

**MARIA-FLORINA BALCAN**  
Carnegie Mellon University  
Pittsburgh, PA 15213-3891  
[ninamf@cs.cmu.edu](mailto:ninamf@cs.cmu.edu)  
[www.cs.cmu.edu/~ninamf](http://www.cs.cmu.edu/~ninamf)

## RESEARCH INTERESTS

Learning Theory, Machine Learning, Theory of Computing, Artificial Intelligence, Algorithmic Economics and Algorithmic Game Theory, Optimization.

## APPOINTMENTS

- **June 2014 – present** Associate Professor, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- **2009 - 2014** Assistant Professor, College of Computing, Georgia Tech, Atlanta, GA.
- **2008 - 2009** Postdoctoral Researcher, Microsoft Research NE, Cambridge, MA.
- **2000 - 2002** Instructor, Computer Science Department, University of Bucharest, Romania.

## EDUCATION

- **Ph.D. 2002 – 2008.** Carnegie Mellon University, Pittsburgh, PA. Computer Science Department.
- **M.S. 2000 – 2002.** University of Bucharest, Romania, Faculty of Mathematics, Computer Science Department, M.S. Degree in Computer Science, GPA 10.00 / 10.00.
- **B.S. 1996 – 2000.** University of Bucharest, Romania. Faculty of Mathematics, Computer Science Dept., GPA 10.00/10.00. “Summa Cum Laude” Diploma.

## HONORS AND AWARDS

### Major Leadership Positions

- General Chair for the 38th International Conference on Machine Learning (**ICML**) 2021.
- 2017, Invited co-organizer for the “Machine Learning and Computational Modelling” session of the Japanese-American-German Kavli Frontiers of Science Symposium, National Academy of Sciences.
- Program Committee Chair for the 33rd International Conference on Machine Learning (**ICML**) 2016.
- Program Committee Chair for the 27th Annual Conference on Learning Theory (**COLT**) 2014.
- Board Member of the International Machine Learning Society, 2011 – present.

### Young Investigator Awards

- 2014, Sloan Research Fellowship.
- 2011, Microsoft Faculty Fellowship.
- 2009, NSF CAREER Award.

### Other Honors

- 2018, Amazon Research Award.
- 2015, Kavli Frontiers of Science Fellow, National Academy of Sciences.
- 2013, Georgia Power Professor of Excellence.
- 2013, Raytheon Faculty Fellowship.
- 2012, Runner Up Best Paper, 25th Annual Conference on Learning Theory (COLT 2012).
- 2011, Google Research Award.
- 2009, CMU School of Computer Science Distinguished Dissertation Award.

- 2008, Mark Fulk Best Student Paper Award, 21st Annual Conference on Learning Theory (COLT 2008).
- 2007 - 2008, IBM Ph.D. Fellowship.
- 2000 - 2001, Romanian Government Merit Fellowship (during my MS studies).
- 1996 - 2000, Romanian Government Merit Fellowship (during my undergraduate studies).
- 2001, World Bank Fellowship, for visiting CNRS, Toulouse, France.
- 1999 - 2000, European Union Erasmus/Socrates scholarship to study at the University of Patras, Greece.

## PROFESSIONAL ACTIVITIES

### Journal Editorial Boards

1. Associate Editor for ACM Transactions on Economics and Computation (**ACM TEAC**), 2018 – present.
2. Associate Editor for SIAM Journal on Data Science (**SIMODS**), 2018 – present.
3. Associate Editor for IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2017 – present.
4. Editor for Communications of the ACM (**CACM**), Research Highlights section, 2015 - present.
5. Action Editor, Machine Learning Journal (**MLJ**), 2015 – 2018.
6. Editor for the Encyclopedia of Algorithms, 2014 - 2015.
7. Editor for SIAM Journal on Computing (**SICOMP**) Special Issue for FOCS 2009.
8. Editor for Journal of Computer and System Sciences (**JCSS**), 2011 - 2013.

### Tutorial/Workshop Chair

1. Workshop Chair for the Annual Symposium on Foundations of Computer Science (**FOCS**) 2019 -2020.
2. Tutorial Chair for the International Conference on Machine Learning (**ICML**) 2019.

### Top Level Program Committee Member

1. Plenary Talks Selection Committee for ACM Symposium on Theory of Computing (**STOC**) 2018.
2. Area Chair for International Conference on Machine Learning (**ICML**) 2018.
3. Annual Conference on Learning Theory (**COLT**) 2018.
4. The 32nd Conference on Artificial Intelligence (**AAAI**) 2018.
5. Senior Area Chair for Neural Information Processing Systems (**NIPS**) 2017.
6. The 49th ACM Symposium on Theory of Computing (**STOC**) 2017.
7. The 18th ACM Conference on Economics and Computation (**EC**) 2017.
8. NIPS Workshop and Symposium Selection Committee 2016.
9. Thirtieth AAAI Conference on Artificial Intelligence (**AAAI**) 2016.
10. ACM-SIAM Symposium on Discrete Algorithms (**SODA**) 2016.
11. Area Chair for Neural Information Processing Systems (**NIPS**) 2015.
12. Area Chair for International Conference on Machine Learning (**ICML**) 2015.
13. The 18th International Conference on Artificial Intelligence and Statistics (**AISTATS**) 2015.
14. Area Chair for the Conference on Uncertainty in Artificial Intelligence (**UAI**) 2015.
15. Area Chair for International Conference on Machine Learning (**ICML**) 2013.
16. The Annual Symposium on Foundations of Computer Science (**FOCS**) 2013.
17. Annual Conference on Learning Theory (**COLT**) 2013.
18. Area Chair for the Conference on Uncertainty in Artificial Intelligence (**UAI**) 2012.
19. Workshop on Internet & Network Economics (**WINE**) 2012.
20. Annual Conference on Learning Theory (**COLT**) 2012.
21. Innovations in Computer Science (**ICS**) 2011.
22. Area Chair for Neural Information Processing Systems (**NIPS**) 2011.
23. Annual Conference on Learning Theory (**COLT**) 2011.
24. Area Chair for Neural Information Processing Systems (**NIPS**) 2010.

25. The 21st International Conference on Algorithmic Learning Theory (**ALT**) 2010.
26. The Annual Symposium on Foundations of Computer Science (**FOCS**) 2009.
27. Annual Conference on Learning Theory (**COLT**) 2009.

### **Events and Workshops Organized**

1. Co-organizer, Simons Symposium on “New Directions in Theoretical Machine Learning”, Spring 2019.
2. Co-organizer, IPAM Workshop on “New Architectures and Algorithms”, Fall 2018.
3. Co-Organizer, “Foundations of Machine Learning”, Semester long Program at the Simons Institute for the Theory of Computing, Spring 2017.
4. Co-organizer, Dagstuhl Workshop on “Game Theory Meets Computational Learning Theory”, Summer 2017.
5. Co-Organizer, Workshop on “Interactive Learning” within the Simons Program on “Foundations of Machine learning”, Spring 2017.
6. Co-Organizer, Workshop on “Learning, Algorithm Design, and Beyond Worst Case Analysis” within the Simons Program on “Algorithms and Uncertainty”, Fall 2016.
7. Co-organizer, Dagstuhl Workshop on “Foundations of Unsupervised Learning”, Fall 2016.
8. Co-organizer, Lorentz Center Workshop on “Theoretical Foundations for Learning from Easy Data”, Fall 2016.
9. Co-Organizer, ICML Workshop on “Active Learning: Towards Bridging Theory and Practice”, 2015.
10. Co-Organizer, IMA Workshop on “Convexity and Optimization”, February 2015.
11. Co-Organizer, Dagstuhl Workshop on “Analysis of Algorithms Beyond the Worst Case”, Fall 2014.
12. Co-Organizer, Campus wide Machine Learning Seminar Series at Georgia Tech, 2014.
13. Co-Organizer, ARC Workshop on “Modern Aspects of Submodularity” 2012.
14. Session organizer: “New Trends in Auction Design,” The 20th International Symposium of Mathematical Programming (ISMP), 2009.
15. Session organizer: “Recent Advances in Machine Learning”, The INFORMS Annual Meeting, 2009.
16. Co-Organizer, NIPS Workshop on “New Challenges in Theoretical Machine Learning: Learning with Data-dependent Concept Spaces”, 2008.

### **Fellowship Selection Committees**

1. Schmidt Science Fellows Selection Committee 2018, 2019.
2. ACM Heidelberg Forum Committee 2016, 2017, 2018.

### **General Program Committee Member**

1. Intl. Workshop on Similarity-Based Pattern Analysis and Recognition (**SIMBAD**) 2011 and 2013.
2. **ECML/PKDD** (Nectar Track) 2010 and 2013.
3. International Conference on Machine Learning (**ICML**) 2007, 2008, and 2012.
4. The 12th ACM Conference on Electronic Commerce (**EC**) 2011.
5. The 19th International World Wide Web Conference (**WWW**) 2010.

### **Journal Refereeing**

Journal of Machine Learning Research, Machine Learning Journal, SIAM Journal on Discrete Mathematics, Transactions on Algorithms, IEEE Transactions on Information Theory, Journal of Artificial Intelligence Research, Artificial Intelligence Journal, Neural Computation, Theoretical Computer Science

### **Conference Refereeing**

FOCS, STOC, SODA, COLT, NIPS, STACS, ESA, APPROX, SPAA, EC, WINE, WAOA, IPCO, WAW.

### Research Project Reviewer and Panelist

1. National Science Foundation. Panelist for Division of Computing and Communications Foundation, AF Career Panel, 2017.
2. National Science Foundation. Reviewer for Division of Information and Intelligent Systems. 2015.
3. National Science Foundation. Panelist for Division of Computing and Communications Foundation, AF Career Panel, 2014.
4. National Science Foundation. Panelist for Division of Computing and Communications Foundation, AF Panel, 2012.
5. Reviewer for the U.S.-Israel Binational Science Foundation. March 2011.
6. National Science Foundation. Panelist for Division of Computing and Communications Foundation, AF Career Panel, 2010.

### LIST OF PUBLICATIONS

#### Books and Chapters

1. **Center Based Clustering: A Foundational Perspective**, Pranjal Awasthi and Maria-Florina Balcan. Book Chapter in “Handbook of Cluster Analysis”, C. Hennig, M. Meila, F. Murtagh, and R. Rocci (Eds.), Chapman and Hall/CRC 2015.
2. **An Augmented PAC Model for Semi-Supervised Learning**, Maria-Florina Balcan and Avrim Blum. Book Chapter in “Semi-Supervised Learning”, Chapelle, O., Zien, A., and Schölkopf, B. (Eds.), MIT Press, 2006.
3. **Search and Knowledge Representation in Artificial Intelligence. Theory and Applications**, Florentina Hristea and Maria-Florina Balcan. University of Bucharest Publishing House, 2005 (in Romanian).

#### Refereed Journal Papers

4. **Submodular Functions: Structure, Learnability, and Optimizations**. Maria-Florina Balcan and Nick Harvey. *SIAM Journal of Computing*, Vol. 47, No. 3, pp. 703–754, 2018.
5. **Nash Equilibria in Perturbation-Stable Games**. Maria-Florina Balcan and Mark Braverman. *Theory of Computing*, Vol. 13, Article 13, 1–31, 2017.
6. **Scalable Influence Maximization for Multiple Products in Continuous-Time Diffusion Networks**. Nan Du, Yingyu Liang, Maria-Florina Balcan, Manuel Gomez-Rodriguez, Hongyuan Zha, and Le Song. *Journal of Machine Learning Research*, 18(2):1–45, 2017.
7. **Local Algorithms for Interactive Clustering**. Pranjal Awasthi, Maria-Florina Balcan, and Konstantin Voevodski. *Journal of Machine Learning Research*, 18(3):1–35, 2017.
8. **The Power of Localization for Learning Linear Separators with Noise**. Pranjal Awasthi, Maria-Florina Balcan, and Phil Long. *Journal of the ACM*, Volume 63, Issue 6, 2017.
9. **Clustering under Perturbation Resilience**, Maria-Florina Balcan and Yingyu Liang. *SIAM Journal of Computing*, Vol. 45, No 1, 102 – 155, 2016.
10. **Statistical Active Learning Algorithms for Noise Tolerance and Differential Privacy**. Maria-Florina Balcan and Vitaly Feldman. *Algorithmica* 72(1), 282 – 315, 2015. Special Issue on Learning Theory. (**Invited**)
11. **Robust Hierarchical Clustering**, Maria-Florina Balcan, Yingyu Liang, and Pramod Gupta. *Journal of Machine Learning Research*, 15, 3831– 3871, 2014.
12. **Near-Optimality in Covering Games by Exposing Global Information**, Maria-Florina Balcan, Sara Krehbiel, Georgios Piliouras, and Jinwoo Shin. *ACM Transactions on Economics and Computation*, Volume 2, Issue 4, 13:1 – 13:22, 2014.
13. **Clustering under Approximation Stability**, Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. *Journal of the ACM*, Volume 60, Issue 2, 8:1- 8:34, 2013.
14. **Surpassing the Price of Anarchy: Leading Dynamics to Good Behavior**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. *SIAM Journal on Computing*, 42(1), 2310 – 264, 2013.

15. **The Price of Uncertainty**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. *ACM Transactions on Economics and Computation*, Volume 1, Number 3, 2013.
16. **Active Clustering of Biological Sequences**, Konstantin Voevodski, Maria-Florina Balcan, Heiko Röglin, Shang-Hua Teng, and Yu Xia. *Journal of Machine Learning Research*, 13, 203 – 225, 2012.
17. **On Nash-Equilibria of Approximation-Stable Games**, Pranjal Awasthi, Maria-Florina Balcan, Avrim Blum, Or Sheffet, and Santosh Vempala. *Current Science Journal*, Volume 103, Issue 9, 1014-1020, 2012. **(Invited)**
18. **The True Sample Complexity of Active Learning**, Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman. *Machine Learning Journal*, 80(2–3), 111 – 139, 2010. Special issue for COLT 2008. **(Invited)**
19. **A Discriminative Framework for Semi-Supervised Learning**, Maria-Florina Balcan and Avrim Blum. *Journal of the ACM*, Volume 57, Issue 3, 2010.
20. **Agnostic Active Learning**, Maria-Florina Balcan, Alina Beygelzimer, and John Langford. *Journal of Computer and System Sciences*, 75 (1): 78–89, 2009. Special Issue on Learning Theory. **(Invited)**
21. **Reducing Mechanism Design to Algorithm Design via Machine Learning**, Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. *Journal of Computer and System Sciences*, 74(8): 1245 – 1270, 2008. Special Issue on Learning Theory. **(Invited)**
22. **On a Theory of Learning with Similarity Functions**, Maria-Florina Balcan, Avrim Blum, and Nathan Srebro. *Machine Learning Journal*, 72(1–2): 89–112, 2008. Special issue of *Machine Learning Journal* for COLT 2007. **(Invited)**
23. **Robust Reductions from Ranking to Classification**, Maria-Florina Balcan, Nikhil Bansal, Alina Beygelzimer Don Coppersmith, John Langford, and Gregory B. Sorkin. *Machine Learning Journal*, 72 (1 – 2): 139 – 153, 2008. Special issue of *Machine Learning Journal* for COLT 2007. **(Invited)**
24. **Approximation Algorithms and Online Mechanisms for Item Pricing**, Maria-Florina Balcan and Avrim Blum. *Theory of Computing*, 3/9 : 179 – 195, 2007.
25. **Kernels as Features: On Kernels, Margins, and Low-dimensional Mappings**, Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. *Machine Learning Journal*, 65(1):79 – 94, 2006.
26. **Approaches to Handwritten/Machine Printed Discrimination Problem**, Maria-Florina Popa and Doru-Cristian Balcan. University of Bucharest Annals, Computer Science, 2000.

#### Refereed Conference Papers

27. **Testing Matrix Rank Optimally**. Maria-Florina Balcan, Yi Li, David Woodruff, and Hongyang Zhang. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2019.
28. **Data Driven Clustering for Parametrized Lloyd’s Families**. Maria-Florina Balcan, Travis Dick, and Colin White. Proceedings of the 32<sup>nd</sup> Annual Conference on Neural Information Processing (NIPS) 2018.
29. **Dispersion for Data-driven Algorithm Design, Online Learning, and Private Optimization**. Maria-Florina Balcan, Travis Dick, and Ellen Vitercik. Proceedings of the 59<sup>th</sup> Annual Symposium on Foundations of Computer Science (FOCS) 2018.
30. **A General Theory of Sample Complexity for Multi-Item Profit Maximization**. Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. The 19<sup>th</sup> ACM Conference on Economics and Computation (ACM-EC) 2018.
31. **Learning to Branch**. Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Proceedings of the 34<sup>th</sup> International Conference on Machine Learning (ICML) 2018.
32. **Diversified Strategies in Games**. Maria-Florina Balcan, Avrim Blum, and Shang Tse-Chen. International Conference on Autonomous Systems (AAMAS) 2018.
33. **Matrix Completion and Related Problems via Strong Duality**. Maria-Florina Balcan, Yingyu Liang, David Woodruff, and Hongyang Zhang. Proceedings of the 9<sup>th</sup> Innovations in Theoretical Computer Science Conference (ITCS) 2018.

34. **Learning-Theoretic Foundations of Algorithm Configuration for Combinatorial Partitioning Problems.** Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, and Colin White. Proceedings of the 30<sup>th</sup> Annual Conference on Learning Theory (COLT) 2017.
35. **Sample and Computationally Efficient Learning Algorithms under s-Concave Distributions.** Maria-Florina Balcan and Hongyang Zhang. Proceedings of the 31<sup>st</sup> Annual Conference On Neural Information Processing Systems (NIPS) 2017.
36. **Risk Bounds for Transferring Representations With and Without Fine-Tuning.** Daniel McNamara and Maria-Florina Balcan. Proceedings of the 33<sup>rd</sup> International Conference on Machine Learning (ICML) 2017.
37. **Differentially Private Clustering in High-Dimensional Euclidean Spaces.** Maria-Florina Balcan, Travis Dick, Yingyu Liang, Wenlong Mou, and Hongyang Zahang. Proceedings of the 33<sup>rd</sup> International Conference on Machine Learning (ICML) 2017.
38. **Lifelong Learning in Costly Feature Spaces.** Maria-Florina Balcan, Avrim Blum, and Vaishnavh Nagarajan. Algorithmic Learning Theory Conference (ALT) 2017.
39. **Data Driven Resource Allocation for Distributed Learning.** Travis Dick, Mu Li, Venkata Krishna Pillutla, Colin White, Maria-Florina Balcan, and Alex Smola. Proceedings of 20<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2017.
40. **Label Efficient Learning by Exploiting Multi-class Output Codes.** Maria-Florina Balcan, Travis Dick, and Yishay Mansour. The Thirty-First AAAI Conference on Artificial Intelligence (AAAI) 2017.
41. **Noise-Tolerant Life-Long Matrix Completion via Adaptive Sampling.** Maria-Florina Balcan and Hongyang Zhang. Proceedings of the 30<sup>th</sup> Annual Conference On Neural Information Processing Systems (NIPS) 2016.
42. **Sample Complexity of Automated Mechanism Design.** Maria-Florina Balcan, Ellen Vitercik, and Tuomas Sandholm. Proceedings of the 30<sup>th</sup> Annual Conference On Neural Information Processing Systems (NIPS) 2016.
43. **Distributed Kernel Principal Component Analysis,** Maria-Florina Balcan, Yingyu Liang, Le Song, David Woodruff, and Bo Xie. Proceedings of the 20<sup>th</sup> ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2016.
44. **Learning Combinatorial Functions from Pairwise Comparisons.** Maria-Florina Balcan, Ellen Vitercik and Colin White. Proceedings of the 29<sup>th</sup> Annual Conference on Learning Theory (COLT) 2016.
45. **Efficient Algorithms for Learning and 1-bit Compressed Sensing under Asymmetric Noise,** Maria-Florina Balcan, Pranjal Awasthi, Nika Haghtalab, and Hongyang Zhang. Proceedings of the 29<sup>th</sup> Annual Conference on Learning Theory (COLT) 2016.
46. **An Improved Gap-Dependency Analysis of the Noisy Power Method.** Maria-Florina Balcan, Simon Du, Yining Wang, and Adams Wei Yu. Proceedings of the 29<sup>th</sup> Annual Conference on Learning Theory (COLT) 2016.
47. **k-Center Clustering under Perturbation Resilience.** Maria-Florina Balcan, Nika Haghtalab, and Colin White. The 43<sup>rd</sup> International Colloquium on Automata, Languages and Programming (ICALP) 2016.
48. **Communication Efficient Distributed Agnostic Boosting.** Shang-Tse Chen, Maria-Florina Balcan, and Polo Chau. Proceedings of 19<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2016.
49. **Active Learning Algorithms for Graphical Model Selection.** Gautam Dasarathy, Aarti Singh, Maria-Florina Balcan, and Jong Park. Proceedings of 19<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2016.
50. **Learning Cooperative Games,** Maria-Florina Balcan, Ariel Procaccia, and Yair Zick. Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI) 2015.
51. **Efficient Learning of Linear Separators under Bounded Noise,** Pranjal Awasthi, Maria-Florina Balcan, Nika Haghtalab, and Ruth Urner. Proceedings of the 28<sup>th</sup> Annual Conference on Learning Theory (COLT) 2015.

52. **Distributed Frank-Wolfe Algorithm: A Unified Framework for Communication-Efficient Sparse Learning**, Aurelien Bellet, Yingyu Liang, Alireza Bagheri Garakani, Maria-Florina Balcan, and Fei Sha. SIAM International Conference on Data Mining 2015.
53. **Efficient Representations for Life-Long Learning and Autoencoding**, Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Proceedings of the 28th Annual Conference on Learning Theory (COLT) 2015.
54. **Commitment Without Regrets: Online Learning in Stackelberg Security Games**. Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel Procaccia. The Sixteenth ACM Conference on Economics and Computation (EC) 2015.
55. **Scalable Kernel Methods via Doubly Stochastic Gradients**, Bo Dai, Bo Xie, Niao He, Yingyu Liang, Anant Raj, Maria-Florina Balcan, and Le Song. Proceedings of the Twenty-Eighth Annual Conference On Neural Information Processing Systems (NIPS) 2014.
56. **Improved Distributed Principal Component Analysis**, Maria-Florina Balcan, Vandana Kanchanapally, Yingyu Liang, and David Woodruff. Proceedings of the Twenty-Eighth Annual Conference On Neural Information Processing Systems (NIPS) 2014.
57. **Learning Time-Varying Coverage Functions**, Nan Du, Yingyu Liang, Maria-Florina Balcan, and Le Song. Proceedings of the Twenty-Eighth Annual Conference On Neural Information Processing Systems (NIPS) 2014.
58. **Learning Economic Parameters from Revealed Preferences**, Maria-Florina Balcan, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V. Vazirani. Proceedings of the 10th Conference on Web and Internet Economics (WINE) 2014.
59. **The Power of Localization for Learning Linear Separators with Noise**, Pranjal Awasthi, Maria-Florina Balcan, and Phil Long. Proceedings of the 46<sup>th</sup> ACM Symposium on Theory of Computing (STOC) 2014.
60. **Local Algorithms for Interactive Clustering**, Pranjal Awasthi, Maria-Florina Balcan, and Konstantin Voevodski. Proceedings of the 31<sup>st</sup> International Conference on Machine Learning (ICML) 2014.
61. **A New Perspective on Learning Linear Separators with Large  $L_p, L_q$  Margins**, Maria-Florina Balcan and Chris Berling. Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics (AISTATS) 2014.
62. **Influence Function Learning in Information Diffusion Networks**, Nan Du, Yingyu Liang, Maria-Florina Balcan, and Le Song. Proceedings of the 31<sup>st</sup> International Conference on Machine Learning (ICML) 2014.
63. **Active Learning and Best-Response Dynamics**, Maria-Florina Balcan, Chris Berling, Avrim Blum, Emma Cohen, Kaushik Patnaik, and Le Song. Proceedings of the Twenty-Eighth Annual Conference On Neural Information Processing Systems (NIPS) 2014.
64. **Statistical Active Learning Algorithms**, Maria-Florina Balcan and Vitaly Feldman. Proceedings of the Twenty-Seventh Annual Conference on Neural Information Processing Systems (NIPS) 2013.
65. **Distributed k-Means and k-Median Clustering on General Topologies**, Maria-Florina Balcan, Steven Ehrlich, and Yingyu Liang. Proceedings of the Twenty-Seventh Annual Conference on Neural Information Processing Systems (NIPS) 2013.
66. **Active and Passive Learning of Linear Separators under Log-concave Distributions**, Maria-Florina Balcan and Phil Long. Proceedings of the 26<sup>th</sup> Annual Conference on Learning Theory (COLT) 2013. Also invited at the Also invited at the 51st Annual Allerton Conference on Communication, Control, and Computing.
67. **Efficient Semi-supervised and Active Learning of Disjunctions**, Maria-Florina Balcan, Chris Berling, Steven Ehrlich, and Yingyu Liang. Proceedings of the 30th International Conference on Machine Learning (ICML) 2013.

68. **Exploiting Ontology Structures and Unlabeled Data for Learning**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Proceedings of the 30th International Conference on Machine Learning (ICML) 2013.
69. **Finding Endogenously Formed Communities**, Maria-Florina Balcan, Christian Borgs, Mark Braverman, Jennifer Chayes, and Shang-Hua Teng. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2013.
70. **Modeling and Detecting Community Hierarchies**, Maria-Florina Balcan and Yingyu Liang. Proceedings of the 2nd International Workshop on Similarity-Based Pattern Analysis and Recognition (SIMBAD) 2013.
71. **Robust Interactive Learning**, Maria-Florina Balcan and Steve Hanneke. Proceedings of the 25th Conference on Learning Theory (COLT) 2012.
72. **Learning Valuation Functions**, Maria-Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. Proceedings of the 25th Conference on Learning Theory (COLT) 2012.
73. **Distributed Learning, Communication Complexity, and Privacy**, Maria-Florina Balcan, Avrim Blum, Shai Fine, and Yishay Mansour. Proceedings of the 25th Conference on Learning Theory (COLT) 2012.  
**Winner of the Runner Up Best Paper Award.**
74. **Clustering under Perturbation Resilience**, Maria-Florina Balcan and Yingyu Liang. Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP) 2012.
75. **Active Property Testing**, Maria-Florina Balcan, Eric Blais, Avrim Blum, and Liu Yang. Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS) 2012.
76. **Near Optimality in Covering and Packing Games by Exposing Global Information**, Maria-Florina Balcan, Sara Krehbiel, Georgios Piliouras, and Jinwoo Shin. Proceedings of the 51st IEEE Conference on Decision and Control (CDC) 2012.
77. **Learning Submodular Functions**, Maria-Florina Balcan and Nicholas J. Harvey. Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC) 2011.  
Also **invited to NECTAR Track for significant machine learning results** at The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), 2012.
78. **Combining Self Training and Active Learning for Video Segmentation**, Alireza Fathi, Maria-Florina Balcan, Xiaofeng Ren, and James M Rehg. Proceedings of the 22nd British Machine Vision Conference (BMVC) 2011.
79. **The Snowball Effect of Uncertainty in Potential Games**, Maria-Florina Balcan, Florin Constantin, and Steven Ehrlich. Proceedings of the 7th Workshop on Internet and Network Economics (WINE) 2011.
80. **Min-Sum Clustering of Protein Sequences with Limited Distance Information**, Konstantin Voevodski, Maria-Florina Balcan, Heiko Röglin, Shang-Hua Teng, and Yu Xia. Proceedings of the 1st International Workshop on Similarity-Based Pattern Analysis and Recognition (SIMBAD) 2011.
81. **Game-theoretic Couplings and Applications**, Maria-Florina Balcan, Florin Constantin, Georgios Piliouras, and Jeff Shamma. Proceedings of the 50th IEEE Conference on Decision and Control (CDC) 2011.
82. **Robust Hierarchical Clustering**, Maria-Florina Balcan, Yingyu Liang, and Pramod Gupta. Proceedings of the 23rd Annual Conference on Learning Theory (COLT) 2010.
83. **Efficient Clustering with Limited Distance Information**, Konstantin Voevodski, Maria-Florina Balcan, Heiko Röglin, Shang-Hua Teng, and Yu Xia. Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI) 2010.
84. **Sequential Item Pricing for Unlimited Supply**, Maria-Florina Balcan and Florin Constantin. Proceedings of the 6th Workshop on Internet and Network Economics (WINE) 2010.



85. **On the Equilibria of Asynchronous Games**, Aaron Roth, Maria-Florina Balcan, Adam Kalai, and Yishay Mansour. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2010.
86. **On Nash Equilibria of Approximation-Stable Games**. Pranjali Awasthi, Maria Florina Balcan, Avrim Blum, Or Sheffet, and Santosh Vempala. Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT) 2010.
87. **Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Proceedings of the Innovations in Computer Science (ICS) 2010.
88. **Finding Low Error Clusterings**, Maria-Florina Balcan and Mark Braverman. Proceedings of the 22nd Annual Conference on Learning Theory (COLT) 2009.
89. **Agnostic Clustering**, Maria-Florina Balcan, Heiko Röglin, and Shang-Hua Teng. Proceedings of the 20th International Conference on Algorithmic Learning Theory (ALT) 2009.
90. **The Price of Uncertainty**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Proceedings of the Tenth ACM Conference on Electronic Commerce (EC) 2009.
91. **Approximate Clustering without the Approximation**, Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2009.
92. **Improved Equilibria via Public Service Advertising**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2009.
93. **The True Sample Complexity of Active Learning**, Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman. Proceedings of the 21st Annual Conference on Learning Theory (COLT) 2008.  
**Winner of the Mark Fulk Best Student Paper Award.**
94. **A Discriminative Framework for Clustering via Similarity Functions**, Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Proceedings of the 40th ACM Symposium on Theory of Computing (STOC) 2008.
95. **Improved Guarantees for Learning via Similarity Functions**, Maria-Florina Balcan, Avrim Blum, and Nathan Srebro. Proceedings of the 21st Annual Conference on Learning Theory (COLT) 2008.
96. **Item Pricing for Revenue Maximization**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Proceedings of the Ninth ACM Conference on Electronic Commerce (EC) 2008.
97. **Clustering with Interactive Feedback**, Maria-Florina Balcan and Avrim Blum. Proceedings of the 19th International Conference on Algorithmic Learning Theory (ALT) 2008.
98. **Margin Based Active Learning**, Maria-Florina Balcan, Andrei Broder, and Tong Zhang. Proceedings of the 20th Annual Conference on Learning Theory (COLT) 2007.
99. **Robust Reductions from Ranking to Classification**, Maria-Florina Balcan, Nikhil Bansal, Alina Beygelzimer, Don Coppersmith, John Langford, and Gregory B. Sorkin. Proceedings of the 20th Annual Conference on Learning Theory (COLT) 2007.
100. **A Theory of Loss-leaders: Making Money by Pricing below Cost**, Maria-Florina Balcan, Avrim Blum, Hubert Chan, and MohammadTaghi Hajiaghayi. Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE) 2007.
101. **Agnostic Active Learning**, Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Proceedings of the 23rd International Conference on Machine Learning (ICML) 2006.
102. **Approximation Algorithms and Online Mechanisms for Item Pricing**, Maria-Florina Balcan and Avrim Blum. Proceedings of the Seventh ACM Conference on Electronic Commerce (EC) 2006.
103. **On a Theory of Learning with Similarity Functions**, Maria-Florina Balcan and Avrim Blum. Proceedings of the 23rd International Conference on Machine Learning (ICML), 2006.
104. **Mechanism Design via Machine Learning**, Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS) 2005.
105. **A PAC-style Model for Learning from Labeled and Unlabeled Data**, Maria-Florina Balcan and Avrim Blum. Proceedings of the 18th Annual Conference on Learning Theory (COLT) 2005.

106. **Kernels as Features: On Kernels, Margins, and Low-dimensional Mappings**, Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Proceedings of the 15th International Conference on Algorithmic Learning Theory (ALT) 2004.
107. **Co-Training and Expansion: Towards Bridging Theory and Practice**, Maria-Florina Balcan, Avrim Blum, and Ke Yang. Proceedings of the Eighteenth Annual Conference on Neural Information Processing Systems (NIPS) 2004.
108. **Handwritten Text Localization in Skewed Documents**, Ergina Kavallieratou, Doru-Cristian Balcan, Maria-Florina Popa, and Nikos Fakotakis. IEEE International Conference on Image Processing (ICIP) 2001.

#### Short Surveys

109. **Active Learning**, Maria-Florina Balcan and Ruth Uerner. Encyclopedia of Algorithms, 2015.
110. **Leading Dynamics to Good Behavior**, Maria-Florina Balcan. SIGecom Exchanges, Volume 10, 2011. (Invited)
111. **Item Pricing for Revenue Maximization**, Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. SIGecom Exchanges, Volume 7.3, 2008. (Invited)
112. **Mechanism Design, Machine Learning, and Pricing Problems**, Maria-Florina Balcan and Avrim Blum. SIGecom Exchanges, Volume 7.1, 2007. (Invited)

#### Open Problems

113. **Better Guarantees for Sparsest Cut Clustering (Open Problem)**, Maria-Florina Balcan. Proceedings of the 22nd Annual Conference on Learning Theory (COLT) Open Problems, 2009.
114. **Open Problems in Efficient Semi-Supervised PAC Learning**, Avrim Blum and Maria-Florina Balcan. Proceedings of 20th Annual Conference on Learning Theory (COLT) Open Problems, 2007.

#### Workshop Papers

115. **Envy-Free Classification**. Maria Florina Balcan, Travis Dick, , Ritesh Noothigattu, Ariel Procaccia. NIPS 2018 Workshop on Ethical, Social, and Governance Issues in AI.
116. **Dispersion for Private Optimization of Piecewise Lipschitz Functions**. Maria Florina Balcan, Travis Dick, and Ellen Vitercik. ICML 2018 Workshop on Privacy in Machine Learning and Artificial Intelligence.
117. **Sample Complexity of Multi-Item Profit Maximization**. Maria Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. EC 2017 Workshop on Algorithmic Game Theory and Data Science.
118. **Differentially Private Algorithm Configuration**. Maria Florina Balcan, Travis Dick, and Ellen Vitercik. ICML 2017 Workshop on Private Secure Machine Learning.
119. **Data Resource Allocation for Distributed Learning**. Travis Dick, Mu Li, Venkata Krishna Pillutla, Colin White, Maria Florina Balcan, and Alex Smola. AAAI 2017 Workshop on Distributed Machine Learning.
120. **On the Geometry of Output Code Multi-class Learning**. Maria-Florina Balcan, Travis Dick, and Yishay Mansour. ICML 2016 Workshop on Data Efficient Machine Learning.
121. **Learning and 1-bit Compressed Sensing under Asymmetric Noise**. Pranjali Awasthi, Maria-Florina Balcan, Nika Haghtalab. ICML 2016 Workshop on Advances in Non-Convex Analysis and Optimization.
122. **Learning Cooperative Games**. Maria-Florina Balcan, Ariel Procaccia, and Yair Zick. The Sixth Workshop on Cooperative Games in Multiagent Systems, CoopMAS-2015.
123. **Distributed Frank-Wolfe Algorithm: A Unified Framework for Communication-Efficient Sparse Learning**. Aurelien Bellet, Yingyu Liang, Alireza Bagheri Garakani, Maria-Florina Balcan and Fei Sha. ICML 2014 workshop on New Learning Frameworks and Models for Big Data.
124. **Distributed PCA and k-Means Clustering**. Yingyu Liang, Maria-Florina Balcan, and Vandana Kanchanapally. NIPS Workshop on Learning Faster from Easy Data, 2013.
125. **On Learning Linear Separators with large  $L_\infty L_1$  Margins**. Maria-Florina Balcan and Chris Berlind. NIPS Workshop on Learning Faster from Easy Data, 2013.

126. **Clustering k-median Perturbation Resilient Instances.** Maria-Florina Balcan and Yingyu Liang. NIPS Workshop on Learning Faster from Easy Data, 2013.
127. **Learning Symmetric Non-monotone Submodular Functions.** Maria-Florina Balcan, Nicholas J. Harvey, and Satoru Iwata. NIPS Workshop on Discrete Optimization in Machine Learning, 2012.
128. **The Weighted Majority Algorithm does not Converge in Nearly Zero-sum Games,** Maria-Florina Balcan, Florin Constantin, and Ruta Mehta. ICML 2012 Workshop on Markets, Mechanisms, and Multi-Agent Models.
129. **Clustering Protein Sequences Given the Approximation Stability of the Min-Sum Objective Function,** Konstantin Voevodski, Maria-Florina Balcan, Heiko Röglin, Shang-Hua Teng, and Yu Xia. The Snowbird Learning Workshop, 2011.
130. **Weighted Neighborhood Linkage,** Maria-Florina Balcan and Pramod Gupta. NIPS Workshop on Robust Statistical Learning, 2010.
131. **Learning with Multiple Similarity Functions,** Maria-Florina Balcan, Avrim Blum, and Nathan Srebro. NIPS Workshop on Kernel Learning: Automatic Selection of Optimal Kernels, 2008.
132. **Similarity-Based Theoretical Foundation for Sparse Parzen Window Prediction,** Maria-Florina Balcan, Avrim Blum, and Nathan Srebro. ICML/UAI/COLT Workshop on Sparse Optimization and Variable Selection, 2008.
133. **Asymptotic Active Learning,** Maria-Florina Balcan, Eyal-Even Dar, Steve Hanneke, Michael Kearns, Yishay Mansour, and Jennifer Wortman. NIPS Workshop on the Principles of Learning Problem Design, 2007.
134. **Sponsored Search Auctions Design via Machine Learning,** Maria-Florina Balcan, Avrim Blum, Jason D. Hartline and Yishay Mansour, ACM-EC Workshop on Sponsored Search Auctions, 2005.
135. **Person Identification in Webcam Images: An Application of Semi-Supervised Learning,** Maria-Florina Balcan, Avrim Blum, Pakyan Choi, John Lafferty, Brian Pantano, Mugizi Robert Rwebangira, and Xiaojin Zhu. ICML Workshop on Learning with Partially Classified Training Data, 2005.

#### Technical Reports

136. **Random Sampling Auctions for Limited Supply.** Maria-Florina Balcan, Nikhil Devanur, Jason Hartline, and Kunal Talwar. Technical Report, CMU-CS-07-154, 2007.

#### Manuscripts

137. **Envy-Free Classification.** Maria-Florina Balcan, Travis Dick, Ritesh Noothigattu, Ariel Procaccia, 2018.
138. **Robust Communication-Optimal Distributed Clustering Algorithms under Approximation Stability.** Ainesh Bakshi, Maria-Florina Balcan, Colin White, and David Woodruff, 2018.

#### Patents

139. **Detecting Overlapping Clusters,** Maria-Florina Balcan, Christian Borgs, Mark Braverman, Jennifer Chayes, and Shang-Hua Teng. US8949237 B2, 2015.

## CONTRIBUTIONS TO EDUCATION

#### Courses Taught

- **Introduction to Machine Learning.** Spring 2019 at Carnegie Mellon University.
- **Advanced Introduction to Machine Learning.** Fall 2018 at Carnegie Mellon University.
- **Introduction to Machine Learning.** Spring 2018 at Carnegie Mellon University.
- **Advanced Introduction to Machine Learning.** Fall 2017 at Carnegie Mellon University.
- **Machine Learning.** Fall 2016 at Carnegie Mellon University, co-teaching with Matt Gormley.
- **Machine Learning.** Spring 2016 at Carnegie Mellon University, co-teaching with William Cohen.

- **Foundations of Machine Learning and Data Science.** Fall 2015 at Carnegie Mellon University, co-teaching with Avrim Blum.
- **Machine Learning.** Spring 2015 at Carnegie Mellon University, co-teaching with Tom Mitchell.
- **Machine Learning Theory.** Fall 2013, Fall 2011, and Spring 2010 (at Georgia Tech).
- **Analysis and Design of Algorithms.** Spring 2014, Spring 2013, Fall 2012, and Spring 2011 (at Georgia Tech).
- **Connections between Learning, Game Theory, and Optimization.** Fall 2010 (at Georgia Tech).
- **Machine Learning Theory.** Spring 2007 at Carnegie Mellon University, co-teaching with Avrim Blum.
- **Artificial Intelligence.** Spring 2002 and Spring 2001 (at University of Bucharest).
- **Introduction to Computer Science.** Fall 2001, Spring 2001, and Fall 2000 (at University of Bucharest).

## Individual Student Guidance and Mentoring

### Postdocs mentored

1. Ruth Urner: October 2013 – December 2014.  
Current Position: Assistant Professor, York University.
2. Florin Constantin: October 2009 - September 2011.  
Current Position: Research Engineer at Google Inc.

### Current PhD Students

1. Dravyansh Sharma. Carnegie Mellon University, January 2019 – present.
2. Mikhail Khodak. Carnegie Mellon University, Sept. 2018 – present (co-advised with Ameet Talwalkar).
3. Ellen Viterick. Carnegie Mellon University, Sept. 2015 – present (co-advised with Tuomas Sandholm).  
National Science Foundation Research Fellowship Winner, 2016–2019.  
Microsoft Research Women Fellowship Winner.
4. Hongyang Zhang. Carnegie Mellon University, Sept. 2015 – present (co-advised with David Woodruff).
5. Travis Dick: Carnegie Mellon University, October 2014 – present.
6. Shang-Tse Chen: Georgia Tech, August 2013– present (co-advised with Polo Chau).  
IBM PhD Fellowship Winner, 2018.

### Graduated PhD Students

1. Colin White: Carnegie Mellon University, October 2014 – October 2018.  
National Defense Science and Engineering Fellowship Winner, 2015–2018.  
PhD Thesis: New Aspects of Beyond Worst-Case Analysis.
2. Steven Ehrlich: August 2010 – July 2016 (co-advised with Jeff Shamma).  
PhD Thesis: Algorithmic, Game Theoretic, and Learning Theoretic Aspects of Distributed Optimization.
3. Chris Berlind: August 2011 – July 2015.  
PhD Thesis: New Insights on the Power of Active Learning.  
Current position: Co-Founder and CTO, Oncora Medical
4. Yingyu Liang: August 2010 – July 2014.  
PhD Thesis: Modern Aspects of Unsupervised Learning.  
Current position: Assistant Professor, Computer Science Department, University of Wisconsin-Madison.

### Graduated MS Students

1. Krishna Pillutla: CSD Master student, Carnegie Mellon University, 2015 -- 2016.  
MS Thesis: Data Driven Resource Allocation for Distributed Machine Learning.  
Current position: PhD student at University of Washington.
2. Pramod Gupta: Master student, Georgia Tech, 2009 – 2011.  
MS Thesis: Robust Clustering Algorithms.

### **Other Students Supervised**

1. Manuel Lang: visiting MS student from the Karlsruhe Institute of Technology (through the interACT program). Summer and Fall 2018.
2. Yue Wu: visiting undergraduate student from Peking University. Summer 2018.
3. Mengxiao Zhang: visiting undergraduate student from Peking University. Summer 2017.
4. Wenlong Mou: visiting undergraduate student from Peking University. Summer 2016.
5. Daniel MCNamara: visiting PhD student (with a Fulbright Postgraduate Scholarship) from Australian National University and Data 61. Fall 2016 – Spring 2017.
6. Yu Wang: undergraduate student, Carnegie Mellon University, Fall 2014.
7. Pranjal Awasthi: visiting PhD student, Summer 2012.
8. Ruta Mehta: visiting student, Summer 2011.
9. Jacob Abernethy: summer intern at MSR-NE, Summer 2009.

### **Member of Ph.D. Examining Committees**

1. Nika Haghtalab, Computer Science Department, CMU. Defended in Summer 2018.  
Principal Advisors: Avrim Blum and Ariel Procaccia.
2. Hassan Ashtiani, School of Computer Science, University of Waterloo. Defended in Spring 2018.  
Principal Advisor: Shai Ben David.
3. Shiva Kaul, Computer Science Department, CMU. Proposed in Summer 2016.  
Principal Advisor: Geoff Gordon.
4. Irina Nicolae, Jean Monnet University of Saint-Etienne. Defended in Fall 2016.  
Principal Advisors: Eric Gaussier and Marc Sebban.
5. Dougal Sutherland, Robotics Institute, CMU. Defended in Fall 2016.  
Principal Advisor: Jeff Schneider.
6. Nisarg Shah, Computer Science Department, CMU. Defended in Summer 2016.  
Principal Advisor: Ariel Procaccia.
7. Akshay Krishnamurthy, Computer Science Department, CMU. Defended in Spring 2015.  
Principal Advisor: Aarti Singh.
8. Aurele Balavoine, School of Electrical and Computer Engineering, Georgia Tech. Defended in Spring 2014.  
Principal Advisors: Justin Romberg and Chris Rozell.
9. Ying Xiao, College of Computing (SCS), Georgia Tech. Defended in Summer 2014  
Principal Advisor: Santosh Vempala.
10. Nicolas Dubeout, School of Electrical and Computer Engineering, Georgia Tech. Defended in Spring 2014.  
Principal Advisor: Jeff Shamma.
11. Niao He, School of Industrial and Systems Engineering (ISYE), Georgia Tech. Defended in Fall 2015.  
Principal Advisor: Arkadi Nemirovski.
12. Krishnakumar Balasubramanian, College of Computing (CSE), Georgia Tech. Defended in Summer 2014.  
Principal Advisor: Guy Lebanon.
13. Hua Ouyang, College of Computing (CSE), Georgia Tech. Defended in Spring 2013.  
Principal Advisor: Alexander Gray.
14. Ravi Sastry, College of Computing (CSE), Georgia Tech. Defended in Fall 2013.  
Principal Advisor: Alexander Gray.
15. Stas Minsker, School of Mathematics, Georgia Tech. Defended in Spring 2012.  
Principal Advisor: Vladimir Koltchinskii.
16. Amanda Pascoe Streib, School of Mathematics, Georgia Tech. Defended in Spring 2012.  
Principal Advisor: Dana Randall.
17. Pushkar Tripathi, College of Computing (ACO program), Georgia Tech. Defended in Spring 2012.  
Principal Advisor: Vijay Vazirani.

18. Michael Fox, School of Electrical and Computer Engineering, Georgia Tech. Defended in Spring 2012.  
Principal Advisor: Jeff Shamma.
19. Yaxian Li, School of Industrial and Systems Engineering (ISYE), Georgia Tech. Defended in Spring 2012.  
Principal Advisor: George Nemhauser.
20. Nishant Mehta, College of Computing (CSE), Georgia Tech. Defended in Spring 2013.  
Principal Advisor: Alexander Gray.
21. Parikshit Ram, College of Computing (CSE), Georgia Tech. Defended in Spring 2013.  
Principal Advisor: Alexander Gray.
22. Liam Mac Dermid, College of Computing (IC), Georgia Tech. Defended in Spring 2013.  
Principal Advisor: Charles Isbell.
23. Shiva Kintali, College of Computing (ACO program), Georgia Tech. Defended in Summer 2011.  
Principal Advisor: Richard Lipton.
24. Lei Wang, College of Computing (ACO program), Georgia Tech (Reader). Defended in Summer 2011.  
Principal Advisor: Vijay Vazirani.
25. Chinmay Karande, College of Computing (ACO program), Georgia Tech (Reader). Defended in Fall 2010.  
Principal Advisor: Vijay Vazirani.
26. Luyi Gui, School of Industrial and Systems Engineering (ISYE), Georgia Tech. Defended in Spring 2013.  
Principal Advisor: Ozlem Ergun.

## INVITED TALKS

### Keynote and Distinguished Lectures

1. **Foundations of Data Driven Algorithm Design.** Keynote Talk at The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD). September 2019 (Forthcoming).
2. **Foundations of Data Driven Algorithm Design.** Plenary Talk at the 24th LIDS Student Conference, MIT, January 2019.
3. **Data Driven Algorithm Design.** USC Computer Science Distinguished Lecture Series. October 2018.
4. **Sample and Computationally Efficient Active Learning.** Information Theory and Applications Workshop (ITA). February 2018
5. **Foundations for Learning in the Age of Big Data.** 9<sup>th</sup> China Theory Week. August 2015.
6. **Learning Submodular Functions with Applications to Multi-Agent Systems.** The 14<sup>th</sup> International Conference on Autonomous Agents and Multiagent System (AAMAS). May 2015.
7. **Statistical Active Learning Algorithms with Connections to Privacy.** 7<sup>th</sup> Workshop for Women in Machine Learning. December 2013.
8. **New Theoretical Frameworks for Modern Learning Paradigms.** Carnegie Mellon University. SCS Distinguished Lecture. March 2010.

### Invited Tutorials and Summer School Lectures

1. **Machine Learning in Automated Mechanism Design for Pricing and Auctions.** Tutorial at 35<sup>th</sup> International Conference on Machine Learning (ICML) 2018.
2. **Generalization Bounds for Passive and Active Learning.** Transylvanian Machine Learning Summer School, Cluj Napoca. July 2018.
3. **Learning and Game Theory.** Summer School on Algorithmic Game Theory, Fudan University, Shanghai. July 2010

### Other Tutorials

1. **New Frontiers of Automated Mechanism Design for Pricing and Auctions.** Tutorial at 33<sup>rd</sup> AAAI Conference on Artificial Intelligence. January 2019.

### Invited Talks at Workshops

1. **Data Driven Algorithm Design.** IPAM Workshop on “New Architectures and Algorithms”. November 2018.
2. **Sample Complexity of Multi-Item Profit Maximization.** TTI-C Workshop on Learning in the Presence of Strategic Behavior. August 2018.
3. **Learning Combinatorial Functions.** The AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence. July 2018.
4. **Foundations of Data Driven Algorithm Design.** Foundations of Machine Learning Reunion, Simons Institute. June 2018.
5. **Foundations of Data Driven Combinatorial Algorithm Selection.** NIPS Workshop on Discrete Structures in Machine Learning. December 2017.
6. **Sample and Computationally Efficient Active Learning Algorithms.** NIPS Workshop on Learning from Limited Labeled Data. December 2017.
7. **Communication Efficient Algorithms for Distributed Machine Learning.** IMA Workshop on Resource Trade-offs: Computation, Communication, and Information. May 2016.
8. **Finding Endogenously Formed Communities.** 3<sup>rd</sup> International Workshop on Similarity-Based Pattern Analysis and Recognition. October 2015.
9. **Distributed Machine Learning.** DIMACS Workshop on Big Data Through the Lens of Sublinear Algorithms. August 2015.
10. **Distributed Machine Learning.** Workshop on Massively Parallel Computation at FCRC 2015. June 2015.
11. **Active Learning of Linear Separators with Noise.** 45h Symposium on the Interface Computing Science and Statistics. June 2015.
12. **Beyond Worst Case Analysis in Clustering: Clustering under Perturbation Resilience.** Learning Theory Workshop at the DALI Meeting. April 2015.
13. **Learning Influence Function from Information Propagation Traces.** Networks: Processes and Causality Workshop at the DALI Meeting. April 2015.
14. **Foundations for Learning in the Age of Big Data.** Workshop on Algorithmic Challenges in Machine Learning. January 2015.
15. **The Power of Localization for Learning with Noise.** Lorentz Workshop on Online Algorithms and Learning. November 2014.
16. **Modern Machine Learning: New Challenges and Connections.** Georgia Scientific Computing Symposium. February 2014.
17. **Finding Endogenously Formed Communities.** Eurandom Workshop on Networks with Community Structure. January 2014.
18. **Statistical Active Learning Algorithms with Connections to Privacy.** SIMONS Workshop on Big Data and Differential Privacy. December 2013.
19. **Distributed Machine Learning.** Simons Workshop on Parallel and Distributed Algorithms for Inference and Optimization. October 2013.
20. **Active and Passive Learning of Linear Separators.** 51st Annual Allerton Conference on Communication, Control, and Computing. October 2013.
21. **Interactive Machine Learning.** STOC'13 Workshop on New (Theoretical) Challenges in Machine Learning. June 2013.
22. **Active and Passive Learning of Linear Separators.** Information Theory and Applications Workshop. February 2013.
23. **Incorporating Unlabeled Data and Interaction in the Learning Process.** Workshop on Provable Bounds in Machine Learning. Princeton. August 2012.
24. **Learning Valuation Functions.** Innovations in Algorithmic Game Theory. Israel, May 2011.

25. **Beyond Worst-Case Analysis in Machine Learning.** Workshop on Beyond Worst-Case Analysis. Stanford, September 2011.
26. **Robust Hierarchical Clustering.** NIPS 2010 Workshop on Robust Statistical Learning. December 2010.
27. **Learning Submodular Functions.** SIAM Conference on Discrete Mathematics. June 2010.
28. **Learning with Similarity Functions.** ICML Workshop on Learning in Non-(geo)metric Spaces. June 2010.
29. **Learning Submodular Functions.** The 5th Bertinoro Workshop on Random(ized) Graphs and Algorithms. May 2010.
30. **Learning Submodular Functions.** NIPS 2009 Workshop on Discrete Optimization in Machine Learning. December 2009.
31. **The True Sample Complexity of Active Learning.** NIPS Workshop on Adaptive Sensing, Active Learning, and Experimental Design. December 2009.
32. **Approximate Clustering without the Approximation.** INFORMS. October 2009.
33. **The Dynamics of Equilibria.** INFORMS. October 2009.
34. **New Theoretical Frameworks for Modern Learning Paradigms.** Barriers in Computational Complexity Workshop, Princeton University. August 2009.
35. **Finding Low Error Clusterings.** Machine Learning Summer School, University of Chicago. June 2009.
36. **A Computational Theory of Clustering.** Information Theory and Applications Workshop. February 2009.
37. **Item Pricing for Revenue Maximization in Combinatorial Auctions.** Dagstuhl Workshop on Computational Social Systems and the Internet. July 2007.
38. **Mechanism Design, Machine Learning, and Pricing Problems.** INFORMS. November 2006.
39. **Mechanism Design, Machine Learning, and Pricing Problems.** Second Bertinoro Workshop on Algorithmic Game Theory. July 2006.

**University Colloquia and Seminars. Industry Presentations.**

40. **Data Driven Algorithm Design.** Caltech Seminar. October 2018.
41. **Foundations of Data Driven Algorithm Design.** TTI-C Seminar, August 2018.
42. **Distributed Machine Learning.** Trends in Optimization Seminar. University of Washington. May 2016.
43. **Learning Submodular Functions.** Operations Research Seminar. Carnegie Mellon University. January 2016.
44. **Distributed Machine Learning.** Harvard CS Colloquium. November 2015.
45. **Learning Submodular Functions.** University of Southern California, CS Colloquium. October 2015.
46. **Learning Submodular Functions.** Machine Learning External Seminar, Gatsby Unit. July 2015.
47. **Learning Submodular Functions.** UC Berkeley EECS Departmental Colloquium Distinguished Lecture Series. April 2015.
48. **Foundations for Learning in the Age of Big Data.** Theory Seminar, Princeton. April 2014.
49. **Foundations for Learning in the Age of Big Data.** Theory of Computation Colloquium, MIT. April 2014.
50. **Foundations for Learning in the Age of Big Data.** Theory Seminar, University of Columbia. March 2014.
51. **Foundations for Learning in the Age of Big Data.** Machine Learning Seminar, Courant Institute of Mathematical Sciences. March 2014.
52. **Modern Machine Learning: New Challenges and Connections.** SCS Special Seminar, Carnegie Mellon University. February 2014.
53. **Foundations for Learning in the Age of Big Data.** TTI-C Colloquium, Toyota Technological Institute at Chicago. January 2014.
54. **Learning Valuation Functions.** CS Theory Seminar, University of Chicago. January 2014.
55. **Foundations for Learning in the Age of Big Data.** ML Seminar, Carnegie Mellon University. November 2013.
56. **Active and Passive Learning of Linear Separators.** Probability Seminar, Georgia Institute of Technology. October 2013.
57. **Active and Passive Learning of Linear Separators.** CATS Seminar, University of Maryland. September 2013.
58. **Learning Valuation Functions.** PRiML Seminar, University of Pennsylvania. April 2013.
59. **Learning Valuation Functions.** Harvard University, Economics and Computer Science Seminar. January 2013.



60. **Learning Valuation Functions.** Microsoft Research New England. August 2011.
61. **Learning Valuation Functions.** LearningTheory@FoCM'11. July 2011.
62. **Learning Submodular Functions.** Princeton University Theory Seminar. April 2011.
63. **Learning Submodular Functions.** The Annual Event of the ARC Center Georgia Institute of Technology. April 2011.
64. **Approximate Clustering without the Approximation.** Microsoft Research Redmond. July 2010.
65. **Approximate Clustering without the Approximation.** IBM Research Yorktown. December 2009.
66. **Approximate Clustering without the Approximation.** Combinatorics Seminar, Georgia Institute of Technology. October 2009.
67. **New Theoretical Frameworks for Modern Learning Paradigms.** Georgia Institute of Technology, College of Computing Colloquium. October 2009.
68. **The Dynamics of Equilibria.** Toyota Technological Institute at Chicago (TTI-C). August 2009.
69. **The Dynamics of Equilibria.** Harvard University, Economics and Computer Science Seminar. May 2009.
70. **A Theory of Learning and Clustering via Similarity Functions.** Google Research New York. April 2008.
71. **A Theory of Learning and Clustering via Similarity Functions.** Ecole Polytechnique Federale de Lausanne, School Seminar. April 2008.
72. **A Theory of Learning and Clustering via Similarity Functions.** Georgia Institute of Technology. April 2008.
73. **A Theory of Learning and Clustering via Similarity Functions.** University of Washington, CSE Colloquium. April 2008.
74. **A Theory of Learning and Clustering via Similarity Functions.** Massachusetts Institute of Technology. CS Special Seminar Series. March 2008.
75. **A Theory of Learning and Clustering via Similarity Functions.** University of Michigan, Ann Arbor. CSE Colloquium. March 2008.
76. **A Theory of Learning and Clustering via Similarity Functions.** Microsoft Research Silicon Valley. March 2008.
77. **A Theory of Learning and Clustering via Similarity Functions.** Stanford, Computer Science Seminar. March 2008.
78. **A Theory of Learning and Clustering via Similarity Functions.** University of Southern California, CS Colloquium Series. March 2008.
79. **A Theory of Learning and Clustering via Similarity Functions.** University of Wisconsin Madison, Computer Sciences Seminar. February 2008.
80. **A Theory of Learning and Clustering via Similarity Functions.** University of Pennsylvania, Wharton School Statistics. February 2008.
81. **A Theory of Learning and Clustering via Similarity Functions.** Microsoft Research Redmond. February 2008.
82. **A Theory of Learning and Clustering via Similarity Functions.** Cornell Theory Seminar. November 2007.
83. **A Theory of Learning and Clustering via Similarity Functions.** China Theory Week, Tsinghua University. September 2007.
84. **A Theory of Learning and Clustering via Similarity Functions.** University of Pennsylvania, Machine Learning Lunch. September 2007.
85. **Mechanism Design, Machine Learning, and Pricing Problems.** University of Southern California, CS Colloquium Series. December 2007.
86. **Mechanism Design, Machine Learning, and Pricing Problems.** Duke University Computer Science Colloquia. December 2007.
87. **Mechanism Design, Machine Learning, and Pricing Problems.** Brown CS Seminar. November 2007.
88. **Mechanism Design, Machine Learning, and Pricing Problems.** University of California, San Diego, Theory Seminar. October 2006.
89. **Mechanism Design, Machine Learning, and Pricing Problems.** Yahoo! Research, Sunnyvale, CA. October 2006.
90. **An Augmented PAC Model for Semi-Supervised Learning.** Toyota Technological Institute at Chicago (TTIC). August 2005.

91. **An Augmented PAC Model for Semi-Supervised Learning.** IBM Research T.J. Watson. June 2005.
92. **An Augmented PAC Model for Semi-Supervised Learning.** Microsoft Research Silicon Valley. May 2005.
93. **Kernels as Features: On Kernels, Margins, and Low-dimensional Mappings.** IBM Research T.J. Watson. July 2005.

## **UNIVERSITY SERVICE**

### **Carnegie Mellon University**

- Member of the Dean Search Advisory Committee 2018.
- Member of the 2018 CSD Faculty Hiring Committee.
- Member of the 2018 and 2019 MLMS Admissions Committees.
- Co-Director of the Master in Machine Learning Program (2017 -- Present).
- Member of the MLD PhD Admissions Committee 2017.
- Member of the SCS Undergraduate Review Committee (2016 -- Present).
- Member of the SCS Council (June 2015 -- Present).

### **Georgia Tech**

- Member, Ph.D Review, School of Computer Science, College of Computing, 2012 - 2014.
- Theory Area PhD Admissions Coordinator, 2013 and 2014.
- Member, ARC Committee for evaluating Graduate Student Fellowship Applications, 2012 and 2013.
- Member, Interim Steering Committee, Center for Data Analytics, 2013.
- Member, Search Committee for the Chair of the School of Computer Science, College of Computing, 2011.
- Main Coordinator for the written Ph.D. Qualifying Exam (Machine Learning area), Spring 2011.
- Member, College of Computing PhD Admissions Committee. 2010, 2011, and 2012.
- Question writer and grader for the Ph.D. Qualifying Exam (Machine Learning area), Fall 2009, Fall 2010, Spring 2011, Fall 2011, Fall 2012, and Spring 2013.