Patrick Hall

Curriculum vitae

Professional Experience

2023–Present **Teaching Faculty**, *The George Washington School of Business*, Washington, DC

- Conducting research in support of the NIST AI Risk Management Framework.
- Teaching graduate and undergraduate classes relating to business analytics, data ethics and responsible AI.

2023–Present **Principal Scientist**, *HallResearch.ai*, Washington, DC

- Founded boutique AI consultancy with focus on responsible AI advice and implementation.
- Advising on traditional model governance and generative AI risk management.
- Auditing high-impact machine learning systems in Consumer Finance, Big Tech and Fin Tech.
- Training corporate personnel and AI community members on AI risk management, explainable AI, model validation, and model governance.

2020–2023 **Visiting Faculty**, *The George Washington School of Business*, Washington, DC

- Teaching and administering MSBA practicum (DNSC 6317) and undergraduate capstone (DNSC 4289) classes and consulting projects.
- Teaching graduate and undergraduate classes, including Analytics Edge and Data Ethics (DNSC 3288W), Analytics Edge and Data Ethics Bootcamp (DNSC 6301), Machine Learning 1 (DNSC 6314), Machine Learning 2 (DNSC 6315) and Responsible Machine Learning (DNSC 6290).
- Advising on curricula in the Department of Decision Sciences.

2019–2023 **Principal Scientist**, *BNH.AI*, Washington, DC

- Co-founded a boutique law firm focused on minimizing the legal and technical risks of artificial intelligence (AI).
- Model audit and certification.
- AI system red teaming and incident response.
- Development of bespoke software offerings for bias testing and data privacy risk controls.
- Assisting with legal guidance related to the use of explainable machine learning (ML), interpretable ML models, generation of adverse action notices, algorithmic discrimination, data privacy, ML security, and advanced model governance.

2015–2020 **Adjunct Faculty**, *The George Washington University*, Washington, DC

- Creating lectures, assignments, and assessments for graduate data mining, ML, and SAS language courses for the Department of Decision Sciences.
- Received the *Outstanding Faculty Award* for the Master of Science in Business Analytics program in 2019.

2017–2020 **Senior Director for Data Science Products**, *H2O.ai*, Mountain View, CA

- Product management and design for interpretable models, ML explanation, model debugging, and discrimination testing in the award-winning Driverless AI AutoML product, resulting in $50M+ revenue and venture funding.

- Research lead for interpretable, secure, and fair ML, resulting in multiple invited and accepted publications and presentations (see below).
- Technical sales lead for responsible ML in US financial services.

2016–2017 **Senior Product Engineer and Data Scientist**, *H2O.ai*, Mountain View, CA

- Implemented prototypes for increased trust and understanding, and reduced discrimination, in black-box AutoML decision-making systems, later incorporated into H2O Driverless AI.
- Advised data science groups across the US financial services industry.
- Co-organized Washington DC AI and Deep Learning meetup group.

2014–2016 **Senior Machine Learning Scientist**, *SAS Institute*, Arlington, VA

- Influenced ML research direction for SAS Enterprise Miner, a leading commercial data mining and ML solution with $1B+ lifetime revenue.
- Advised Fortune 500 customers on real-world, large-scale ML problems; involved in numerous successful $1M+ customer engagements.
- Enabled technical sales teams and user communities by creating white papers, blogs, and YouTube videos, and developing and maintaining open source, public GitHub repositories to share advanced ML use cases.
- Collaborated with colleagues to design an automated system to clean and cluster a dataset for market segmentation purposes (US Patent Numbers 9202178, 9367602, 9367799, 9471869, and 9489621).

2014 **Research Statistician Developer**, *SAS Institute*, Cary, NC

- Became the 11th person worldwide to be certified as a Cloudera data scientist.

2012–2014 **Senior Associate Research Statistician Developer**, *SAS Institute*, Cary, NC

- Co-invented the aligned box criterion (ABC) for estimating the number of clusters in a dataset (US Patent Number 9424337). This method was later implemented in two commercial distributed, in-memory data mining and ML solutions.
- Developed SAS Enterprise Miner R Integration, enabling R models to be deployed in-database, in-Hadoop, and through web services and to be monitored using industry standard model management tools.
- Teamed up with colleagues to introduce deep learning research into SAS Enterprise Miner R & D; awarded *Coolest Idea* prize with two colleagues in SAS Advanced Analytics Science Fair for *Deep Learning in Enterprise Miner* project.

2011–2012 **Practicum Project Team Leader**, *M&T Bank*, Buffalo, NY

- Directed strategy, managed day-to-day tasks, and handled communications with corporate sponsors for a team of five master's degree students; provided technical mentorship.
- Constructed risk-adjusted response models for credit risk analytics applications using SAS Enterprise Miner.

2011 **Software Engineer**, *Larson Technologies Inc.*, Carrboro, NC

- Developed traditional and cloud-based Java applications for physical security hardware.
- Interfaced Java and C++ software libraries with smart card technologies.
- Administered basic Ubuntu Linux Server/Apache Tomcat/MySQL database systems.

2010 **Linux Software Intern**, *Tetrark Inc.*, Chapel Hill, NC

- Enhanced custom CUDA-based software on an Ubuntu Linux platform for a POC integrating Tetrark's GPU libraries with a future customer's solutions; Tetrark later became part of the graphics acceleration for Google Chrome.

2003–2007 **Technical Director**, *World Camp Inc.*, Lilongwe, Malawi

- Processed and analyzed a two-year backlog of data for World Camp, leading to future grant funding and automated statistical protocols to evaluate teacher performance.

- Managed daily operations of rural education programs, including procuring food and supplies for hundreds of students, arranging safe transportation for over 30 World Camp personnel, and overseeing 10 teaching teams.
- Fostered relationships with study abroad programs at Duke University and Middlebury College; recruited volunteer teachers from 7 East Coast colleges.

Education

2011–2012 **Master of Science in Analytics (MSA)**, *North Carolina State University*
- Director's Prize, 4.0 GPA.
- Research: *Let the Data Paint the Picture: Data-Driven, Interactive and Animated Visualizations Using SAS, Java and the Processing Graphics Library*.

2010–2011 **Master's studies, Computer Science**, *North Carolina State University*
- Classwork: object-oriented and procedural programming, databases.

2005–2006 **PhD studies, Computational Chemistry**, *University of Illinois*
- Roger Adams Fellow.
- Research: *Isoprenoid Biosynthesis as a Drug Target: Bisphosphonate Inhibition of Escherichia coli K12 Growth and Synergistic Effects of Fosmidomycin*.
- Classwork: quantum and statistical mechanics.

1999-2004 **Bachelor of Arts in Mathematics (BA)**, *University of North Carolina*
- NSF REU Fellow.
- Research: *Undergraduate Introductory Quantitative Chemistry Laboratory Course: Interdisciplinary Group Projects in Phytoremediation*.

Patents

17/888,109 Testing and documenting algorithmic systems (application).
11922283B2 Model interpretation.
11386342B2 Model interpretation.
9495414 Cluster computation using random subsets of variables.
9489621 Graph based selection of decorrelated variables.
9471869 Determination of composite clusters.
9424337 Number of clusters estimation.
9367799 Neural network based cluster visualization that computes pairwise distances between centroid locations, and determines a projected centroid location in a multidimensional space.
9367602 Probabilistic cluster assignment.
9202178 Computerized cluster analysis framework for decorrelated cluster identification in datasets.

Selected Publications

NIST (2024). "AI 600-1, Generative AI Risk Profile". In: *nist.gov*. `https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf`.

Schwartz, Reva et al. (2024a). "The Assessing Risks and Impacts of AI (ARIA) Program Evaluation Design Document". In: *nist.gov*. `https://ai-challenges.nist.gov/aria/docs/ARIA_Program_Companion_Document_Dec20.pdf`.

Schwartz, Reva et al. (2024b). "The NIST Assessing Risks and Impacts of AI (ARIA) Pilot Evaluation Plan". In: *nist.gov*. `https://ai-challenges.nist.gov/aria/docs/evaluation_plan.pdf`.

Atherton, Daniel et al. (2023). *NIST AI 100-3: The Language of Trustworthy AI: An In-Depth Glossary of Terms*. URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-3.pdf.

Hall, Patrick, James Curtis, and Parul Pandey (2023). *Machine Learning for High-risk Applications*. URL: https://www.oreilly.com/library/view/machine-learning-for/9781098102425/. Sebastopol, CA: O'Reilly.

NIST (2023). *NIST AI 100-1: AI Risk Management Framework*. URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

Brennen, Andrea et al. (2022). *AI Assurance Audit of RoBERTa, an Open source, Pretrained Large Language Model*. URL: https://assets.iqt.org/pdfs/IQTLabs_RoBERTaAudit_Dec2022_final.pdf/web/viewer.html.

Hall, Patrick (2022a). *The Data Scientific Method vs. The Scientific Method*. URL: https://opendatascience.com/the-data-scientific-method-vs-the-scientific-method/.

– (2022b). *What We Learned Auditing Sophisticated AI for Bias*. URL: https://www.oreilly.com/radar/what-we-learned-auditing-sophisticated-ai-for-bias/.

Schwartz, Reva et al. (2022). *NIST SP1270: Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*. URL: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf.

Buehler, Kevin et al. (2021). *Getting to know—and manage—your biggest AI risks*. URL: https://www.mckinsey.com/capabilities/quantumblack/our-insights/getting-to-know-and-manage-your-biggest-ai-risks.

Hall, Patrick and Ayoub Ouederni (2021). *Seven Legal Questions for Data Scientists*. URL: https://www.oreilly.com/radar/seven-legal-questions-for-data-scientists/.

Hall, Patrick et al. (2021). "A United States Fair Lending Perspective on Machine Learning". In: *Frontiers in Artificial Intelligence* 4. URL: https://www.frontiersin.org/articles/10.3389/frai.2021.695301/full.

Leong, Brenda and Patrick Hall (2021). *Five Things Lawyers Should Know About Artificial Intelligence*. URL: https://www.abajournal.com/columns/article/5-things-lawyers-should-know-about-artificial-intelligence/.

Burt, Andrew and Patrick Hall (2020). *What to Do When AI Fails*. URL: https://www.oreilly.com/radar/what-to-do-when-ai-fails/.

Darnell, Dan, Rafael Coss, and Patrick Hall (2020). *The Future of Analytics: The New Landscape of Artificial Intelligence and Machine Learning Applications*. URL: http://info.h2o.ai/rs/644-PKX-778/images/OReilly-The-Future-of-Analytics-eBook.pdf.

Hall, Patrick (2020a). *Explaining Machine Learning Models to the Business*. URL: https://www.infoworld.com/article/3533369/explaining-machine-learning-models-to-the-business.html.

Hall, Patrick and Andrew Burt (2020a). "Build Your AI Incident Response Plan... Before It's Too Late". In: *Cyber Defense eMagazine* (August). URL: https://www.cyberdefensemagazine.com/build-your-ai-incident-response-plan/.

– (2020b). *Sample AI Incident Response Checklist*. URL: https://bnh-ai.github.io/resources/.

Hall, Patrick, Navdeep Gill, and Benjamin Cox (2020). *Responsible Machine Learning: Actionable Strategies for Mitigating Risks and Driving Adoption*. URL: http://info.h2o.ai/rs/644-PKX-778/images/OReilly_Responsible_ML_eBook.pdf.

Hall, Patrick and Michael Proksch (2020). *From GLM to GBM – Part 2: How an Economics Nobel Prize could revolutionize insurance and lending*. URL: https://www.h2o.ai/blog/from-glm-to-gbm-part-2/.

Hall, Patrick and Collin Starkweather (2020). "AI Incidents in Financial Services – Parts 1 and 2". In: *Thompson-Reuters Regulatory Intelligence*.

Hall, Patrick et al. (2020). "A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing". In: *Information* 11 (3). URL: https://www.mdpi.com/2078-2489/11/3/137.

Hall, Patrick et al. (2020b). *Machine Learning: Considerations for fairly and transparently expanding access to credit*. URL: https://bit.ly/33cIJf5.

Mugdal, Yogesh et al. (2020). *Artificial Intelligence Risk & Governance*. URL: https://ai.wharton.upenn.edu/artificial-intelligence-risk-governance/.

Proksch, Michael and Patrick Hall (2020). *From GLM to GBM – Part 1: How an Economics Nobel Prize could revolutionize insurance and lending*. URL: https://www.h2o.ai/blog/from-glm-to-gbm-part-1/.

Hall, Patrick (2019a). "On the Art and Science of Machine Learning Explanations". In: *KDD XAI Workshop*. URL: https://arxiv.org/pdf/1810.02909.pdf.

– (2019b). *Proposals for Model Vulnerability and Security*. URL: https://www.oreilly.com/ideas/proposals-for-model-vulnerability-and-security.

– (2019c). *Real-world Strategies for Model Debugging*. URL: https://towardsdatascience.com/strategies-for-model-debugging-aa822f1097ce.

Hall, Patrick and Andrew Burt (2019). *Why You Should Care About Debugging Machine Learning Models*. URL: https://www.oreilly.com/radar/why-you-should-care-about-debugging-machine-learning-models/.

Hall, Patrick and Navdeep Gill (2019). *An Introduction to Machine Learning Interpretability: An Applied Perspective on Fairness, Accountability, Transparency, and Explainable AI, 2nd Edition*. URL: https://www.h2o.ai/wp-content/uploads/2019/08/An-Introduction-to-Machine-Learning-Interpretability-Second-Edition.pdf.

Hall, Patrick, Navdeep Gill, and Nicholas Schmidt (2019). "Proposed Guidelines for the Responsible Use of Explainable Machine Learning". In: *NeurIPS Robust AI in Financial Services Workshop*. URL: https://arxiv.org/pdf/1906.03533.pdf.

Stalla-Bourdillon, Sophie et al. (2019). *Warning Signs: Privacy and Security in an Age of Machine Learning*. URL: https://fpf.org/wp-content/uploads/2019/09/FPF_WarningSigns_Report.pdf.

Hall, Patrick and Navdeep Gill (2018). *An Introduction to Machine Learning Interpretability: An Applied Perspective on Fairness, Accountability, Transparency, and Explainable AI*. URL: https://pages.dataiku.com/hubfs/ML-interperatability.pdf.

Hall, Patrick, Navdeep Gill, and Lingyao Meng (2018). *Testing Machine Learning Explanation Techniques*. URL: https://www.oreilly.com/ideas/testing-machine-learning-interpretability-techniques.

Hall, Patrick and Navdeep Gill (2017). "Debugging the Black-Box COMPAS Risk Assessment Instrument to Diagnose and Remediate Bias". In: *Open Review Preprint*. URL: https://openreview.net/pdf?id=r1iWHVJ7Z.

Hall, Patrick, Megan Kurka, and Angela Bartz (2017). *Using H2O Driverless AI*. URL: http://docs.h2o.ai/driverless-ai/latest-lts/docs/booklets/DriverlessAIBooklet.pdf.

Hall, Patrick, Wen Phan, and Sri Satish Ambati (2017). *Ideas on Interpreting Machine Learning*. URL: https://www.oreilly.com/ideas/ideas-on-interpreting-machine-learning.

Hall, Patrick et al. (2017). *Machine Learning Interpretability with H2O Driverless AI*. URL: http://docs.h2o.ai/driverless-ai/latest-stable/docs/booklets/MLIBooklet.pdf.

Hall, Patrick (2016a). *Predictive Modeling: Striking a Balance Between Accuracy and Interpretability*. URL: https://www.oreilly.com/ideas/predictive-modeling-striking-a-balance-between-accuracy-and-interpretability.

– (2016b). *The Preoccupation with Test Error in Applied Machine Learning*. URL: https://www.oreilly.com/ideas/the-preoccupation-with-test-error-in-applied-machine-learning.

Hall, Patrick, Wen Phan, and Katie Whitson (2016). *The Evolution of Analytics*. URL: https://pdfs.semanticscholar.org/cc62/c04074334d1d39b1c9f6a47b1ada99858529.pdf.

Hall, Patrick et al. (2016). "An Efficient Pattern Recognition Approach with Applications". In: *Proceedings of SAS Global Forum*. URL: https://pdfs.semanticscholar.org/f3fd/f7815ffbef2e06d57c1b8c8705091bf3d12a.pdf.

Holdaway, Keith et al. (2016). "Surface Reservoir Characteristics from Subsurface Seismic Images with Deep Learning Methodologies". In: *International Conference and Exhibition, Barcelona, Spain, 3-6 April 2016*. URL: https://library.seg.org/doi/10.1190/ice2016-6319324.1, pp. 239–239.

Wujek, Brett, Patrick Hall, and Funda Günes (2016). "Best Practices for Machine Learning Applications". In: *Proceedings of SAS Global Forum*. URL: https://support.sas.com/resources/papers/proceedings16/SAS2360-2016.pdf.

Hall, Patrick et al. (2014). "An Overview of Machine Learning with SAS Enterprise Miner". In: *Proceedings of SAS Global Forum*. URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.648.3176&rep=rep1&type=pdf.

Snyder, Ryan and Patrick Hall (2012). "Let the Data Paint the Picture: Data-Driven, Interactive and Animated Visualizations Using SAS, Java and the Processing Graphics Library". In: *Proceedings of SAS Global Forum*. URL: http://support.sas.com/resources/papers/proceedings12/275-2012.pdf.

Van Engelen, Debra L. et al. (2007). "Undergraduate Introductory Quantitative Chemistry Laboratory Course: Interdisciplinary Group Projects in Phytoremediation". In: *Journal of Chemical Education* 84.1. URL: https://pubs.acs.org/doi/pdf/10.1021/ed084p128, p. 128.

Leon, Annette et al. (2006). "Isoprenoid Biosynthesis as a Drug Target: Bisphosphonate Inhibition of Escherichia coli K12 Growth and Synergistic Effects of Fosmidomycin". In: *Journal of Medicinal Chemistry* 49.25. URL: https://pubs.acs.org/doi/abs/10.1021/jm060492b, pp. 7331–7341.

Selected Conferences, Presentations and Podcasts

- *Poor Measurement in AI*, GMU Guest Lecture for Professor Missy Cummings
- *Overcoming Organizational Challenges When Adopting The NIST AI Risk Management Framework*, GWU TRAILSCon 2025
- *ACM FAccT Conference 2025* (Program Committee)
- *The AI Revolution in Consumer Finance* (Invited, Session chair)
  Philadelphia Federal Reserve, 2024
- *Dispelling magical thinking and the black box of AI with GWU's Patrick Hall*
  Humans of AI Podcast, 2024
- *De-risking Generative AI for Adoption*, Databricks AI Summit, 2023
- *Practice Tutorial: Using the NIST AI Risk Management Framework* (Tutorial)
  ACM FAccT Conference, 2023
- ACM FAccT Conference 2023 (Program Committee)
- *Machine Learning Explainability Master Class* (Invited)
  UNCC Analytics Frontiers Conference, 2023
- *Explainable Machine Learning Models* (Invited)
  NSCU Institute for Advanced Analytics Guest Lecture, 2023
- *Quality AI/ML Models by Design with PiML* (Invited)
  OpenTeams TechShares, 2023

- *Emerging ML/AI Topics in Consumer Credit* (Invited, Session chair)
  Philadelphia Federal Reserve, 2023
- *Seven Legal Questions for Data Scientists*
  GWU SEAS Trustworthy AI PhD Seminar, 2022
- *Workshop on the Ethical Design of AIs (EDAIs)*
  National Science Foundation Ideation Workshop, 2022
- *A Tutorial on Contemporary Machine Learning Risk Management* (Invited)
  NSCU Institute for Advanced Analytics Guest Lecture, 2022
- *AI and Hiring Tech Panel* (Invited)
  DEFCON, 2022
- *Interpretable and Explainable AI and Machine Learning* (Invited)
  National Academy of Science, 2022
- *An Overview of the U.S. National Institute of Standards and Technology (NIST) Socio-technical AI Bias Guidance*
  Frankfurt AI Risk Roundtable, 2022
- *AI RMF Function Govern* (Invited, Session chair)
  National Institute of Standards and Technology, 2022
- *Fairness, Governance, and the Future of Artificial Intelligence* (Invited)
  Elon Analytics Day: Analytics in Society, 2022
- *A Tutorial on Contemporary Machine Learning Risk Management* (Tutorial, Invited)
  ODSC, 2022
- *Does your AI system comply with the law?*
  ZDNet Podcast, 2021
- *AI Superstream Series: Responsible AI* (Invited)
  AI Superstream Series: Responsible AI, 2021
- *Real-world Strategies for Debugging Machine Learning Systems*
  DataBricks AI Summit, 2021
- *What to Do When AI Fails with Andrew Burt and Patrick Hall*
  O'Reilly Meet Experts, 2021
- *Bias in AI: Understand the impact on marketing and advertising* (Invited)
  Quantcast Webinar, 2021
- *What To Do When AI Fails: AI Incident Response*
  MLConf, 2021
- *AI Regulation in Finance Panel* (Invited)
  QuantUniversity, 2021
- *Responsible AI* (Invited)
  Georgetown University Master's in Business Analytics Guest Lecture, 2021
- *AI Incident Response: What to Do When AI Fails* (Invited)
  RSQRD Meetup, 2021
- *Real-world Strategies for Model Debugging* (Invited)
  Duke University Master's of Statistics Proseminer Series, 2021
- *What To Do When AI Fails: AI Incident Response* (Invited)
  QuantUniversity, 2021
- *What does a risk management strategy look like for explainable AI?* (Invited)
  National Institute of Standards and Technology, 2021
- *Algorithmic Decision Making: Exploring Practical Approaches to Liability, Fairness, and Explainability*
  Toronto Machine Learning Series, 2020
- *US Perspective on Avoiding, Detecting and Responding to the Liabilities of AI and Analytics* (Invited)
  Frankfurt AI Risk Roundtable, 2020
- *Managing AI Liabilities: A Tech-Literate Legal Perspective on Machine Learning*
  DC Data Science Meetup, 2020
- *What to Do When AI Fails*

Mind Hack, 2020
- *Mitigating the Liabilities of Artificial Intelligence in NLP*
  NLP Summit, 2020
- *Real-world Strategies for Model Debugging*
  Toronto Machine Learning Series, April 2020
- *Interpretable and Fair Machine Learning in Finance*
  ASA Symposium on Data Science and Statistics, 2020 (Session chair and organizer)
- *Toward Human-Centered Machine Learning*
  CrunchConf, 2019 (Invited)
- *Proposed Guidelines for the Responsible Use of Explainable Machine Learning*
  KDD, 2019 (Invited)
- *Increasing Trust and Interpretability in Machine Learning with Model Debugging*
  JSM, 2019 (Invited)
- *Practical Techniques for Interpretable Machine Learning* (Tutorial)
  O'Reilly Strata, 2019
- *A Discussion of Model Explanation Tools*
  NYC Deep Learning and Big Data Science Meetup, 2018
- *On the Art and Science of Machine Learning Explanations*
  JSM, 2018 (Invited)
- *Machine Learning Interpretability: The Good, the Bad, and the Ugly*
  NYC Deep Learning and Big Data Science Meetup, 2018
- *Practical Techniques for Interpreting Machine Learning Models: Introductory Open Source Examples Using Python, H2O, and XGBoost* (Tutorial)
  FATML, 2018
- *Gradient Boosting Machines with Sparkling Water*
  Spark Saturday, 2017
- *Interpretability in Conversation with Patrick Hall and Sameer Singh*
  Fast Forward Labs Webinar, 2017
- *Interpretable AI: Not Just For Regulators!*
  O'Reilly AI, 2017
- *A Survival Guide for Machine Learning: Top 10 Tips from a Battle-tested Solution*
  O'Reilly Strata + Hadoop World, 2016
- *Deep Learning and Real World Experience*
  Data Science Melbourne Meetup, 2015
- *Playing Nice: Using PMML, Python, R, and SAS for Production Analytics*
  DC Statistical Programming Meetup, 2015
- *An Overview of Machine Learning with SAS Enterprise Miner*
  Informs Big Data, 2014

Professional Certifications

- Cloudera Certified Data Scientist, CCP-DS: 11, 2014
- SAS Certified Advanced Programmer for SAS 9, 2011
- SAS Certified Predictive Modeler for SAS Enterprise Miner 6, 2011
- Brody Professional Communications, 2011

Selected Awards

- H2o.ai AI 100 List, 2024
- George Washington University School of Business, Outstanding Faculty Award, 2019, 2022, 2023
- NC State Institute for Advanced Analytics, Director's Prize, 2012
- University of Illinois, Roger Adams Fellowship, 2005
- National Science Foundation, REU Fellowship, 2001

Selected Volunteer Experience

2021–Present **Board of Directors**, *AI Incident Database*, Sacramento, CA
- Currently serving as President of the Board of Directors.

2023–Present **Board of Advisors**, *Underwriters Laboratories Digital Safety Institute*, Chicago, IL

2024–Present **Board of Advisors**, *Humane Intelligence*, San Francisco, CA

Selected Press Appearances

- *Elon Musk's AI chatbot, Grok, started calling itself 'MechaHitler'*, National Public Radio, 2025
- *Google's AI Overview has no opt-out. That's making some people unhappy*, National Public Radio, 2024
- *California Governor Vetoes Sweeping A.I. Legislation*, The New York Times, 2024
- *California AI bill veto opens door to bank-unfriendly revisions*, American Banker 2024
- *Google Rolls Back A.I. Search Feature After Flubs and Flaws*, The New York Times, 2024
- *How to detect AI-generated content*, TechTarget, 2023
- *New York has new rules on AI hiring tools. Here's how the changes might help or hinder your job prospects*, WNYC Gothamist, 2023
- *Why red team exercises for AI should be on a CISO's radar*, CISO Magazine, 2023
- *Plan to open-source Twitter algorithm raises transparency, revenue questions*, S&P Global Market Intelligence, 2022
- *Twitter whistleblower report complicates odds of legal victory over Musk*, S&P Global Market Intelligence, 2022
- *Twitter's Photo-Cropping Algorithm Favors Young, Thin Females*, Wired, 2021
- *Law firms are building A.I. expertise as regulation looms*, Fortune, 2021
- *AI Bias Problem Needs More Academic Rigor, Less Hype*, Datanami, 2021
- *AI in banking: Where it works and where it doesn't*, American Banker, 2021
- *AI Compliance Playbook: Seven Questions to Ask Before Regulators or Reporters Do*, Cybersecurity Law Report, 2021
- *What algorithm auditing startups need to succeed*, VentureBeat, 2021
- *Immuta Chief Legal Officer Joins Data Scientist to Create Boutique AI Law Firm*, Law.com, 2020
- *Grilling the answers: How businesses need to show how AI decides*, Computer Weekly, 2020
- *Data 2020 Outlook Part II: Explainable AI and Multi-model Databases*, ZDNet, 2020
- *Companies want explainable AI, vendors respond*, SearchCIO, 2017
- *Understanding cognitive computing: How much is enough?*, KMWorld, 2016
- *Why The Golden Age Of Machine Learning is Just Beginning* The Next Platform, 2015

Data Mining Competitions

- Kaggle Grupo Bimbo Inventory Demand Forecasting Contest, Top 1%, 14th of 1969 teams
- Kaggle Acquire Valued Shoppers Contest, Top 10%, 81st of 952 teams
- Cloudera Data Science Challenge 2: Medicare Anomaly Detection, 1 of 11 accepted solutions

Social Media

- GitHub: 3,000+ stars
- LinkedIn: 10,000+ connections
- Quora: 1,000,000+ content views