Knowledge Rules,” in *Proceedings of the 26th International Conference on Database and Expert Systems Applications (DEXA 2015)* (pp. 94–108), Valencia, Spain, 2015. **Best Paper Award.**
31. R. Motamedi, R. Rejaie, W. Willinger, D. Lowd, and R. Gonzalez, “Inferring Coarse Views of Connectivity in Very Large Graphs,” in *Proceedings of the ACM Conference on Online Social Networks (COSN)* (pp. 191–202), Dublin, Ireland, 2014. (Acceptance rate: 16%.)
32. M. Torkamani and D. Lowd, “On Robustness and Regularization of Structural Support Vector Machines,” in *Proceedings of the 31st International Conference on Machine Learning (ICML)* (pp. 577–585), Beijing, China, 2014. (Acceptance rate: 25%.)
33. A. Rooshenas and D. Lowd, “Learning Sum-Product Networks with Direct and Indirect Interactions,” in *Proceedings of the 31st International Conference on Machine Learning (ICML)* (pp. 710–718), Beijing, China, 2014. (Acceptance rate: 25%.)
34. A. Bates, R. Leonard, H. Pruse, D. Lowd, and K. Butler, “Leveraging USB to Establish Host Identity Using Commodity Devices,” in *Proceedings of the 21st ISOC Network and Distributed System Security Symposium (NDSS)* (14 pages), San Diego, CA, USA, 2014. (Acceptance rate: 19%.)

35. D. Lowd and A. Rooshenas, “Learning Markov Networks with Arithmetic Circuits,” in *Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)* (pp. 406–414), Scottsdale, AZ, USA, 2013. (Acceptance rate: 34%.)
36. M. Torkamani and D. Lowd, “Convex Adversarial Collective Classification,” in *Proceedings of the 30th International Conference on Machine Learning (ICML)* (pp. 642–650), Atlanta, GA, USA, 2013. (Oral. Acceptance rate: 12% for oral presentations.)
37. S. Jiang, D. Lowd, and D. Dou, “Learning to Refine an Automatically Extracted Knowledge Base using Markov Logic,” in *Proceedings of the IEEE International Conference on Data Mining (ICDM)* (pp. 912–917), Brussels, Belgium, 2012. (Acceptance rate: 20%.)
38. D. Lowd, “Closed-Form Learning of Markov Networks from Dependency Networks,” in *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI-12)* (pp. 533–542), Catalina Island, CA, 2012. (Spotlight. Acceptance rate: 32% overall; 14% with a talk or spotlight.)
39. D. Lowd and A. Shamaei, “Mean Field Inference in Dependency Networks: An Empirical Study,” in *Proceedings of the 25th Conference on Artificial Intelligence (AAAI-11)* (pp. 404–410), San Francisco, CA, 2011. (Acceptance rate: 25%.)
40. D. Lowd and P. Domingos, “Approximate Inference by Compilation to Arithmetic Circuits,” in *Advances in Neural Information Processing Systems (NIPS) 24* (pp. 1477–1485), Vancouver, BC, 2010. (Acceptance rate: 24%.)
41. D. Lowd and J. Davis, “Learning Markov Network Structure with Decision Trees,” in *Proceedings of the 10th IEEE International Conference on Data Mining (ICDM)* (pp. 334–343), Sydney, Australia, 2010. (Full paper. Acceptance rate: 19% overall; 9% for full papers.)
42. S. Natarajan, T. Khot, D. Lowd, K. Kersting, P. Tadepalli and J. Shavlik, “Exploiting Causal Independence in Markov Logic Networks: Combining Undirected and Directed Models,” in *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)* (pp. 434–450), Barcelona, Spain, 2010. (Acceptance rate: 16%.)
43. D. Lowd and N. Kushmerick, “Using Saliency to Segment Desktop Activity into Projects,” in *Proceedings of the Thirteenth International Conference on Intelligent User Interfaces* (pp. 463–468), Sanibel Island, Florida, 2009. (Short paper. Acceptance rate: 31%.)
44. D. Lowd and P. Domingos, “Learning Arithmetic Circuits,” in *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence* (pp. 383–392), Helsinki, Finland, 2008.
45. D. Lowd and P. Domingos, “Efficient Weight Learning for Markov Logic Networks,” in *Proceedings of the Eleventh European Conference on Principles and Practice of Knowledge Discovery in Databases* (pp. 200–211), Warsaw, Poland, 2007.
46. D. Lowd and P. Domingos, “Recursive Random Fields,” in *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence* (pp. 950–955), Hyderabad, India, 2007. (Oral presentation. Acceptance rate: 16%.)
47. D. Lowd and P. Domingos, “Naive Bayes Models for Probability Estimation,” in *Proceedings of the Twenty-Second International Conference on Machine Learning* (pp. 529–536), Bonn, Germany, 2005. (Acceptance rate: 27%.)
48. D. Lowd and C. Meek, “Good Word Attacks on Statistical Spam Filters,” in *Proceedings of the Second Conference on Email and Anti-Spam* (8 pages), Palo Alto, CA, 2005.

49. D. Lowd and C. Meek, “Adversarial Learning,” in *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 641–647), Chicago, IL, 2005. (Acceptance rate: 21.2%.)

REFEREED WORKSHOP AND POSTER PUBLICATIONS

50. A. Kabir, C. Tankala, and D. Lowd, “On the Practicality of Differential Privacy for Knowledge Tracing,” in *Proceedings of the 18th International Conference on Educational Data Mining (EDM 2025)*, Palermo, Italy, 2025. (Poster paper.) **Nominee for Best Poster Paper Award.**
51. W. You and D. Lowd, “The Ultimate Cookbook for Invisible Poison: Crafting Subtle Clean-Label Text Backdoors with Style Attributes,” in *The Third Workshop on New Frontiers in Adversarial Machine Learning (AdvML-Frontiers) at NeurIPS 2024*, 2024.
52. J. Heibel and D. Lowd, “MaPPing Your Model: Assessing the Impact of Adversarial Attacks on LLM-based Programming Assistants,” in *ICML 2024 Workshop on Trustworthy Multi-modal Foundation Models and AI Agents (TiFA)*, 2024.
53. W. You, Z. Hammoudeh, and D. Lowd, “Large Language Models Are Better Adversaries: Exploring Generative Clean-Label Backdoor Attacks Against Text Classifiers,” in *The Second Workshop on New Frontiers in Adversarial Machine Learning*, 2023.
54. Z. Hammoudeh and D. Lowd, “Feature Partition Aggregation: A Fast Certified Defense Against a Union of ℓ_0 Attacks,” in *The Second Workshop on New Frontiers in Adversarial Machine Learning*, 2023.
55. W. You and D. Lowd, “Towards Stronger Adversarial Baselines Through Human-AI Collaboration,” in *Proceedings of NLP Power! The First Workshop on Efficient Benchmarking in NLP*, (pp.11–20), 2022.
56. Z. Xie, J. Brophy, A. Noack, W. You, K. Asthana, C. Perkins, S. Reis, Z. Hammoudeh, D. Lowd, and S. Singh, “What Models Know About Their Attackers: Deriving Attacker Information From Latent Representations,” in *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, (pp. 69–78), 2021.
57. Z. Hammoudeh and D. Lowd, “Simple, Attack-Agnostic Defense Against Targeted Training Set Attacks Using Cosine Similarity,” in *ICML Workshop on Uncertainty and Robustness in Deep Learning* (6 pages plus appendix), 2021.
58. J. Brophy and D. Lowd, “EGGS: A Flexible Approach to Relational Modeling of Social Network Spam,” in the *Ninth International Workshop on Statistical Relational AI (StarAI 2020)* (10 pages), New York, NY, 2020.
59. J. Brophy and D. Lowd, “Collective Classification of Social Network Spam,” in the *AAAI-17 Workshop on Artificial Intelligence for Cyber Security (AICS2017)* (pp. 162–169), San Francisco, CA, 2017.
60. I. Burago and D. Lowd, “Automated Attacks on Compression-Based Classifiers,” in *Proceedings of the 2015 ACM Workshop on Artificial Intelligence and Security (AISec)* (pp. 69–80), Denver, CO, 2015.
61. S. Jiang, D. Lowd, and D. Dou, “A Probabilistic Approach to Knowledge Translation,” in *Fifth International Workshop on Statistical Relational AI (StarAI 2015)* (7 pages), 2015.
62. S. Jiang, D. Lowd, and D. Dou, “Ontology Matching with Knowledge Rules,” in *Fifth International Workshop on Statistical Relational AI (StarAI 2015)* (7 pages), 2015.

63. M. Torkamani and D. Lowd, “Applying Dropout Regularization to Support Vector Machines,” in *NIPS 2014 Workshop on Perturbation, Optimization, and Statistics*, Montreal, Quebec, Canada (5 pages), 2014.
64. D. Lowd, B. Lessley, and M. De Raj, “Towards Adversarial Reasoning in Statistical Relational Domains,” in *AAAI-14 Workshop on Statistical Relational AI (StarAI 2014)* (pp. 57–59), 2014.
65. M. Torkamani and D. Lowd, “On Robustness and Regularization of Structural Support Vector Machines,” in *NIPS 2013 Workshop on Perturbation, Optimization, and Statistics* (5 pages), Stateline, NV, USA, 2013.
66. A. Rooshenas and D. Lowd, “Learning Sum-Product Networks with Direct and Indirect Interactions,” in *NIPS 2013 Workshop on Deep Learning* (7 pages), Stateline, NV, USA, 2013.
67. D. Stevens and D. Lowd, “On the Hardness of Evading Combinations of Linear Classifiers,” in *Proceedings of the 2013 ACM Workshop on Artificial Intelligence and Security (AISec)* (pp. 77–86), Berlin, Germany, 2013.
68. A. Rooshenas and D. Lowd, “Learning Tractable Graphical Models Using Mixtures of Arithmetic Circuits,” in *Late-Breaking Developments in the Field of Artificial Intelligence: Papers Presented at the Twenty-Seventh AAAI Conference on Artificial Intelligence* (pp. 104–106), Bellevue, WA, USA, 2013.
69. D. Lowd and A. Rooshenas, “Learning Markov Networks with Arithmetic Circuits,” in *ICML-13 Workshop on Infering* (6 pages), Atlanta, GA, USA, 2013.
70. M. Torkamani and D. Lowd, “Convex Adversarial Collective Classification,” in *UAI-12 Workshop on Statistical Relational AI* (8 pages), Catalina Island, CA, 2012.
71. S. Jiang, D. Lowd, and D. Dou, “Using Markov Logic to Refine an Automatically Extracted Knowledge Base,” in *UAI-12 Workshop on Statistical Relational AI* (8 pages), Catalina Island, CA, 2012.
72. S. Natarajan, T. Khot, D. Lowd, P. Tadepalli, K. Kersting, and J. Shavlik, “Exploiting Causal Independence in Markov Logic Networks: Combining Undirected and Directed Models,” in *AAAI-10 Workshop on Statistical Relational AI* (6 pages), Atlanta, Georgia, 2010.
73. D. Lowd, C. Meek and P. Domingos, “Foundations of Adversarial Machine Learning,” in *NIPS-2007 Workshop on Machine Learning in Adversarial Environments for Computer Security* (2 pages), Vancouver, Canada, 2007.
74. D. Lowd and P. Domingos, “Recursive Random Fields,” in *Proceedings of the ICML-2006 Workshop on Open Problems in Statistical Relational Learning* (8 pages), Pittsburgh, PA, 2006.

INVITED PAPERS

75. P. Domingos, D. Lowd, S. Kok, A. Nath, H. Poon, and P. Singla, “Unifying Logical and Statistical AI,” in *ACM/IEEE Symposium on Logic in Computer Science* (pp. 1–11), New York City, NY, 2016.

GENERAL AUDIENCE ARTICLES

76. D. Lowd, “Can Facebook use AI to fight online abuse?” in *The Conversation*, June 12, 2018. Reprinted by Scientific American, Salon, the Seattle P.I., and many other venues. Available online at:
<https://theconversation.com/can-facebook-use-ai-to-fight-online-abuse-95203>.

PRESS COVERAGE

- Quoted in news article about AI and mental health: “Will AI replace therapists?” in *Newsweek*, January 8, 2025.
Available online at:
<https://www.newsweek.com/will-ai-replace-therapists-chatgpt-mental-health-usa-2010934>
- Live broadcast appearance on navigating AI in 2025, *Global News*, January 4, 2025.
Available online at:
<https://www.youtube.com/watch?v=LP0xNwilLeY>
- Quoted in news article about AI chatbot safety for children: “Are AI Chatbots Safe for Children? Experts Weigh in After Teen’s Suicide” in *Newsweek*, November 13, 2024.
Available online at:
<https://www.newsweek.com/are-ai-chatbots-safe-children-experts-weigh-after-teens-suicide-8/8/28134>
- Work covered and public talk announced: “Science talk takes on ‘fake news’” in *The Nugget Newspaper*, April 16, 2019.
Available online at:
<https://nuggetnews.com/MobileContent/Education/Education/Article/Science-talk-takes-on-fake-news-8/8/28134>
- Quoted in news article about filtering violence online: “Why AI is still terrible at spotting violence online” in *CNN Business*, March 18, 2019.
Available online at:
<https://www.cnn.com/2019/03/16/tech/ai-video-spotting-terror-violence-new-zealand/index.html>
- Quoted in news article about OpenAI: “Nonprofit OpenAI looks at the bill to craft a Holy Grail AGI, gulps, spawns commercial arm to bag investors’ mega-bucks” in *The Register*, March 13, 2019.
Available online at:
https://www.theregister.co.uk/2019/03/13/openai_nonprofit_status/.
- Work covered in news article: “Potato, potato. Toma6to, I’m going to kill you... How a typo can turn an AI translator against us” in *The Register*, June 28, 2018.
Available online at:
https://www.theregister.co.uk/2018/06/28/machine_translation_vulnerable/.
- Quoted in news article about adversarial artificial intelligence: “The tiny changes that can cause AI to fail” in *BBC Future Now*, April 11, 2017.
Available online at:
https://www.bbc.com/future/article/20170410-how-to-fool-artificial-intelligence?ocid=global_future_rss

SOFTWARE RELEASED

Libra: Machine learning toolkit for Bayesian networks, Markov networks, and arithmetic circuits
<http://libra.cs.uoregon.edu>

(Published in the Journal of Machine Learning Research Machine Learning Open-Source Software (JMLR MLOSS).)

Alchemy: Algorithms for statistical relational AI

<http://alchemy.cs.washington.edu>

(Along with various other contributors.)

NBE: A Bayesian learner with very fast inference

<http://www.cs.washington.edu/ai/nbe>

CoFE: COllaborative Filtering Engine

<http://eecs.oregonstate.edu/iis/CoFE/>

(Along with various other contributors. No longer supported.)

INVITED TALKS

“The Manchurian Classifier: Invisible Triggers and Vulnerabilities in Text Models”

Oregon State University AI Seminar, Corvallis, OR, March, 2025.

“Training Data Influence Analysis and Estimation”

Machine Learning for Simulation Science (MLS), Institute for Artificial Intelligence, University of Stuttgart, Germany (remote), February, 2023.

“When Can We Trust Artificial Intelligence?”

Sisters Middle School, Sisters, OR, April, 2019.

“Algorithms and Artificial Intelligence: Science takes on fake news”

Sisters Science Club, Sisters, OR, April, 2019.

“When Can We Trust Artificial Intelligence?”

University of Oregon, Machine Learning Meetup, February, 2019.

“When Can We Trust Artificial Intelligence?”

QuackChat, Eugene, OR, January, 2019.

“A Whirlwind Tour of Machine Learning”

University of Oregon, ITS/HET Seminar, October, 2017.

“Adversarial Statistical Relational AI”

Sixth International Workshop on Statistical Relational AI (StarAI 2016), July, 2016.

“Adversarial Machine Learning in Relational Domains”

University of Texas Dallas (UTDallas), March, 2016.

“Adversarial Machine Learning in Relational Domains”

University of Utah, March, 2016.

“Adversarial Machine Learning in Relational Domains”

University of Maryland Baltimore County (UMBC), March, 2016.

“Adversarial Machine Learning in Relational Domains”

Tufts University, February, 2016.

“Adversarial Machine Learning in Relational Domains”

Michigan State University, February, 2016.

“Structure Learning for Sum-Product Networks”

Oregon State University, October, 2015.

“Structure Learning for Sum-Product Networks”

University of California Santa Cruz, November, 2014.

“Adversarial Machine Learning in Relational Domains”
 LinkedIn, November, 2014.

“Adversarial Machine Learning in Relational Domains”
 United Technologies Research Center Berkeley, November, 2014.

“Structure Learning for Sum-Product Networks”
 United Technologies Research Center New Haven, November, 2014.

“Adversarial Machine Learning in Relational Domains”
 Amazon.com, Inc., October, 2014.

“Machine Learning with Evasive Adversaries”
 University of Washington. September, 2014.

“Learning Tractable Probabilistic Models” (with Pedro Domingos)
 Tutorial for the 2014 Conference on Uncertainty in Artificial Intelligence (UAI 2014). July, 2014.

“Using Dependency Networks To Learn Markov Networks”
 University of Wisconsin-Madison. October, 2013.

“Learning Relational Classifiers for Adversarial Domains”
 University of Indiana-Bloomington. October, 2013.

“Learning Relational Classifiers for Adversarial Domains”
 SRI International. October, 2013.

“Better Learning and Inference with Dependency Networks”
 University of Maryland, College Park. January, 2013.

“Learning Relational Classifiers for Adversarial Domains”
 LinkedIn. January, 2013.

“Convex Adversarial Collective Classification”
 SRI International. January, 2013.

“Convex Adversarial Collective Classification”
 Dagstuhl Perspectives Workshop on Machine Learning Methods for Computer Security. October, 2012.

“Convex Adversarial Collective Classification”
 University of Washington. September, 2012.

“Toward Adversarial Collective Classification”
 2nd ARO Workshop on Moving Target Defense, George Mason University. October, 2011.

“Better Learning and Inference with Dependency Networks”
 University of Washington. September, 2011.

“Mean Field Inference in Dependency Networks: An Empirical Study”
 Oregon State University. July, 2011.

“Inference Complexity As Learning Bias”
 Oregon State University. January, 2010.

“Adversarial Machine Learning”
 Portland State University. November, 2009.

“Markov Logic: Representation, Inference and Learning”
University of Michigan. March, 2009.

“Markov Logic: Representation, Inference and Learning”
University of Oregon. April, 2009.

“Foundations of Adversarial Machine Learning”
University of Cagliari. July, 2008.

“Adversarial Learning”
Oregon State University. June, 2006.

GRANTS AND OTHER FUNDING

6/2023 – 5/2028: *INVITE: Inclusive and Innovative Intelligent Technologies for Education*, National Science Foundation (NSF), \$1,109,778 (Oregon’s share of a collaborative grant led by UIUC); total budget: \$19,562,906.

10/2020 – 12/2024: Reverse Engineering Adversarially Constructed Text (REACT), Defense Advanced Research Projects Agency (DARPA), \$1,555,853 (PI with UO as lead institution and UC Irvine as subcontractor).

2/2020 – 2/2021: *Relating learning and intelligence across artificial and human agents*, 2019 Data Science Initiative Seed Funding convening award, University of Oregon, \$10,000 (co-PI with Ben Hutchinson).

6/2016 – 6/2020: *AMIA: Automated Media Integrity Assessment*, Defense Advanced Research Projects Agency, Media Forensics (MediFor) Program (DARPA-BAA-15-58), Technical Area TA2, \$807,000 (subcontract from SRI International; total budget: \$9,490,000).

6/2015 – 6/2018: *Inferring Trustworthiness and Deceit with Adversarial Relational Models*, Army Research Office, Young Investigator Program, \$357,573.

9/2014 – 8/2016: *EAGER: Machine Learning to Combat Adversarial Attacks*, National Science Foundation, \$104,997.

9/2013 – 8/2014: *Learning Tractable Graphical Models with Latent Variables*, Google Faculty Research Award, \$56,856.

5/2013 – 2/2016: *Understanding the Mechanism of Social Network Influence in Health Outcomes through Multidimensional and Semantic Data Mining Approaches*, National Institutes of Health, \$1,500,000 (co-PI, with Dejing Dou and 6 others).

7/2011 – 6/2015: *Statistical Knowledge Translation and Knowledge Integration Using Markov Logic*, National Science Foundation, \$495,000 (co-PI, with Dejing Dou).

1/2011 – 6/2014: *A Unified Approach to Abductive Inference*, Multidisciplinary University Research Initiative, Army Research Office, \$232,998 (subcontract from UW).

COURSES TAUGHT

CS 670: Data Science. Spring 2021, Spring 2022, Winter 2023, Spring 2024, Winter 2025.

CS 610pm: Probabilistic Methods for AI. Spring 2012.

CS 610: Welcome to Grad School. Fall 2021-2022, 2024. (2-credit seminar.)

CS 610: Teaching Effectiveness / Welcome to Grad School. Fall 2023. (2-credit seminar, co-taught.)
 CS 473/573: Probabilistic Methods for AI. Winter 2010, Fall 2010, Spring 2013—2018, Spring 2025.
 CS 472/572: Machine Learning. Winter 2013, Winter 2015—2018, Winter 2020.
 CS 471/571: Introduction to Artificial Intelligence. Fall 2012, Fall 2016.
 CS 413/513: Advanced Data Structures. Spring 2010, Fall 2011.
 CS 399: Probability and Statistics for Computer Science. Fall 2017.
 CS 313: Intermediate Data Structures. Winter 2014, Fall 2015, Fall 2019–2024.
 CS 211: Computer Science II. Winter 2011, Spring 2011, Winter 2012.

STUDENTS

Ph.D. advisor (graduated)

Wencong You (Amazon)

Dissertation title: *The Dual Game of Clean-Label Backdoor Attacks and Defenses in NLP*.
 Summer 2025.

Zayd Hammoudeh

Dissertation title: *Certified and Forensic Defenses Against Poisoning and Backdoor Attacks*.
 Fall 2023.

Jonathan Brophy

Dissertation title: *Understanding and Adapting Tree Ensembles: A Training Data Perspective*.
 Fall 2022.

Javid Ebrahimi (Visa) (co-advised with Dejing Dou)

Dissertation title: *Robustness of Neural Networks for Discrete Input: An Adversarial Perspective*.
 Fall 2018.

Amirmohammad Rooshenas (University of Massachusetts, Amherst)

Dissertation title: *Learning Tractable Graphical Models*.
 Winter 2017.

MohamadAli Torkamani (Cambia)

Dissertation title: *Robust Large Margin Approaches for Machine Learning in Adversarial Settings*.
 Summer 2016.

Shangpu Jiang (Google) (co-advised with Dejing Dou)

Dissertation title: *Knowledge Base Refinement and Knowledge Translation with Markov Logic Networks*.
 Fall 2015.

M.S. thesis advisor (graduated)

Bharath Kumar Nachenahalli Bhuthegowda

Thesis title: *Methods for Analyzing the Evolution of Email Spam*. Summer 2018.

Jonathan Brophy (University of Oregon)

Thesis title: *Collective Classification of Social Network Spam*. Spring 2017.

Igor Burago (University of California Irvine)

Thesis title: *Automated Attacks on Compression-Based Classifiers*. Spring 2014.

Arash Shamaei (Oregon State University)

Thesis title: *Fast Inference Algorithms in Dependency Networks*. Summer 2011.

B.S. thesis advisor (graduated)

Sam Nelson

Thesis title: *Predicting SoundCloud Spammers*. Spring 2016.

PROFESSIONAL SERVICE

Editorial Board:

Journal of Machine Learning Research (JMLR), Editorial Board, 2020—Present
Springer Data Mining and Knowledge Discovery (DMKD), Editorial Board, 2019—2021
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI),
Associate Editor, 2018—2020
ECML PKDD 2018 Journal Track, Guest Editorial Board, 2018

Associate Program Chair:

AAAI Conference on Artificial Intelligence (AAAI) 2023

Workshop Program Co-Chair:

AAAI Conference on Artificial Intelligence (AAAI) 2015, 2016

Proceedings Chair:

Conference on Uncertainty in Artificial Intelligence (UAI) 2014, 2015

Workshop Co-Chair:

ICML 2019 3rd Workshop on Tractable Probabilistic Modeling
IJCAI-ECAI-ICML 2018 Workshop on Tractable Probabilistic Models
New Perspectives for Relational Learning, Banff International Research Station, April 2015
ICML 2014 Workshop on Learning, Security and Privacy
ICML 2014 Workshop on Learning Tractable Probabilistic Models

Program Committees:

International Conference on Learning Representations (ICLR) 2018, 2019;
International Joint Conference on Artificial Intelligence (IJCAI) 2011, 2013, 2017;
2015–2016, 2018–2019, 2022 (Senior Program Committee)
AAAI Conference on Artificial Intelligence (AAAI) 2006, 2010–2015;
2017, 2019–2020 (Senior Program Committee), 2024
Neural Information Processing Systems (NeurIPS) conference, 2006, 2012–2018 (Reviewer),
2020 (Area Chair), 2022–2023 (Area Chair, Datasets and Benchmarks track), 2024 (Area
Chair)
Uncertainty in Artificial Intelligence (UAI) 2010–2018
International Conference on Machine Learning (ICML) 2010–2014,
2016 (**Outstanding Reviewer Award**), 2017, 2019 (Area Chair), 2020 (Expert Re-
viewer)
International Conference on Artificial Intelligence and Statistics (AISTATS) 2012,
2013, 2017, 2019
Conference on Decision and Game Theory for Security (GameSec) 2019
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2011,
2015–2016
ECML-PKDD 2012

Asian Conference on Machine Learning (ACML) 2014
 ACM Workshop on Artificial Intelligence and Security (AISec) 2013–2016, 2018, 2019, 2022
 AAAI-16 Workshop on Artificial Intelligence for Cyber Security (AICS2016)
 Workshop on Statistical Relational AI (StarAI) 2013—2018
 ICML Workshop on Structured Learning (SLG) 2013
 ICML Workshop on Statistical Relational Learning 2012
 AAAI Nectar 2011
 ECML/PKDD Workshop on Privacy and Security issues in Data Mining
 and Machine Learning, 2010

Journal Reviewer:

Journal of Machine Learning Research (JMLR)
 IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
 Machine Learning Journal (MLJ)
 Journal of Artificial Intelligence Research (JAIR)
 Artificial Intelligence Journal (AIJ)
 Data Mining and Knowledge Discovery (DMKD) journal
 Transactions on Dependable and Secure Computing (TDSC)
 IEEE Security & Privacy
 ACM SIGKDD Explorations
 Journal of Systems and Software
 International Journal of Pattern Recognition and Artificial Intelligence

Conference Reviewer:

International Conference on Machine Learning (ICML) 2008
 ACM Conference on Electronic Commerce (EC) 2006

Other:

AI Curriculum Task Force, 2026
 CAS Launch Committee Convener, 2024–2025
 DEI Committee Chair, 2022–2025
 UO School of Computer and Data Science (SCDS) Strategic Implementation Team, 2023
 UO School of Computer and Data Science (SCDS) Planning Committee, 2022
 Hiring committee: University of Oregon CIO, 2022
 Hiring committee, 2021–2022 biomedical data science search
 Ripple Fellowship Committee (2022)
 Director of Graduate Studies, Dept. of Computer Science (2019–2021)
 Member of Oregon’s Artificial Intelligence Taskforce Education Sub-Committee, 2020