

Foster Provost

fprovost@stern.nyu.edu

RESEARCH INTERESTS

Data Science
Machine Learning
Artificial Intelligence

Crowdsourcing & Human-AI Loops
Advertising and Advertising Technology
Network Data

EDUCATION

Ph.D. Computer Science, University of Pittsburgh, 1993
M.S. Computer Science, University of Pittsburgh, 1988
B.S. Physics and Mathematics, Duquesne University, 1986

HONORS & AWARDS

European Research Paper of the Year 2017 (AIS, CIONET – for a paper in *MISQ*)
Best Paper Award, *Information Systems Research* (2016)
INFORMS President’s Pick, for paper published in *Information Systems Research* (Fall 2015)
2014 NYU/Stern MSBA Best Teacher Award
3 of the Top-5 Most Downloaded Papers (including #1) in the journal *Big Data* (as of Nov. 2015)
Nominated for 2014 NYU/Stern Professor of the Year by the MBA student body (one of 31 nominees)
Nominated for 2013 NYU/Stern Professor of the Year by the MBA student body (one of 23 nominees)
Best Paper Award, ACM SIGKDD 2012, Industry Track
People’s Choice Award, Runner-up, 2012 Future of Advertising Conf. on Empirical Generalizations in Advertising
Wharton Customer Analytics Initiative Awards, 2011, 2012, 2014 (one of several)
Marketing Science Institute Award, 2009 (one of several)
The INFORMS Design Science Award, December 2009
Best Paper Award Runner-up, ACM SIGKDD 2008
Elected as one of ten founding board members of the Intl. Machine Learning Society, Fall 2001 (Reelected 2006)
Second Place in KDDCUP-2003; Third Place in KDDCUP-2003
IBM Faculty Awards, 2000 & 2001
Best Paper Award, Third International Conference on Knowledge Discovery and Data Mining (KDD-97)
President's Award, NYNEX Science and Technology (now Verizon), 1995 (for a data analytics model that prioritizes network elements in need of rehabilitation)
Graduated First in Class, Duquesne University, 1986. Summa Cum Laude (4.0 GPA), and awarded the Certificates of General Excellence, Excellence in Mathematics, and Excellence in Physics

EMPLOYMENT HISTORY

Professor of Data Science, New York University, 2015-present
Interim Director, Center for Data Science, NYU, 2015-2016
Professor of Information Systems, Stern School of Business, New York University, 2008-present

Associate Professor of Information Systems, Stern School of Business, New York University, 2001-2008
Assistant Professor of Information Systems, Stern School of Business, New York University, 1999-2001

Research Scientist, NYNEX Science and Technology, Inc. (now Verizon), 1994-1999

ADVISORY POSITIONS

Scientific Advisory Board, ISI Foundation, 2015-

Academic Advisory Committee, NYU Skirball Center for the Performing Arts, 2015-

NYU Venture Fund Investment Review Board, 2013-

Wharton Customer Analytics Research Council, 2011-

Academic Advisor, NYC Media Lab, 2011-2014

Chief Scientist (Consulting), Coriolis Ventures, 2008-2011

DARPA Workshop on Network Science, at the intersection of Economics and Machine Learning, October 2005.

U.S. Treasury Department (future efforts towards Resilient Financial Information Systems), March 2005.

National Academies' Computer Science and Telecommunications Board planning workshop on information fusion for counter-terrorism, June 2002.

White House Office of Science and Technology Policy (data mining for counter-terrorism), February 2002

COMPANIES CoFOUNDED (in one sense or another)

Dstillery (data-science-based advertising; formerly Media6Degrees), 2008-
2014 Forbes' America's Most Promising Companies (#47)

Integral Ad Science (data science/analytics; formerly AdSafe Media), 2009-
2015 Forbes' America's Most Promising Companies (#15)
Crain's New York Business 6th fastest growing company in NY (2013)
2012 Product of the Year - Business Intelligence Group Awards (Professional Services)

Everyscreen Media (data-driven mobile advertising), 2011-2013
Acquired by Media6Degrees, May 2013

Predicube (data-driven advertising in Europe), 2014-
2016 Annual Master of Media Award - First place for Best Innovation (United Media Agencies, Belgium)

Detectica (detecting and controlling financial malfeasance), 2015-

UNIVERSITY/SCHOOL/DEPARTMENTAL SERVICE

University: Interim Director of NYU's Center for Data Science (2015-2016); NYU Committee on Intellectual Property Policy (2013-); Steering Committee for NYU Moore Sloan Data Science Environment Initiative--\$37.8M program with Berkeley and U.Wash. (2013-2016); NYU Venture Fund Investment Review Board (2013-); Skirball Center Academic Advisory Committee (2015-); NYU MSDS Curriculum Committee (2014-); University Working Group for 2014 Middle States Self-Study (2012-2014); University Working Group on Data Science Center and Data Science degree programs (2012-2013);

School/Department: Co-designed & Hosted MSBA Business Analytics Symposium, "Competing in the Connected Economy: Applications & Implications" (Fall 2016); Design & Host Stern Biz Analytics Club event, "Whiskey Analytics (with tasting)" (each semester); Stern MSBA curriculum committee (2015-2016); IOMS/IS curriculum committee (2014-); Co-organized the IOMS Dept. Research Fest, Fall 2014; IS Group faculty recruiting committee (2012-2013); Statistics Group faculty recruiting committee (2010-2011); Deans' Faculty Advisory Committee (2007-2010); IS Group faculty recruiting committee (2000-2010); Stern Ph.D. Program Oversight Committee (2008, 2004); Cochair of faculty recruiting committee for IS Group (2007-2008); Chair of faculty recruiting for IS Group (2003-2005); Ph.D. recruiting committee for IS group (2007); Led the design of proposal for school-wide incentive policy for faculty procurement of external funding (2007-2009); MBA Portfolio Committee (2004-2005); U.G. Dean's Faculty Advisory Committee (2004-2005); *and others going back further.*

SELECTED RECENT TEACHING & COURSE DEVELOPMENT

MSBA Class: Introduction to Business Analytics, Developed & Taught 2017.
TRIUM Exec Education Class: Data Analytics, Developed & Taught 2015-2016
MBA Class: Data Science for Business Analytics, Developed & Taught Spring 2017
MSDS/MBA Class: Intro to Data Science/Data Mining for Biz Analytics, Developed & Taught Fall 2013, 2015.
MSBA Class: Data Science for Business Analytics: Developed & Taught 2013, 2014, 2016
MBA Class: Practical Data Science: Developed & Taught (with Josh Attenberg) Fall 2012.
MBA Class: Data Mining for Business Analytics: Developed & Taught Fall 2005, Spring 2007-2013.
Ugrad Class: Data Mining for Business Analytics: Developed & Taught Spring 2007-2011.
MBA/PhD Class: Networks, Crowds & Markets; Developed and Taught (with lots of help) S2011. [sold out 1st run]
Ph.D. Seminar: Data Science Research. Developed & Taught: S2012, S2013, S2014, S2017.
Ph.D. Seminar: Technical Information Systems Research. Developed & Taught: S2000, etc.
MBA Core Class: Information Technology and Electronic Commerce. Taught: S2000, S&F2001, S&F2002.
MBA Class: Information and Internet Technologies. Developed (w/ V.Vassalos): S2000. Taught: S2000, F2000.

GRANTS

- PI: Marketing Science Institute Award (\$6500) "On-line brand advertising using social networks induced from user-generated content," 2009.
- Co-PI: IBM OCR Award (\$200,000) "Data and Risk." 2008-2010.
- PI: DOD/AFRL Grant (\$186,925) Evidence Assessment, Grouping, Linking, and Evaluation Program, 2004-5.
- PI: NYSIA Grant (\$35,000), "Algorithms and software for suspicion scoring," 2004-2005.
- PI: NSF Grant IIS-0329135 (\$67,087), "Generation & Evaluation of Confidence Bounds for ROC Curves," 2003-2004.
- PI: DARPA Grant (\$825,000), Evidence Extraction and Link Discover Program, 2001-2003.
- PI: IBM Faculty Awards (\$80,000), "Extending the Reach of Data Mining," 2000, 2001.
- Co-PI: NSF Grant NSF-IRI-9412549 (\$240,101), "Using Parallelism to Scale Up Machine Learning to Large Data-Analysis Problems" 1994.

PATENTS

- System, method, and computer-accessible medium for the prediction of user demographics on online news items (pending)
- Methods, Systems, and Media for Identifying Errors in Predictive Models using Annotators (#9,311,599 B1)
- Privacy-sensitive Methods, Systems, and Media for Geo-Social Targeting (pending)
- Methods for explaining document classifications (allowed by Patent Office; pending issuance)
- Methods, computer-accessible medium and systems for construction and inference of networked data in a bank setting (pending)
- Methods, systems, and media for providing hybrid learning approaches (pending)
- Methods, systems, and media for determining location information from real-time bid requests (#9,014,717) (#9,179,264)
- Methods, systems, and media for applying scores and ratings to web pages, web sites, and content for safe and effective online advertising (#8,732,017)
- Privacy-sensitive methods, systems, and media for targeting online advertisements using brand affinity modeling (pending)
- System, process and software arrangement for assisting with a knowledge discovery process (#7,167,846)
- Automated design of fraud detection systems (#5,790,645)

FACULTY FELLOWSHIPS

- Andre Meyer Faculty Fellow (2014-2017)
- NEC Faculty Fellow (2011-2014)
- NEC Faculty Fellow (2008-2011)
- Paduano Fellow in Business Ethics (2007-2008)

Ph.D. STUDENTS & POSTDOCS

Maytal Saar-Tsechansky, Associate Professor, University of Texas at Austin
Stern School, NYU, Ph.D. 2002

Outstanding Winner (first place), eBRC/SAP Doctoral Award, 2001

Claudia Perlich, Chief Scientist, Dstillery (formerly at IBM Research)
Stern School, NYU, Ph.D. 2005

Winner (one of four), eBRC/SAP Doctoral Award, 2002

Shawndra Hill, Senior Research Scientist, Microsoft Research (formerly at Wharton)
Stern School, NYU, Ph.D. 2007

Winner of the Herman E. Kroos Award for the best dissertation at NYU's Stern School, 2007

Finalist, George B. Dantzig Dissertation Award (INFORMS)

Josh Attenberg, CEO/CTO Detectica (formerly Director, Data Science at Etsy)
NYU/Poly Computer Science Ph.D., 2014

Xiaohan Zhang, Senior Data Scientist, Integral Ad Science
NYU/Stern Ph.D., 2014

Jessica Clark, Assistant Professor, University of Maryland (starting Fall 2017)
Stern School, NYU, Ph.D. 2017 (expected)

Robert Moakler, Ph.D. student, NYU/Stern (expected to complete in 2017)

Enric Junqué de Fortuny, Assistant Professor, NYU Shanghai
Univ. Antwerp Ph.D. 2014 (Co-advised with Prof. D. Martens of Univ. Antwerp)

Gary Weiss, Associate Professor, Fordham University, Dept. of Computer and Information Sciences
Ph.D. Rutgers University 2003 (Co-advised with Prof. H. Hirsh of Rutgers University)

Sofus Macskassy, Head of Data Analytics, Branch Metrics (Ph.D. Rutgers University 2003)
Formerly Director, Fetch Labs, Fetch Technologies
Associate Research Scientist, Stern School. 2004-2005

Victor Shengli Sheng, Associate Professor, Univ. Central Arkansas (Ph.D. Univ. W. Ontario, 2007),
Associate Research Scientist, Stern School. 2007-2009.

PROFESSIONAL SERVICE

Editor-in-Chief:

Machine Learning journal (January 2004 – June 2010)

Founding Board Member (elected):

International Machine Learning Society (2001-2010)

Program Chair/Organizer:

The 1st, 2nd, 3rd, 4th, 5th and 6th Workshops on Information in Networks (WIN), Sept 2009-2013, 2015. (Social Networks Summit among researchers from Computer Science, Economics, Management, Physics, Sociology, and other disciplines.) (co-organizer/co-chair with Sinan Aral & Arun Sundararajan)

The 1st & 2nd Human Computation Workshops (HCOMP-2009, -2010) (First workshops on “human computation” – crowdsourcing, micro-outsourcing, etc. for solving large or difficult problems; co-located with SIGKDD 2009 & SIGKDD 2010) (co-founder/co-organizer/co-chair with several others)

The 1st Workshop on Social Media Analytics (SOMA-2010) (co-organizer/co-chair with Prem Melville & Jure Leskovec)

Fourth Symposium on Statistical Challenges in Electronic Commerce Research (SCECR 2008) (co-organizer/co-chair with A. Ghose, and A. Sundararajan)

Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2001) (program co-chair with R. Srikant)

Editorial Boards:

Data Mining and Knowledge Discovery journal [2007-present]

Machine Learning journal [199?-present] (Action Editor, 2001-2003)(Editor-in-Chief, 2004-2010)

Journal of Machine Learning Research (JMLR) [2000-2010]

Journal of Artificial Intelligence Research (JAIR) [2002-2005]

Guest Editor:

Data Mining and Knowledge Discovery 5 (1/2) (January/March 2001). Special issue on "Applications of Data Mining to Electronic Commerce" (with Ron Kohavi).

Machine Learning 30 (2/3) 1998. Special issue on “Applications of Machine Learning and the Knowledge Discovery Process” (with Ron Kohavi).

PUBLICATIONS

PAPERS SUBMITTED/UNDER REVISION/IN PREPARATION

1. “Big Data, Data Science, and Civil Rights.” Barocas, Bradley, Honavar, and Provost. *Invited paper for the Computing Community Consortium of the Computing Research Association (CRA)*.
2. “Classification over bipartite graphs through projection.” M. Stankova, D. Martens, and F. Provost. *Under revision*.
3. “Examining dimensionality reduction for predictive modeling.” Clark & Provost. *Under review*.
4. “Benchmarking Study of Classification Techniques for Behavioral Data.” De Cnudde, Martens, Evgeniou, Provost. *Under review*.
5. “Who’s Watching TV?” J. Clark, J.-F. Paiement, F. Provost. *Under review*.
6. “In Pursuit of Enhanced Customer Retention Management.” E. Ascarza, S. Neslin, O. Netzer, Z. Anderson, P. Fader, S. Gupta, B. Hardie, A. Lemmens, B. Libai, D. Neal, F. Provost, R. Schrif. *Under Review*.
7. “Privacy, transparency and control for predictive analytics from massive fine-grained personal data.” Chen, Fraiberger, Moakler & Provost. *In preparation*.
8. “The value of data in predictive modeling: A case study in online advertising.” J. Moeyersoms, B. Dalessandro, F. Provost, and D. Martens. *In preparation*.
9. “Big Data is Small Data.” J. Clark & F. Provost. *Nascent*.
10. “Estimating Audience Interest Distributions based on Audience Web Behavior.” X. Zhang, F. Provost, K. Tsemekhman. *In limbo*.

MUSIC ALBUM

11. Hurvich, C. and F. Provost. *Mean Reversion*. (Original Art Rock) Pre-release February 2015. To be released by Kloster Art Music.

BOOKS

12. Provost, F. and T. Fawcett. *Data Science for Business: What you need to know about data mining and data analytic thinking*. O'Reilly Media, 2013. **Already adopted by >150 universities in at least 36 countries**, including Cambridge, Carnegie Mellon, Duke, Georgia Tech, Harvard, INSEAD, MIT, Stanford, U. Minnesota, U. Texas Austin, Wharton, and others. (see <http://data-science-for-biz.com> for a partial list). **One of Fortune Magazine's five "must read books for MBAs", December 2014.**
13. Provost, F. and R. Srikant (Eds.). *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM Press, 2001.
14. Kohavi, R. and F. Provost (Eds.). *Applications of Data Mining to Electronic Commerce*. Kluwer Academic Publishers, 2001.

JOURNAL ARTICLES

15. "Cost-Effective Quality Assurance in Crowd Labeling." J. Wang, P. Ipeirotis, and F. Provost (2017). *Information Systems Research*, in press. <http://pubsonline.informs.org/doi/10.1287/isre.2016.0661>
16. "Wallenius Naïve Bayes." E. Junqué de Fortuny, D. Martens, and F. Provost. *Accepted to Machine Learning journal, subject to minor revisions*.
17. "Mining Massive Fine-Grained Behavior Data to Improve Predictive Analytics." Martens, D., F. Provost, J. Clark, E. Junqué de Fortuny. *MIS Quarterly*, Vol. 40 No. 4, pp. 869-888/December 2016. **European Research Paper of the Year 2017 (organized by AIS and CIONET)**
18. "Finding Mobile Consumers with a Privacy-Friendly Geo-Similarity Network." Provost, F., D. Martens, and A. Murray. *Information Systems Research* 26(2), 243-265, 2015. **Best Paper Award – ISR 2016. Also, INFORMS President's Pick (Fall 2015).**
19. "Evaluating and Optimizing Online Advertising: Forget the click, but there are good proxies." B. Dalessandro, R. Hook, C. Perlich, F. Provost. (Preliminary version won **People's Choice Award** Runner-up at the 2012 Empirical Generalizations in Advertising Conference.) *Big Data* 3(2), 90-102, 2015.
20. "Beat the Machine: Challenging Humans to Find a Predictive Model's 'Unknown Unknowns.'" Attenberg, J., P. Ipeirotis, F. Provost. *ACM Journal of Data and Information Quality* 6(1) 2015.
21. "A Data Scientist's Guide to Startups." F. Provost, G. Webb, R. Bekkerman, O. Etzioni, U. Fayyad, C. Perlich. *Big Data* 2(3):117-128, September 2014. (**The journal's 5th most downloaded paper, with over 11,000 downloads** as of November 2015).
22. "Machine learning for targeted display advertising: Transfer learning in action." C. Perlich · B. Dalessandro · O. Stitelman · T. Raeder · F. Provost. *Machine Learning* 95(1):103-127, 2014.
23. "Explaining Data-Driven Document Classifications." Martens, D. and F. Provost. *MIS Quarterly* 38(1), 73-99, 2014.
24. "Repeated Labeling Using Multiple Noisy Labelers." Ipeirotis, P., F. Provost, V. Sheng, J. Wang. *Data Mining and Knowledge Discovery* 28(2), 402-441, 2014.

25. "Predictive Modeling with Big Data: Is Bigger Really Better?" E. Junqué de Fortuny, D. Martens, and F. Provost. *Big Data* 1(4), 215-226, 2013 (**the journal's 3rd most downloaded paper, with over 14,000 downloads as of Nov. 2015**).
26. "Information in Digital, Economic and Social Networks." Sundararajan, A., F. Provost, G. Oestreicher-Singer, and S. Aral. *Information Systems Research* 24(4), 883-905, 2013.
27. "Data Science and its Relation to Big Data and Data-driven Decision Making." F. Provost and T. Fawcett. *Big Data* 1(1), 51-59, 2013 (**the journal's #1 most downloaded paper, with over 27,000 downloads as of November 2015**).
28. "Inactive Learning? Difficulties Employing Active Learning in Practice." Attenberg, J. and F. Provost. *SIGKDD Explorations* 12(2), 2010.
29. Saar-Tsechansky, M., P. Melville, and F. Provost. "Active Feature-value Acquisition." *Management Science* 55(4): 664-684, 2009.
30. Saar-Tsechansky, M. and F. Provost. "Handling Missing Features when Applying Classification Models." *Journal of Machine Learning Research* 8(Jul):1623-1657, 2007.
31. Macskassy, S. and F. Provost. "Classification in Networked Data: A toolkit and a univariate case study." *Journal of Machine Learning Research* 8(May):935-983, 2007.
32. Saar-Tsechansky, M. and F. Provost. "Decision-centric Active Learning of Binary-Outcome Models" *Information Systems Research* 18(1), March 2007, pp. 4-22.
33. Hill, S., F. Provost, and C. Volinsky. "Network-based marketing: Identifying likely adopters via consumer networks." *Statistical Science* 21(2), 256-276, 2006.
34. Perlich, C. and F. Provost. "Distribution-based Aggregation for Relational Learning from Identifier Attributes." *Machine Learning* 62 (1/2) 65-105, 2006.
35. Bernstein, A., F. Provost and S. Hill. "Intelligent Assistance for the Data Mining Process: An Ontology-based Approach." *IEEE Transactions on Knowledge and Data Engineering* 17(4), pp. 503-518. April 2005.
36. Saar-Tsechansky, M. and F. Provost. "Active Sampling for Class Probability Estimation and Ranking." *Machine Learning* 54(2) 2004, 153-178.
37. Hill, S. and F. Provost. "The Myth of the Double-Blind Review? Author Identification Using Only Citations." *SIGKDD Explorations* 5(2) 2003, 114-119.
38. Perlich, C., F. Provost, and S. Macskassy. "Predicting citation rates for physics papers: Constructing features for an ordered probit model." *SIGKDD Explorations* 5(2) 2003, 89-90.
39. Weiss, G. and F. Provost. "Learning when Training Data are Costly: The effect of class distribution on tree induction." *Journal of Artificial Intelligence Research* 19 (2003) 315-354.
40. Perlich, C., F. Provost, and J. Simonoff. "Tree Induction vs. Logistic Regression: A Learning-curve Analysis." *Journal of Machine Learning Research* 4 (2003) 211-255.
41. Provost, F. and P. Domingos. "Tree Induction for Probability-based Ranking." *Machine Learning* 52:3, 199-215, 2003.
42. Provost, F. and T. Fawcett. "Robust Classification for Imprecise Environments." *Machine Learning* 42, 203-231, 2001.

43. Dhar, V., D. Chou, and F. Provost, "Discovering Interesting Patterns for Investment Decision Making with GLOWER -- A Genetic Learner Overlaid With Entropy Reduction," *Data Mining and Knowledge Discovery* 4(4): 251-280, 2000.
44. Provost, F. and V. Kolluri, "A Survey of Methods for Scaling Up Inductive Algorithms." *Data Mining and Knowledge Discovery* 3 (2): 131-169, 1999.
45. Provost, F. and A. Danyluk, "Problem Definition, Data Cleaning and Evaluation: A Classifier Learning Case Study." *Informatica* 23: 123-136, 1999.
46. Fawcett, T. and F. Provost, "Adaptive Fraud Detection." *Data Mining and Knowledge Discovery* 1 (3):291-316, 1997.
47. Provost, F. and J. Aronis, "Scaling Up Machine Learning with Massive Parallelism." *Machine Learning* 23 (1):33-46, 1996.
48. Provost, F. and B. Buchanan, "Inductive Policy: The Pragmatics of Bias Selection." *Machine Learning* 20 (1/2):35-61, 1995
49. Krenzelok, E. and F. Provost, "The Ten Most Common Plant Exposures Reported to Poison Information Centers in the United States." *Journal of Natural Toxins* 4 (2): 195-202, 1995.
50. Provost, F. and R. Melhem, "A Distributed Algorithm for Embedding Trees in Hypercubes with Modifications for Run-Time Fault Tolerance." *Journal of Parallel and Distributed Computing* 14 (1): 85-89, 1992.
51. Clearwater, S., W. Cleland, F. Provost, E. Stern and Z. Zhang, "A Real-Time Expert System for Trigger Logic Monitoring." *Nuclear Instruments and Methods in Physics Research* A293: 491-495, 1990.

OTHER JOURNAL PAPERS

52. Provost, F., Invited comment on Bolton and Hand's "Statistical Fraud Detection." *Statistical Science* 17 (2003).
53. Kohavi, R. and F. Provost, "Applications of Data Mining to Electronic Commerce." Guest editorial in *Data Mining and Knowledge Discovery* 5 (1/2) 2001.
54. Provost, F. and R. Kohavi, "On Applied Research in Machine Learning." Guest editorial in *Machine Learning* 30 (2/3) 1998.

SELECTED CONFERENCE PAPERS (HIGHLY SELECTIVE CONFERENCES)

* A star indicates that the content of a paper is included completely or to some extent in a previously mentioned published paper, as detailed in braces following each case.

55. "Iteratively Refining SVMs using Priors." E. Junque de Fortuny, T. Evgeniou, F. Provost, and D. Martens. *IEEE International Conference on Big Data (IEEE BigData 2015)*.
56. "Measuring causal impact of online actions via natural experiments: application to display advertising." D. Hill, R. Moakler, A. Hubbard, V. Tsemekhman, F. Provost, K. Tsemekhman. In the Proceedings of the Twenty-first ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (**KDD 2015**).

57. "Corporate Residence Fraud Detection." D. Martens, et al. In *Proceedings of the Twentieth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2014)*.
58. "Scalable Hands-Free Transfer Learning for Online Advertising." B. Dalessandro, et al. In *Proceedings of the Twentieth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2014)*.
59. "Using Co-visitation Networks For Classifying Non-Intentional Traffic." Stitelman, O., C. Perlich, B. Dalessandro, R. Hook, T. Raeder, and F. Provost. In *Proceedings of the Nineteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2013)*.
60. "Scalable Supervised Dimensionality Reduction Using Clustering." Raeder, T., C. Perlich, B. Dalessandro, O. Stitelman, and F. Provost. In *Proceedings of the Nineteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2013)*.
61. "Bid Optimizing and Inventory Scoring in Targeted Online Advertising." Perlich, C., B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. In *Proceedings of the Eighteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2012)*. **Best Paper Award**, Industry & Government Track.
62. "Design Principles of Massive, Robust Prediction Systems." Raeder, T., O. Stitelman, B. Dalessandro, C. Perlich, and F. Provost. In *Proceedings of the Eighteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2012)*.
63. "Online Active Inference and Learning." Attenberg, J. and F. Provost. In *Proceedings of the Seventeenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2011)*.
64. "Why label when you can search? Alternatives to active learning for applying human resources to build classification models under extreme class imbalance." Attenberg, J. and F. Provost. In *Proceedings of the Sixteenth ACM SIGKDD International Conf. on Knowledge Discovery and Data Mining (KDD 2010)*.
65. Attenberg, J. P. Melville and F. Provost. "A Unified Approach to Active Dual Supervision." In *Proceedings of the European Conference on Machine Learning and Principles of Knowledge Discovery in Databases (ECML PKDD 2010)*.
66. Provost, F., B. Dalessandro, R. Hook, X. Zhang & A. Murray. Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting. In *Proceedings of the Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2009)*.
67. *Sheng, S., F. Provost, and P. Ipeirotis. "Get Another Label? Improving Data Quality and Data Mining Using Multiple, Noisy Labelers." In *Proceedings of the Fourteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2008)*. **Best Paper Award Runner-up**.
68. Provost, F., P. Melville, and M. Saar-Tsechansky. "Data acquisition and cost-effective predictive modeling: targeting offers for electronic commerce." In *Proceedings of the Ninth International Conference on Electronic Commerce*, August 2007.
69. Melville, P., M. Saar-Tsechansky, F. Provost, and R. Mooney, "An Expected Utility Approach to Active Feature-value Acquisition." *Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM-2005)*, pp. 483-486. [short paper]
70. Macskassy, S., F. Provost, S. Rosset. "ROC Confidence Bands: An Empirical Evaluation." In *Proceedings of the 22nd International Conference on Machine Learning (ICML-2005)*. [acceptance rate: 27%] [Also appears in the ICML-2005 Workshop on ROC Analysis in Machine Learning (ROCML-2005).]

71. Macskassy, S. and F. Provost. "Suspicion scoring based on guilt-by-association, collective inference, and focused data access." In *Proceedings of the 2005 International Conference on Intelligence Analysis (IA '05)*. [acceptance rate: 25%]
72. Melville, P., M. Saar-Tsechansky, F. Provost, and R. Mooney, "Active Feature-Value Acquisition for Classifier Induction." *Proceedings of the Fourth IEEE International Conference on Data Mining (ICDM-2004)*, pp. 483-486. [short paper]
73. *Perlich, C., and F. Provost. "Aggregation-based Feature Invention and Relational Concept Classes." In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2003)*. [acceptance rate: 13%] {The second half of this paper is extended in paper 34.}
74. Macskassy, S., H. Hirsh, F. Provost, R. Sankaranarayanan, V. Dhar. "Intelligent Information Triage." In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-2001)*. [acceptance rate: 23%]
75. *Saar-Tsechansky, M. and F. Provost. "Active Learning for Class Probability Estimation and Ranking." In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01)*. [acceptance rate: 24%] {Much of the work reported in this paper is included in paper 36.}
76. Provost, F., D. Jensen and T. Oates, "Efficient Progressive Sampling." *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-99)*. [acceptance rate: 9%]
77. Fawcett, T. and F. Provost, "Activity Monitoring: Noticing Interesting Changes in Behavior." *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-99)*. [acceptance rate: 9%]
78. *Provost, F. and T. Fawcett, "Robust Classification Systems for Imprecise Environments." In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*. [acceptance rate: 30%] {The work reported in this paper is included in paper 42.}
79. Provost, F., T. Fawcett, and R. Kohavi "The Case Against Accuracy Estimation for Comparing Classifiers." *Proceedings of the Fifteenth Intl. Conference on Machine Learning (ICML-98)*. [acceptance rate: 31%]
80. *Provost, F. and T. Fawcett, "Analysis and Visualization of Classifier Performance: Comparison under Imprecise Class and Cost Distributions." In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining (KDD-97)*. **Best Paper Award Winner**. [acceptance rate: 10%] {The work reported in this paper is included in paper 42.}
81. *Provost, F. and V. Kolluri, "Scaling Up Inductive Algorithms: An Overview." In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining (KDD-97)*. [short paper; total acceptance rate (long+short): 41%] {The work reported in this paper is included in paper 44.}
82. Aronis, J. and F. Provost, "Increasing the Efficiency of Inductive Learning with Breadth-first Marker Propagation." In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining (KDD-97)*. [short paper; total acceptance rate (long+short): 41%]
83. Provost, F. and D. Hennessy, "Scaling Up: Distributed Machine Learning with Cooperation." In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*. [acceptance rate: 31%]
84. *Fawcett, T. and F. Provost, "Combining Data Mining and Machine Learning for Effective User Profiling." In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)*. [acceptance rate: 20%] {The work reported in this paper is included in paper 46.}

85. Aronis, J., F. Provost, and B. Buchanan, "Exploiting Background Knowledge in Automated Discovery." In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)*. [short paper; total acceptance rate (long+short): 33%]
86. Provost, F., "Iterative Weakening: Optimal and Near-Optimal Policies for the Selection of Search Bias." In *Proceedings of the Eleventh National Conf. on Artificial Intelligence (AAAI-93)*. [acceptance rate: 25%]
87. Danyluk, A. and F. Provost, "Small Disjuncts in Action: Learning to Diagnose Errors in the Telephone Network Local Loop." In *Proceedings of the Tenth International Conference on Machine Learning (ICML-93)*. [acceptance rate: 27%]
88. *Provost, F. and B. Buchanan, "Inductive Policy." In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*. [acceptance rate: 21%] {Some of the work reported in this paper is included in paper 48.}

ADDITIONAL CONFERENCE PAPERS

89. "Revealing Life Events from Inferred Customer Similarity: A Predictive Modeling Approach." T. Nian, Junqué de Fortuny, F. Provost. Proceedings of WITS 2014. [acceptance rate: 38%]
90. "Hyperlocal: Inferring Location of IP Addresses in Real-time Bid Requests for Mobile Ads." L. Long, T. Eliassi-Rad, F. Provost, and L. Moores. In Proceedings of the 6th ACM SIGSPATIAL International Workshop on Location-based Social Networks (LBSN 2013).
91. *Macskassy, S. and F. Provost. "Suspicion scoring of networked entities based on guilt-by-association, collective inference, and focused data access." In Annual Conference of the North American Association for Computational Social and Organizational Science (NAACSOS), 2005. [This is a followup paper to paper 71, with new results but considerable overlap.]
92. Macskassy, S. and F. Provost. "NetKit-SRL: A Network Learning Toolkit and its use for classification of networked data." In Annual Conference of the North American Association for Computational Social and Organizational Science (NAACSOS), 2005.
93. Kolluri, V., F. Provost, B. Buchanan, and D. Metzler. "Knowledge Discovery Using Concept-Class Taxonomies." In *Lecture Notes in Computer Science* 3339. *AI 2004: Advances in Artificial Intelligence: 17th Australian Joint Conference on Artificial Intelligence*. Springer-Verlag Heidelberg.
94. Provost, F. and D. Hennessey, "Distributed Machine Learning: Scaling up with Coarse-grained Parallelism." In *Proc of the Second International Conference on Intelligent Systems for Molecular Biology (ISMB-94)*.
95. Provost, F., "ClimBS: Searching the Bias Space." In *Proceedings of the Fourth International IEEE Conference on Tools with Artificial Intelligence (TAI-92)*. [acceptance rate: 36%]
96. Provost, F. and R. Melhem, "Embedding Rings in Hypercubes for Run-Time Fault Tolerance." In *Proceedings of the Fourth ISMM/IASTED Intl. Conference on Parallel and Distributed Computing and Systems*, 1991.
97. Clearwater, S. and F. Provost, "RL4: A Tool for Knowledge-Based Induction." In *Proceedings of the Second International IEEE Conference on Tools for Artificial Intelligence (TAI-90)*. [acceptance rate: 40%]

OTHER WORKSHOP/SYMPOSIUM PAPERS

98. “Enhancing Transparency and Control when Drawing Data-Driven Inferences about Individuals.” Daizhuo Chen, Samuel P. Fraiberger, Robert Moakler, Foster Provost. ICML-2016 Workshop on Human Interpretability in Machine Learning (WHI 2016). <https://arxiv.org/abs/1606.08063>
99. “Explaining Classification Models Built on High-Dimensional Sparse Data.” Julie Moeyersoms, Brian d’Alessandro, Foster Provost, David Martens. ICML-2016 Workshop on Human Interpretability in Machine Learning (WHI 2016). <https://arxiv.org/abs/1607.06280>
100. “Pleasing the advertising oracle. Probabilistic prediction from sampled, aggregated ground truth.” M. Williams, C. Perlich, B. Dalessandro, F. Provost. In Proceedings of the Eighth International Workshop on Data Mining for Online Advertising (ADKDD 2014).
101. Moakler, R., et al. “Causal impact of online advertisements using viewability as a method of treatment.” Winter Conference on Business Intelligence, Feb. 2014.
102. “Estimating Audience Interest Distributions based on Audience Web Behavior.” X. Zhang, F. Provost, K. Tsemekhman. Winter Business Intelligence Conference, 2012.
103. “Causally Motivated Attribution for Online Advertising.” B. Dalessandro, O. Stitelman, C. Perlich, and F. Provost. *ADKDD* 2012.
104. “Geo-social Network Advertising.” F. Provost, D. Martens and A. Murray. Winter Business Intelligence Conference, 2012.
105. “Estimating the Effect of Online Display Advertising on Browser Conversion.” O. Stitelman, B. Dalessandro, C. Perlich, and F. Provost. *ADKDD* 2011.
106. “Beat the Machine: Challenging crowdsourcing workers to find the unknown unknowns.” Attenberg, J., P. Ipeirotis & F. Provost. *Proceedings of the Third Human Computation Workshop (HCOMP 2011)*.
107. “Managing Crowdsourcing Workers.” with Jing Wang and Panos Ipeirotis. Winter Business Intelligence Conference, 2011.
108. “Quality management on Amazon Mechanical Turk,” P. Ipeirotis, F. Provost, and J. Wang. *Proceedings of the Second Human Computation Workshop (KDD-HCOMP 2010)*.
109. Attenberg, J. and F. Provost. “Active inference and learning for classifying streams.” In *ICML-2010 Workshop on Budgeted Learning*.
110. Attenberg, J., P. Melville and F. Provost. “Guided feature labeling for budget-sensitive learning under extreme class imbalance.” In *ICML-2010 Workshop on Budgeted Learning*.
111. Zheng, R., F. Provost, and A. Ghose. “Social Network Collaborative Filtering: Preliminary Results.” The Sixth Workshop on e-Business (WeB 2007).
112. Hill, S., F. Provost and C. Volinsky (2007). “Learning and Inference in Massive Social Networks.” The 5th International Workshop on Mining and Learning with Graphs, August 2007. Extended Abstract.
113. *Macskassy, S, and F. Provost (2006). “A brief survey of machine learning methods for classification in networked data and an application to suspicion scoring.” E.M. Airoidi et al. (Eds.): ICML 2006 Ws, *LNCS 4503*, pp. 172-175. Springer-Verlag. (Originally appeared at the Workshop on Statistical Network Learning at 23rd International Conference on Machine Learning (ICML 2006), Pittsburgh, 29 June 2006.) {There is considerable overlap between this paper and paper 31.}

114. *Melville, P., M. Saar-Tsechansky, F. Provost, R. Mooney. "Economical Active Feature-value Acquisition through Expected Utility Estimation." In the KDD-05 Workshop on Utility-based Data Mining. {There is considerable overlap between this paper and paper 55.}
115. Giraud-Carrier, C. and F. Provost. "Toward a Justification of Meta-learning: Is the No Free Lunch Theorem a Show-stopper?" In the ICML-2005 Workshop on Meta-Learning.
116. Macskassy, S., F. Provost, S. Rosset. "Pointwise ROC Confidence Bounds: An Empirical Evaluation." In the ICML-2005 Workshop on ROC Analysis in Machine Learning (ROCML-2005).
117. Hill, S., F. Provost and C. Volinsky. "The Gift of Gab: Evidence Tele-Commerce Firms can Profit from Viral Marketing." (Two-page extended abstract) First Interdisciplinary Symposium between Information Systems, Statistics and Related Fields. Decision and Information Technologies Department, Robert H. Smith School of Business, Univ. of Maryland. May 2005.
118. Macskassy, S. and F. Provost. "Confidence Bands for ROC Curves: Methods and an Empirical Study." First Workshop on ROC Analysis in AI. August 2004.
119. Macskassy, S. and F. Provost. "A Simple Relational Classifier." In the Second Workshop on Multi-Relational Data Mining (MRDM-2003) at KDD-2003.
120. Bernstein, A., S. Clearwater, and F. Provost. "The Relational Vector-space Model and Industry Classification." To appear in *Proceedings of the IJCAI-2003 Workshop on Learning Statistical Models from Relational Data*.
121. *Provost, F., C. Perlich, and S. Macskassy. "Relational Learning Problems and Simple Models." To appear in *Proceedings of the IJCAI-2003 Workshop on Learning Statistical Models from Relational Data*. {Some of the work reported in this position paper is included in paper 98.}
122. *Perlich, C. and F. Provost. "Aggregation and Concept Complexity in Relational Learning." To appear in *Proceedings of the IJCAI-2003 Workshop on Learning Statistical Models from Relational Data*. {Some of the work reported in this position paper is included in paper 55.}
123. Bernstein, A., S. Clearwater, S. Hill, C. Perlich, and F. Provost. "Discovering Knowledge from Relational Data Extracted from Business News." In *Proceedings of the KDD-2002 Workshop on Multi-Relational Data Mining*, 2002.
124. *Bernstein, A. and F. Provost. "An Intelligent Assistant for the Knowledge Discovery Process." In *Proceedings of IJCAI-01 Workshop on Wrappers for Performance Enhancement in Knowledge Discovery in Databases*, 2001. {The work reported in this paper is included in working paper 12.}
125. *Macskassy, S., H. Hirsh, F. Provost, R. Sankaranarayanan, V. Dhar. Information Triage using Prospective Criteria. In *Proceedings of the Eighth International Conference on User Modeling (UM-2001), Workshop on Machine Learning, Information Retrieval and User Modeling*, 2001. {The work reported in this paper also appears in paper 74.}
126. Aronis, J., V. Kolluri, F. Provost, and B. Buchanan, "The WoRLD: Knowledge Discovery from Multiple Distributed Databases." In *Proc. of the Florida Artificial Intelligence Research Symposium (FLAIRS-97)*.
127. Provost, F. and A. Danyluk, "Learning from Bad Data." In *Proceedings of the ML-95 Workshop on Applying Machine Learning in Practice*, 1995.

128. Krenzelok, E., F. Provost, T. Jacobsen, J. Aronis, B. Buchanan, "Assessing Patient Referral Patterns to a Health Care Facility in Plant Exposure Patients Using Computer Artificial Intelligence." European Assn. of Poison Centres and Clinical Toxicologists Scientific Meeting. May 18-20, 1995, Krakow, Poland.
129. Krenzelok, E., F. Provost, T. Jacobsen, J. Aronis, B. Buchanan, "Poinsettia (Euphorbia pulcherrima) Exposures Have Good Outcomes...Just As We Thought." European Association of Poison Centres and Clinical Toxicologists Scientific Meeting, 1995.
130. Aronis, J. and F. Provost, "Efficiently Constructing Relational Features from Background Knowledge for Inductive Machine Learning" In *Proceedings of the AAAI-94 Workshop on Knowledge Discovery in Databases*, (KDD-94).
131. Provost, F., "Goal-Directed Inductive Learning: Trading Off Accuracy for Reduced Error Cost." In *Proceedings of the AAAI Spring Symposium on Goal-Directed Learning*, 1994.
132. Danyluk, A. and F. Provost, "Adaptive Expert Systems: Applying Machine Learning to NYNEX MAX." In *Proceedings of the AAAI-93 Workshop: AI in Service and Support--Bridging the Gap between Research and Applications*, 1993.
133. Provost, F., "A Baseline Taxonomy of Bias Adjustment Policies." In *Proceedings of the ML-92 Workshop on Biases in Learning*, 1992.

BOOK CHAPTERS & OTHER PUBLICATIONS

134. "Selective Data Acquisition for Machine Learning." Attenberg, J., P. Melville, F. Provost and M. Saartsehansky. In B. Krishnapuram, S. Yu, and B. Rao (Eds.), *Cost-Sensitive Machine Learning*. Chapman & Hall/CRC, 2011.
135. Perlich, C. and F. Provost. "Aggregation for Predictive Modeling with Relational Data." *Encyclopedia of Data Warehousing and Mining*, 2005.
136. *Provost, F. and V. Kolluri, "Scalability." In W. Kloesgen and J. Zytkow (eds.), *Handbook of Data Mining and Knowledge Discovery*, Oxford University Press, 2002, pp. 418-433. {Some of the work reported in this paper is included in paper 44.}
137. Danyluk, A. and F. Provost, "Telecommunications Network Diagnosis." In W. Kloesgen and J. Zytkow (eds.), *Handbook of Data Mining and Knowledge Discovery*, Oxford University Press, 2002, pp. 897-902.
138. *Fawcett, T. and F. Provost, "Fraud Detection." W. Kloesgen and J. Zytkow (eds.), *Handbook of Data Mining and Knowledge Discovery*, Oxford University Press, 2002, pp. 726-731. {Some of the work reported in this paper is included in paper 46.}
139. Provost, F. "Learning with Imbalanced Data Sets 101." Invited paper for the AAAI-2000 Workshop on Learning with Imbalanced Data Sets, 2000.
140. *Provost, F., D. Jensen and T. Oates, "Progressive Sampling." In H. Liu and H. Motoda (eds.), *Instance Selection and Construction, A Data Mining Perspective*, Kluwer Academic Publishers, 2000. {The work reported in this paper is included in paper 76.}
141. *Provost, F., "Distributed Data Mining: Scaling up and beyond." In H. Kargupta and P. Chan (eds.), *Advances in Distributed Data Mining*, San Francisco, CA: Morgan Kaufmann, 2000. {Some of the work reported in this paper is included in paper 44.}

142. Danyluk, A., T. Fawcett, and F. Provost, "AI Approaches to Time-series Problems." Workshop report in *AI Magazine*, 1999.
143. Fawcett, T., I. Haimowitz, F. Provost, and S. Stolfo, "AI Approaches to Fraud Detection and Risk Management." (Extended) workshop report in *AI Magazine*, 1998.
144. Provost, F. and B. Buchanan, "Inductive Strengthening: The effects of a simple heuristic for restricting hypothesis space search." In K.P. Jantke (ed.), *Analogical and Inductive Inference (Lecture Notes in Artificial Intelligence 642)*. Springer-Verlag, 1992.
145. *Clearwater, S., W. Cleland, F. Provost, E. Stern and Z. Zhang, "A Real-Time Expert System for Experimental High Energy/Nuclear Physics." In D. Perrett-Gallix and W. Wojcik (eds.), *New Computing Techniques in Physics Research*. Paris: Centre National de la Recherche Scientifique, 1990. {The work reported in this paper is included in paper 51.}
146. *Provost F. and R. Melhem, "Distributed Fault Tolerant Embedding of Trees and Rings in Hypercubes." In I. Koren (ed.) *Defect and Fault Tolerance in VLSI systems Volume 1*. New York, NY: Plenum Press, 1989. {Some of the work reported in this paper is included in paper 50 and some is reported in paper 96.}

OTHER WORKING PAPERS

147. "Social Network Collaborative Filtering." Zheng, R., D. Wilkinson, and F. Provost. Working paper CeDER-8-08. Center for Digital Economy Research, Stern School of Business, New York University. September 2008.
148. *Provost, F. and P. Domingos. "Well-trained PETs: Improving Probability Estimation Trees" CeDER Working Paper #IS-00-04. {Early version of paper 41.}
149. Provost, F., J. Aronis, and B. Buchanan. "Rule-space Search for Knowledge-based Discovery." CIO Working Paper #IS 99-012.

SELECTED TUTORIALS

- "Predictive Modeling with Social Networks." Tutorial (invited) at the Conference on Complex Networks: from theory to interdisciplinary applications, July 2016.
- "Predictive Modeling with Social Networks." Tutorial (invited) at the Third International AAAI Conference on Weblogs and Social Media (with Jennifer Neville), May 2009.
- "Predictive Modeling with Social Networks." Tutorial at the Fourteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (with Jennifer Neville), August 2008.
- "Social Network Mining: A Tutorial on Inference and Learning with Social Network Data." Tutorial (invited) at the Twenty-Third AAAI Conf. on Artificial Intelligence (with Jennifer Neville), July 2008.
- "Inference and Learning with Networked Data." Invited tutorial lecture for the NATO Advanced Study Institute on Mining Massive Data Sets for Security, September 2007.
- "Modeling Complex Networks for Electronic Commerce." Tutorial at the ACM Conference on Electronic Commerce (with Arun Sundararajan), June 2007.
- "Evaluating Machine Learning, Knowledge Discovery, and Data Mining" (with David Jensen). Tutorial presented at **KDD-98**, at **AAAI-99** and at **IJCAI-99**.

SELECTED INVITED TALKS (Etc.)

- “Using Data Analytics to Enhance Compliance.” Invited panel presentation for Compliance: New Risks, Challenges, and Approaches, Conference held by NYU Program on Corporate Compliance and Enforcement, October 2016.
- “The Predictive Power of Massive Data about our Fine-grained Behavior.” Keynote talk for the 15th International Symposium on Intelligent Data Analysis (IDA 2016), October 2016.
- “Data-driven Models: Can we explain why they do what they do?” Invited talk for *Science Crossroads*, annual meeting of the Institute for Scientific Interchange Foundation, Turin Italy, October 2016.
- “Competing with Analytics in the Connected Economy.” Keynote talk for NYU/Stern’s Business Analytics Symposium, September 2016.
- “Data Science for Connected Commerce.” Invited Lecture at MIT Sloan School, September 2016.
- “Academia & Entrepreneurship.” Invited talk for Univ. of Alberta’s Machine Intelligence Institute, Sept. 2016.
- “The Predictive Power of Massive Data about our Fine-grained Behavior.” Keynote talk for the 9th ACM International Conference on Web Search and Data Mining (WSDM), February 2016.
- “Privacy, Confidentiality, Transparency & Control: From a fundamental tension in data science to the ‘explanation trick.’” Invited panel talk & discussion, FATML workshop at NIPS 2014.
- “Teaching Data Science.” Invited panel on Data Science Education, Palo Alto, December 2014.
- “Understanding Decisions Driven by Big Data: From Analytics Management to Privacy-friendly Cloaking Devices.” Keynote lecture at StrataEU, November 2014.
- “Predictive Analytics with Fine-grained Behavior Data.” Invited talk, Data-driven Business Day, StrataEU, November 2014.
- “Data Science for Business.” Invited Lecture at MIT Sloan School, October 2014.
- “Whiskey Analytics.” Lecture and whiskey tasting, inaugural event for Stern MBA Business Analytics Club, September 2014.
- “A Conversation with Foster Provost.” NYU President’s Council & NYU Giving Society Event, May 2014.
- “Predictive Models from Massive Fine-Grained Data Work Great, But Do You Understand What They're Doing?” Invited talk for MIT’s Workshop on Big Data Analytics, March 2014
- “Display Ad Targeting based on Fine-Grained Consumer Behavior.” Invited talk at the 2013 Marketing in Israel Conference, December 2013.
- “Data Science for Business.” Invited lecture (based on the book) at the Institute for Big Data Analytics at Dalhousie University, Nov 2013.
- “Big Data: Is Bigger Really Better? Predictive Modeling with Fine-Grained Behavior Data.” Keynote talk at Strata NY, Oct 2013
- “Mining (Massive) Consumer Behavior Data for Predictive Marketing.” Invited talk at CMU, Nov. 2013.
- “Mining (Massive) Consumer Behavior Data for Predictive Marketing.” Invited talk at Rutgers U., May 2013.
- “Mining (Massive) Consumer Behavior Data for Predictive Marketing.” Invited talk at the Microsoft Research Machine Learning Summit, Paris, April 2013.

- “Mining (Massive) Consumer Behavior Data for Marketing.” Keynote talk at the IEEE International Conference on Data Mining (ICDM), Industry Track, Dec 2012.
- “Mining (Massive) Consumer Behavior Data for Marketing.” Invited talk at IBM Research, Oct 2012.
- “Social Targeting for Privacy-Friendly On-line Advertising.” Invited talk at University of Antwerp, Nov. 2011.
- “Social Targeting for Privacy-Friendly On-line Advertising.” Invited talk for the 2011 International Conference on Weblogs and Social Media, July 2011.
- “Machine Learning for Display Advertising.” Keynote lecture for the 2010 Workshop on Machine Learning for Online Advertising, December 2010.
- “(Privacy-friendly!) Social Network Targeting for On-line Advertising” Invited talk for KDD-2010.
- “Get Another Label? Improving Data Quality and Machine Learning using Multiple, Noisy Labelers.” Invited talk for NYU/Poly, November 2010.
- “Get Another Label? Improving Data Quality and Machine Learning using Multiple, Noisy Labelers.” Invited talk for MIT Operations Research Center, November 2010.
- “Get Another Label? Improving Data Quality and Machine Learning using Multiple, Noisy Labelers.” Invited talk for Columbia University Center for Computational Learning Systems, April 2010.
- “Social Targeting for Privacy Friendly Online Advertising.” Invited talk for NYU/Stern/L2 Social Commerce Clinic, October 2010.
- “Social-Media-Based Targeting for Privacy-Friendly On-Line Advertising.” Edison Lecture, Notre Dame University, October 2010.
- “Active Inference and Learning for Data Streams.” Invited Talk for INFORMS, November 2010.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk for Wharton WIMI Conference on the Emergence and Impact of User-Generated Content, December 2009.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk for Tufts University, October 2009.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk at the INFORMS annual meeting, October 2009.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk at the Conference on Marketing meets Data Mining, August, 2009.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk at the Third International Workshop on Data Mining and Audience Intelligence for Advertising, June 2009.
- “Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting.” Invited talk for Carnegie Mellon University, April 2009.
- “Inference and Learning with (Social) Network Data.” Invited talk for Williams College, April 2008.
- “Inference and Learning with Networked Data.” Invited talk for KXEN Corp. Paris, January 2008.
- “Inference and Learning with Networked Data.” Invited talk reprise, Dept. of Applied Economic Sciences, Katholieke Universiteit Leuven, January 2008.
- “Inference and Learning with Networked Data.” Invited talk, Univ. Conn. Business School, November 2007
- “Inference and Learning with Networked Data.” Invited talk, Dept. of Applied Economic Sciences, Katholieke Universiteit Leuven, September 2007.
- “Inference and Learning with Networked Data.” Invited talk (one of two) for the Sixth Annual Computer Science Day, Fordham University, April 2007.
- “Inference and Learning with Networked Data.” Invited talk, INSEAD, March 2007.
- “Classification in Networked Data.” Invited talk, University of Zurich, November 2006.
- “Mining Social Networks for Fun and Profit.” Invited talk, University of Pittsburgh, October 2006.
- “Data Mining for Business Intelligence.” Invited talk for New York CTO Club, October 2005.
- “Economic Machine Learning and Utility-based Data Mining” Invited talk (one of three) KDD-2005 Workshop on Utility-based Data Mining.

- “Toward a Justification of Meta-learning: Is the No Free Lunch Theorem a Show-stopper?” Invited talk (one of one) ICML-2005 Workshop on Meta-Learning.
- “Classification and Learning with Networked Data” University of Notre Dame. April 2005.
- “Classification and Learning with Networked Data” Brigham Young University. March 2005.
- “Classification and Learning with Networked Data” Google Research. February 2005.
- “Defining the Research Agenda for Business Intelligence.” Invited panel presentation for the 2005 Utah Winter Information Systems Conference. David Eccles School of Business, University of Utah. March 2005.
- “Classification and Learning with Networked Data: Some observations and results.” Invited talk for the Red McComb’s School of Business and the Department of Computer Science, University of Texas at Austin, November 2004.
- “Active Information Acquisition for Machine Learning.” Invited talk for Amazon.com’s data mining group, August 2004.
- “Obtaining External Funding.” Invited talk for the Ph.D. Project in Information Systems Conference, August 2004.
- “The role of applications in the science of machine learning.” Invited talk (one of three) for the Twentieth International Conference on Machine Learning, August 2003.
- “Marginal class distribution and budget-sensitive learning.” Invited talk (one of three) for the ICML-2003 Workshop on Learning from Unbalanced Data Sets, August 2003.
- “Research in Data Mining and Machine Learning.” Invited talk for the Ph.D. Project in Information Systems Conference, August 2003.
- “The Machine Learning Journal.” Invited panel presentation for the Ph.D. Project in Information Systems Conference (journal editors panel), August 2003.
- “Data mining and cross-selling.” Invited talk for Pitney Bowes’ Advanced Concepts and Technology division, July 2003.
- “Data mining and business intelligence.” Invited talk for Pitney Bowes’ Advanced Concepts and Technology division, January 2003.
- “Adaptive Fraud Detection.” Invited talk for the Rutgers University Electronic Commerce Association, December 2002.
- “Tree induction versus logistic regression.” Keynote talk (one of four) at SAS’s annual Data Mining Conference, October 2002.
- “Tree induction versus logistic regression.” Invited talk at the annual meeting of the Classification Society of North America, June 2002.
- “Adaptive Fraud Detection.” Invited talk at Stevens Institute of Technology, April 2002.
- “Intelligent Assistance for the Knowledge Discovery Process.” Invited talk at INFORMS, Nov. 2001.
- “Reviving the Learning Curve: A critical study of machine learning performance.” Invited talk at the University of Texas at Austin, Data Mining Lecture Series, Oct. 2001.
- “Wrappers: Intelligent Support for the Knowledge Discovery Process.” Invited talk (one of two) for the IJCAI-01 Workshop on Wrappers for Performance Enhancement in KDD, August 4, 2001.

- “Machine Learning for Information Monitoring and Triage.” Invited talk (one of three) for the 14th Canadian Conference on Artificial Intelligence, June 2001.
- “Monitoring and Triage of Business News.” Invited talk for the DIMACS Working Group on Data Analysis and Digital Libraries, DIMACS Center, Rutgers University, May 17, 2001.
- "Robust Classification for Imprecise Environments." Invited presentation at INFORMS, Nov. 2000.
- "Information Monitoring and Triage for Mobile Alerting and Timely Decision Making." Mathematical Sciences Department, IBM T.J. Watson Research Center, October, 2000.
- "Tutorial Overview: The Cost-Sensitive Learning Problem(s) and a Review of the Main Existing Approaches to Solving Them." Invited (keynote) talk for the ICML Workshop on Cost-sensitive Learning, 2000.
- "Learning with Imbalanced Data Sets 101." Invited talk for the AAAI Workshop on Learning with Imbalanced Data Sets, 2000.
- "Well-behaved PETs: Improving Probability Estimation Trees." Invited presentation at University of Pittsburgh and at Washington State University, Fall 1999.
- “Rolling the Dice with Data Mining Success.” KDD-98 Workshop on the Keys to the Commercial Success of Data Mining, August 1998.
- “Distributed Data Mining: Scaling up and beyond.” KDD-98 Workshop on Distributed Data Mining, Aug 1998.
- “Evaluating Data Mining Results.” Annual Symposium of the ACM Central Ohio Chapter, 1998.
- “How Well Does Data Mining Work? Evaluating the Knowledge Discovered.” Third Annual Data Mining Summit, March 1998.
- “Comparing Classifiers under Imprecise Business Criteria.” Mathematical Sciences Department, IBM T.J. Watson Research Center, September 1997.
- “Fraud Detection Systems: Combining Data Mining and Machine Learning.” Second Annual Data Mining Summit, February 1997.
- “Automated Design of Fraud Detection Systems.” Brown University Industrial Partners Program, Nov. 1996.
- "Pushing Machine Learning into the Real World." NYNEX Science and Technology, May 1994.
- "What Do Learning Goals Look Like?" (Panel) AAAI Spring Symposium on Goal-Directed Learning, 1994.
- "Machine Learning & ‘Real-World’ Complications." NYNEX Science and Technology, May 1993.
- "Policies for Machine Learning." Mathematical Sciences Dept., IBM T.J. Watson Research Center, April 1993.
- "Policies for Inductive Bias Selection." NYNEX Science and Technology, May 1992.