

CONTACT INFORMATION Amazon Research Cambridge
One the Square, Station Road
Cambridge, CB1 2GA (UK)

e-mail: borja.balle@gmail.com
web: borjaballe.github.io
skype: borja.balle
DBLP: <http://dblp.uni-trier.de/pers/hd/b/Balle:Borja>

QUICK SUMMARY **Background:** PhD in Computer Science, MSc in Applied Mathematics, BSc's in Mathematics and Telecommunication Engineering

Trajectory: UPC (Spain) – NYU (USA) – McGill University (Canada) – Lancaster University (UK) – Amazon Research (UK)

Research topics: Privacy-preserving machine learning with differential privacy and secure multi-party computation / Statistical and algorithmic aspects of machine learning / Applications to natural language processing and reinforcement learning / Analytical aspects of automata theory

Research statistics: 13 invited talks and tutorials, 32 peer-reviewed papers in top-tier journals and conferences, over 40 co-authors

Service highlights: United Nations task team on privacy-preserving technologies / Workshops chair (NIPS 2015) / Area chair (NIPS 2014, 2018) / ICGI steering committee / Co-organizer of 7 international workshops

Awards: best paper (EACL 2012), best student paper (ICGI 2012), best student paper runner-up (NIPS 2012), UPC extraordinary PhD thesis award

EDUCATION **Universitat Politècnica de Catalunya (UPC)**, Barcelona, Spain

PhD Computer Science — July 2013

- *Thesis:* Learning Finite-State Machines: Algorithmic and Statistical Aspects
- *Advisors:* Jorge Castro and Ricard Gavaldà

MSc Applied Mathematics — March 2009

- *Thesis:* On Whitehead's Algorithm for Equivalence Testing in Free Groups
- *Advisor:* Enric Ventura

BSc (5-year degree) Telecommunication Engineering — December 2007

- *Thesis:* An Algebraic Method for Constructing Cycles on DeBruijn's Graph and Applications to Shaft Encoding
- *Advisors:* Josep M. Fuertes and Enric Ventura

BSc (5-year degree) Mathematics — February 2007

- *Merits:* Participation in an exclusive program for interdisciplinary studies (CFIS)

AWARDS AND HONORS

Awards

- Extraordinary PhD Thesis Award, Universitat Politècnica de Catalunya, 2015
- Best Paper Award at EACL, 2012
- Best Student Paper Award at ICGI, 2012

- Runner-up for Best Student Paper Award at NIPS, 2012
- Best Undergraduate Thesis in Engineering, Catalan Technological Society, 2008

Honors

- Oral presentation at NIPS 2012 (5.4% of accepted papers)
- Runner-up of the Valortec Competition on business plans for technology start-ups
- 2nd place in the Zulu Competition 2010 on active learning for finite automata
- My NIPS 2012 paper was featured in Google’s list “Excellent Papers for 2012”

GRANTS AND FELLOWSHIPS

Competitive Fellowships

- FPU Doctoral Fellowship, Spanish Ministry of Education
August 2009 – January 2013
- FPU Travel Award (for visiting NYU), Spanish Ministry of Education
November 2011 – June 2012
- FI Doctoral Fellowship, Catalan Government
January 2009 – July 2009
- Undergraduate Research Fellowship, Spanish Ministry of Education
September 2005 – June 2006

POSITIONS HELD

Amazon Research, Cambridge, UK

Machine Learning Scientist, Core Machine Learning Group (April 2017 – ongoing)

- Research and development of privacy-aware machine learning systems

Lancaster University, Lancaster, UK

Lecturer, Department of Mathematics and Statistics (October 2015 – March 2017)

- Joint affiliation with the Data Science Institute

McGill University, Montreal, Canada

Postdoctoral Fellow, School of Computer Science (October 2013 – September 2015)

- Member of the “Reasoning and Learning Laboratory” working with professors Prakash Panangaden, Joelle Pineau, and Doina Precup

Universitat Politècnica de Catalunya, Barcelona, Spain

Research and Teaching Assistant, Department of Computer Science (January 2009 – July 2013)

- Member of the “Laboratory for Relational Algorithmics, Complexity and Learning”

RESEARCH VISITS New York University (NYU), New York, NY

- *Department:* Courant Institute for Mathematical Sciences
- *Host:* Mehryar Mohri
- *Period:* November 2011 – June 2012

OTHER INDUSTRIAL EXPERIENCE

Quipu, Barcelona, Spain

Advisor (January 2017 – ongoing)

Board Member (September 2013 – December 2016)

Freelance Consultant, Catalunya, Spain

Statistics and Information Technology (January 2008 – September 2013)

- Consulted for: RCR Arquitectes, Intelligent Pharma, Espuña, Ajuntament de Barcelona, Series.ly, and AqSense

Ericsson Research, Madrid, Spain

External Collaborator (September 2011 – February 2012)

- Participate in a joint research project with UPC to design and develop a prototype system to predict customer churning for telecommunication providers

Intelligent Pharma, Barcelona, Spain

Data Scientist (March – September 2008)

- Analyze results from large-scale in-silico experiments for drug discovery
- Write special-purpose tools to solve geometric tasks in computational chemistry

PUBLICATIONS

Journal Publications

- [1] B. Balle and M. Mohri. Generalization Bounds for Learning Weighted Automata. *Theoretical Computer Science*, 716:89–106, 2018. ISSN 0304-3975. (SJR Ranking: Q1).
- [2] A. Gascón, P. Schopmann, B. Balle, M. Raykova, S. Zahur, J. Doerner, and D. Evans. Privacy-Preserving Distributed Linear Regression on High-Dimensional Data. *Proceedings on Privacy Enhancing Technologies*, (4):345–364, 2017. (Acceptance Rate: 22%).
- [3] B. Balle, X. Carreras, F. M. Luque, and A. Quattoni. Spectral Learning of Weighted Automata: A Forward-Backward Perspective. *Machine Learning*, 96:33–63, 2014. ISSN 0885-612. (SJR Ranking: Q1).
- [4] B. Balle, J. Castro, and R. Gavaldà. Adaptively Learning Probabilistic Deterministic Automata from Data Streams. *Machine Learning*, 96:99–127, 2014. ISSN 0885-612. (SJR Ranking: Q1).
- [5] B. Balle, J. Castro, and R. Gavaldà. Learning Probabilistic Automata: A Study In State Distinguishability. *Theoretical Computer Science*, 473:46–60, 2013. ISSN 0304-3975. (SJR Ranking: Q1).
- [6] J. M. Fuertes, B. Balle, and E. Ventura. Absolute-Type Shaft Encoding Using LFSR Sequences With a Prescribed Length. *IEEE Transactions on Instrumentation and Measurement*, 57(5):915–922, 2008. ISSN 0018-9456. (SJR Ranking: Q1).

Peer-Reviewed Conference Publications

- [7] B. Balle and Y.-X. Wang. Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising. In *Proceedings of the 35th International Conference on Machine Learning (ICML) (Stockholm, Sweden, July 2018)*, 2018. (CORE Rating: A*; Acceptance Rate: 25%).
- [8] Y. Grinberg, H. Aboutaleb, M. Lyman-Abramovitch, B. Balle, and D. Precup. Learning Predictive State Representations from Non-uniform Sampling. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI (New Orleans, Louisiana, February 2018)*. AAAI Press, 2018. (Oral presentation; CORE Rating: A*; Acceptance Rate: 25%).

- [9] M. Ruffini, G. Rabusseau, and B. Balle. Hierarchical Methods of Moments. In *Advances in Neural Information Processing Systems (NIPS) (Long Beach, USA, December 2017)*, volume 30 of *Advances in Neural Information Processing Systems (ISSN 1049-5258)*, pages 1901–1911, 2017. (CORE Rating: A*; Acceptance Rate: 21%).
- [10] G. Rabusseau, B. Balle, and J. Pineau. Multitask Spectral Learning of Weighted Automata. In *Advances in Neural Information Processing Systems (NIPS) (Long Beach, USA, December 2017)*, volume 30 of *Advances in Neural Information Processing Systems (ISSN 1049-5258)*, pages 2588–2597, 2017. (CORE Rating: A*; Acceptance Rate: 21%).
- [11] B. Balle and O.-A. Maillard. Spectral Learning from a Single Trajectory under Finite-State Policies. In *Proceedings of the 34rd International Conference on Machine Learning (ICML) (Sydney, Australia, August 2017)*, volume 70 of *Proceedings of Machine Learning Research (PMLR)*, pages 361–370, 2017. (CORE Rating: A*; Acceptance Rate: 25%).
- [12] B. Balle, P. Gordeau, and P. Panangaden. Bisimulation Metrics for Weighted Automata. In *Proceedings of the 44rd International Colloquium on Automata, Languages, and Programming (ICALP) (Warsaw, Poland, July 2017)*, volume 80 of *LIPICs*, pages 103:1–103:14, 2017. (CORE Rating: A; Acceptance Rate: 30%).
- [13] B. Balle, M. Gomrokchi, and D. Precup. Differentially Private Policy Evaluation. In *Proceedings of the 33rd International Conference on Machine Learning (ICML) (New York, U.S., June 2016)*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2130–2138, 2016. (CORE Rating: A*; Acceptance Rate: 24%).
- [14] L. Langer, B. Balle, and D. Precup. Learning Multi-Step Predictive State Representations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI) (New York, U.S., July 2016)*, pages 1662–1668. IJCAI/AAAI Press, 2016. ISBN 978-1-57735-770-4. (CORE Rating: A*; Acceptance Rate: 25%).
- [15] G. Rabusseau, B. Balle, and S. Cohen. Low-Rank Approximation of Weighted Tree Automata. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) (Cadiz, Spain, May 2016)*, volume 51 of *JMLR Workshop and Conference Proceedings*, pages 839–847, 2016. (CORE Rating: A; Acceptance Rate: 31%).
- [16] C. Zhou, B. Balle, and J. Pineau. Learning Time Series Models for Pedestrian Motion Prediction. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (Stockholm, Sweden, May 2016)*, pages 3323–3330. IEEE, 2016. ISBN 978-1-4673-8026-3. (CORE Rating: B; Acceptance Rate: 35%).
- [17] B. Wang, B. Balle, and J. Pineau. Multitask Generalized Eigenvalue Program. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI (Phoenix, Arizona, February 2016)*, pages 2115–2121. AAAI Press, 2016. (CORE Rating: A*; Acceptance Rate: 26%).
- [18] B. Balle, R. Eyraud, F. M. Luque, A. Quattoni, and S. Verwer. Results of the sequence prediction challenge (spice): a competition on learning the next symbol in a sequence. In *Proceedings of the 13th International Conference on Grammatical Inference, ICGI 2016, Delft, The Netherlands, October 5-7, 2016*, volume 57 of *JMLR Workshop and Conference Proceedings*, pages 132–136. JMLR.org, 2016. (Invited Paper).

- [19] B. Balle and M. Mohri. On the Rademacher Complexity of Weighted Automata. In *Proceedings of the 26th International Conference on Algorithmic Learning Theory, ALT (Banff, Canada, October 2015)*, volume 9355 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2015. ISBN 978-3-319-24485-3. (CORE Rating: B; Acceptance Rate: 50%).
- [20] B. Balle and M. Mohri. Learning Weighted Automata. In *Proceedings of the 6th International Conference on Algebraic Informatics, CAI (Stuttgart, Germany, September 2015)*, volume 9270 of *Lecture Notes in Computer Science*, pages 1–21. Springer, 2015. ISBN 978-3-319-23020-7. (Invited Paper).
- [21] P. L. Bacon, B. Balle, and D. Precup. Learning and Planning with Timing Information in Markov Decision Processes. In *Proceedings of the 31th Uncertainty in Artificial Intelligence Conference, UAI (Amsterdam, The Netherlands, July 2015)*, pages 111–120. AUAI Press, 2015. ISBN 978-0-9966431-0-8. (CORE Rating: A*; Acceptance Rate: 34%).
- [22] B. Balle, P. Panangaden, and D. Precup. A Canonical Form for Weighted Automata and Applications to Approximate Minimization. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS (Kyoto, Japan, July 2015)*, pages 701–712. IEEE Computer Society, 2015. ISBN 978-1-4799-8875-4. (CORE Rating: A*; Acceptance Rate: 35%).
- [23] A. Quattoni, B. Balle, X. Carreras, and A. Globerson. Spectral Regularization for Max-Margin Sequence Tagging. In *Proceedings of the 31th International Conference on Machine Learning, ICML (Beijing, China, June 2014)*, volume 32 of *JMLR Proceedings (ISSN 1938-7228)*, pages 1386–1394, 2014. (CORE Rating: A*; Acceptance Rate: 25%).
- [24] B. Balle, W. L. Hamilton, and J. Pineau. Methods of Moments for Learning Stochastic Languages: Unified Presentation and Empirical Comparison. In *Proceedings of the 31th International Conference on Machine Learning, ICML (Beijing, China, June 2014)*, volume 32 of *JMLR Proceedings (ISSN 1938-7228)*, pages 1710–1718, 2014. (CORE Rating: A*; Acceptance Rate: 25%).
- [25] B. Balle, B. Casas, A. Catarineu, R. Gavaldà, and D. Manzano-Macho. The Architecture of a Churn Prediction System Based on Stream Mining. In *Artificial Intelligence Research and Development - Proceedings of the 16th International Conference of the Catalan Association for Artificial Intelligence (Vic, Spain, October 2014)*, volume 256 of *Frontiers in Artificial Intelligence and Applications*, pages 157–166. IOS Press, 2013. ISBN 978-1-61499-319-3. (CORE Rating: N/A).
- [26] B. Balle and M. Mohri. Spectral Learning of General Weighted Automata via Constrained Matrix Completion. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012 (Lake Tahoe, USA, December 2012)*, Advances in Neural Information Processing Systems (ISSN 1049-5258), pages 2168–2176, 2012. (CORE Rating: A*; Acceptance Rate: 25%).
- [27] B. Balle, J. Castro, and R. Gavaldà. Bootstrapping and Learning PDFAs in Data Streams. In *Proceedings of the 11th International Conference on Grammatical Inference, ICGI (Washington, USA, September 2012)*, volume 21 of *JMLR Proceedings (ISSN 1938-7228)*, pages 34–48, 2012. (CORE Rating: N/A).
- [28] B. Balle, A. Quattoni, and X. Carreras. Local Loss Optimization in Operator Models: A New Insight Into Spectral Learning. In *Proceedings of the 29th*

International Conference on Machine Learning, ICML (Edinburgh, UK, June 2012), 2012. (CORE Rating: A*; Acceptance Rate: 27%).

- [29] F. M. Luque, A. Quattoni, B. Balle, and X. Carreras. Spectral Learning for Non-Deterministic Dependency Parsing. In *EACL 2012, 13th Conference of the European Chapter of the Association for Computational Linguistics (Avignon, France, April 2012)*, pages 409–419. The Association for Computer Linguistics, 2012. ISBN 978-1-937284-19-0. (CORE Rating: A; Acceptance Rate: 27%).
- [30] B. Balle, A. Quattoni, and X. Carreras. A Spectral Learning Algorithm for Finite State Transducers. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML-PKDD. Proceedings, Part I (Athens, Greece, September 2011)*, volume 6911 of *Lecture Notes in Computer Science*, pages 156–171. Springer, 2011. ISBN 978-3-642-23779-9. (CORE Rating: A; Acceptance Rate: 20%).
- [31] B. Balle, J. Castro, and R. Gavaldà. A Lower Bound for Learning Distributions Generated by Probabilistic Automata. In *Proceedings of Algorithmic Learning Theory, 21st International Conference, ALT (Canberra, Australia, October 2010)*, volume 6331 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2010. ISBN 978-3-642-16107-0. (CORE Rating: B).
- [32] B. Balle, J. Castro, and R. Gavaldà. Learning PDFAs with Asynchronous Transitions. In *Grammatical Inference: Theoretical Results and Applications, 10th International Colloquium. Proceedings (Valencia, Spain, September 2010)*, volume 6339 of *Lecture Notes in Computer Science*, pages 271–275. Springer, 2010. ISBN 978-3-642-15487-4. (Short paper; CORE Rating: N/A).
- [33] B. Balle, E. Ventura, and J.M. Fuertes. An Algorithm to Design Prescribed Length Codes for Single-tracked Shaft Encoders. In *IEEE International Conference on Mechatronics (ICM) (Malaga, Spain, April 2009)*, 2009. ISBN 978-1-4244-4195-2. (CORE Rating: N/A).

Workshop Contributions

- [34] P. Schoppmann, A. Gascón, and B. Balle. Secure Text Classification in Federated Databases. *Presentation at Workshop Theory and Practice of Multi-Party Computation (Aahrus, Denmark, May 2018)*, 2018.
- [35] B. Balle. Revisiting the Gaussian Mechanism. *Poster at CCS Workshop Theory and Practice of Differential Privacy (Dallas, TX, U.S.A., October 2017)*, 2017.
- [36] P. Schoppmann, A. Gascón, and B. Balle. Private Nearest Neighbors Classification in Federated Databases. *Poster at ICML Workshop Secure and Private Machine Learning (Sydney, Australia, July 2017)*, 2017.
- [37] A. Gascón, P. Schoppmann, B. Balle, M. Raykova, S. Zahur, J. Doerner, and D. Evans. Privacy-Preserving Distributed Linear Regression on High-Dimensional Data. *Poster at ICML Workshop Secure and Private Machine Learning (Sydney, Australia, July 2017)*, 2017.
- [38] A. Gascón, P. Schoppmann, B. Balle, M. Raykova, S. Zahur, J. Doerner, and D. Evans. Privacy-Preserving Distributed Linear Regression on High-Dimensional Data. *Poster at NIPS Workshop Private Multi-Party Machine Learning (Barcelona, Spain, December 2016)*, 2016.
- [39] B. Balle, M. Gomrokchi, and D. Precup. Differentially Private Policy Evaluation. *Poster at ICML Workshop Theory and Practice of Differential Privacy (New York, NY, U.S.A., June 2016)*, 2016.

- [40] P. Schoppmann, A. Gascón, and B. Balle. Solving Private Systems of Linear Equations with Garbled Circuits. *Poster at the 37th IEEE Symposium on Security and Privacy (San Jose, CA, U.S.A., May 2016)*, 2016.
- [41] L. Langer, B. Balle, and D. Precup. Multi-Step Predictive State Representations. *Poster at NIPS Workshop on Time Series (Montreal, Canada, December 2015)*, 2015.
- [42] B. Wang, B. Balle, and J. Pineau. Multitask Generalized Eigenvalue Program. *Poster at NIPS 2015 Workshop on Transfer and Multi-Task Learning: Trends and New Perspectives (Montreal, Canada, December 2015)*, 2015.
- [43] P. L. Bacon, B. Balle, and D. Precup. Learning and Planning with Timing Information in Markov Decision Processes. *Poster at 2nd Multi-disciplinary Conference on Reinforcement Learning and Decision Making, RLDM (Edmonton, Canada, June 2015)*, 2015.
- [44] P. L. Bacon, B. Balle, and D. Precup. Learning and Planning with Timing Information in Markov Decision Processes. *Talk and Poster at NIPS Workshop From Bad Models to Good Policies (Montreal, Canada, December 2014)*, 2014.
- [45] A. M.S. Barreto, B. Balle, J. Pineau, and D. Precup. Starting to Uncover the Relationship Between Stochastic Factorization and Hidden Markov Models. *Poster at NIPS Workshop Novel Trends and Applications in RL (Montreal, Canada, December 2014)*, 2014.
- [46] B. Balle, J. Castro, and R. Gavaldà. Learning Markovian Models from Time-Evolving Data Streams. *Talk at From Data to Knowledge: Machine-Learning with Real-time and Streaming Applications (Berkeley, California, May 2012)*, 2012.
- [47] B. Balle, A. Quattoni, and X. Carreras. Local Loss Optimization in Operator Models: A New Insight into Spectral Learning. *Poster at New England Machine Learning Day (Cambridge, Massachusetts, May 2012)*, 2012.
- [48] B. Balle. Implementing Kearns-Vazirani Algorithm for Learning DFA Only with Membership Queries. *Talk and Paper at Zulu Competition Workshop (Valencia, Spain, September 2010)*, 2010.

Preprints and Technical Reports

- [49] P. Schoppmann, A. Gascón, and B. Balle. Private Nearest Neighbors Classification in Federated Databases. *Under Review, Cryptology ePrint 2018/289*, 2018.
- [50] B. Balle, P. Panangaden, and D. Precup. Singular Value Automata and Approximate Minimization. *Under Review, Arxiv 1711.05994*, 2017.
- [51] L. Addario-Berry, B. Balle, and G. Perarnau. Diameter and Mixing Time in Random d -out Digraphs. *Preprint, Arxiv 1504.06840*, 2015.
- [52] B. Balle. Ergodicity of Random Walks on Random DFA. *Technical Report, Arxiv 1311.6830*, 2013.

TEACHING AND
SUPERVISION
EXPERIENCE

Cambridge University, Cambridge, United Kingdom

- Supervisor (MPhil in Machine Learning, Speech and Language Technology final project) — Summer 2018
- Alexandru Coca, *Differentially Private Generative Models*

University College London (UCL), London, United Kingdom

- Co-supervisor (MSc in Information Security final project) — Summer 2018
 - Qinwu Wang, *Differentially Private Boolean Circuits*

Summer School on Foundations of Programming and Software Systems (FoPSS), Oxford, United Kingdom

- Invited lecturer — July 2018
 - Introductory graduate class on *Automata Learning* (3 hours)

Oxford University, United Kingdom

- Visiting lecturer — February 2016
 - Advanced graduate class on *Automata and Learning* (6 hours)

Lancaster University, United Kingdom

- Co-supervisor (MSc Data Science final project) — Summer 2016
 - Syed Jawaad, *Optimizing Heating Policies with Reinforcement Learning*
- Supervisor (EPSRC-funded undergraduate research internship) — Fall 2016
 - Aidan Metcalfe, *The Mathematics of Privacy*

McGill University, Montreal, Canada

Reading Group Organizer — Summer 2014

- *Topic*: Spectral learning algorithms
- *Students*: ~10

Supervisor (undergraduate honors project) — Summer/Fall 2014

- Chenghui Zhou, *A Spectral Method for Learning Motion Models with Applications to Human Tracking*

Supervisor (undergraduate summer research internships)

- Chenghui Zhou (2014, 2015)
- Lucas Langer (2015)
- Pascale Gourdeau (2015)

Universitat Politècnica de Catalunya, Barcelona, Spain

Supervisor (undergraduate final year project) — Spring 2013

- Albert Santiago, *Implementation of a Spectral Method for Learning Finite-State Machines*

Lecturer (Theory of Computation) — Spring 2011

- *Syllabus*: Introduction to language theory, finite automata, context-free grammars, and Turing machines
- *Students*: ~50

Teaching Assistant (Information Retrieval) — Spring 2011

- *Responsibilities*: Design and mark computer assignments; explain relation between lectures and assignments
- *Students*: ~30

Tutorials

- Tutorial on “Singular Value Automata” (9th International Workshop Weighted Automata: Theory and Applications (WATA), Leipzig, May 2018)
- Tutorial on “Differential Privacy for Machine Learning” co-presented with S. Kaviswanathan and Y.-X. Wang (Amazon Machine Learning Conference, May 2018)
- Tutorial on “Differential Privacy” (The Alan Turing Institute, January 2018)
- Tutorial on “(Co)-Algebraic and Analytical Aspects of Weighted Automata Minimisation and Equivalence” co-presented with A. Silva (13th International Workshop on Coalgebraic Methods in Computer Science (CMCS), Eindhoven, April 2016)
- Tutorial on “Spectral Learning Techniques for Weighted Automata, Transducers, and Grammars” co-presented with X. Carreras and A. Quattoni (Conference on Empirical Methods in Natural Language Processing (EMNLP), Qatar, October 2014)

Invited Talks

- Invited talk on “Implementation and Deployment Challenges in Private Machine Learning Systems” (Royal Society Workshop on Privacy-Enhancing Technologies, July 2018)
- Invited talk on “Learning and Automata” (Logic and Learning Workshop, The Alan Turing Institute, January 2018)
- Invited talk on “Machine Learning Algorithms for Weighted Automata” (5th Highlights of Logic, Games and Automata Conference, London, September 2017)
- Invited talk on “Privacy-Preserving Distributed Linear Regression on High-Dimensional Data” (Workshop on Decentralized Machine Learning, Optimization and Privacy, INRIA Lille, September 2017)
- Invited talk on “Secure Multi-Party Linear Regression on High-Dimensional Data” (Advances in Data Science, Manchester, May 2017)
- Invited talk on “Secure Multi-Party Linear Regression on High-Dimensional Data” (Workshop on Privacy in Statistical Analysis, Imperial College London, May 2017)
- Invited talk on “Approximation Algorithms for Weighted Automata” (Bellairs Workshop on Algorithmic Aspects of Dynamical Systems, Barbados, March 2017)
- Invited talk on “Differential Privacy and Secure Multi-Party Computation in Linear Regression” (QARMA Workshop on Machine Learning, Marseille, October 2016)
- Invited talk on “Theoretical Guarantees for Learning Weighted Automata” (13th International Conference on Grammatical Inference (ICGI), Delft, October 2016)
- Invited talk on “Spectral Algorithms for Learning Predictive State Representations with Prior Information” (9th Barbados Workshop on Reinforcement Learning, Barbados, April 2015)

Seminar Talks

- Imperial College London (March 2018)
- Cambridge University (October 2017)
- University of Kent (February 2017)
- Trinity College Dublin (January 2017)
- Imperial College London (October 2016)
- University of Leicester (June 2016)
- The University of Edinburgh (March 2016)
- Xerox Research Center Europe (November 2015)
- INRIA Lille (February 2015)
- Université d’Aix-Marseille (September 2014)
- Oxford University (September 2014)
- Toyota Technical Institute Chicago (April 2013)
- Google Mountain View (February 2013)
- McGill University (September 2012)
- Massachusetts Institute of Technology (May 2012)
- Google New York (February 2012)

- New York University (November 2011)
- Tokyo Institute of Technology (November 2010)
- The University of Waikato (October 2010)

SERVICE

Committees and Advisory Boards

- Member of the “Task Team on Privacy Preserving Technologies” within the United Nations’ “Global Working Group on Big Data for Official Statistics” (2018-onwards)
- Member of the Steering Committee for the “International Conference in Grammatical Inference” (2016-onwards)

Organization of Workshops and Other Events

- Workshop on “Privacy in Machine Learning and Artificial Intelligence” at ICML 2018
- Workshop on “Learning and Automata” at LICS 2017
- Workshop on “Privacy and Fairness in Machine Learning” at DALI 2017
- Workshop on “Private Multi-Party Machine Learning” at NIPS 2016
- Machine learning competition “Sequence Prediction Challenge (SPICE)” within ICGI 2016
- Workshop on “Methods of Moments and Spectral Learning” at ICML 2014
- Workshop on “Spectral Learning” at NIPS 2013
- Workshop on “Spectral Learning” at ICML 2013

Conference and Workshop Program Committees

- 2019: ALT
- 2018: NIPS (area chair), ICML, ALT, ICGI, CCS (reviewer), CONCUR (reviewer), LearnAut (FLOC-LICS workshop)
- 2017: AISTATS, EACL, ICML, COLT (reviewer), NIPS, PSML (ICML workshop)
- 2016: ICML, NIPS, ACL, EMNLP, AISTATS, ICGI, StatFSM (ACL workshop), CAp
- 2015: NIPS (workshops chair), ICML, IJCAI, VSM-NLP (NAACL workshop)
- 2014: NIPS (area chair), ICML, EMNLP, EACL, ICGI
- 2013: ICML, NIPS
- 2012: NIPS
- 2011: ICALP (reviewer), ICML (reviewer)

Reviewer for Journals

- Theoretical Computer Science
- Machine Learning
- Journal of Machine Learning Research

Other

- Invited to participate in the Royal Society Workshop on Privacy Enhancing Technologies (May 2018)
- Reviewer for the French National Research Agency evaluation panel on Mathematics and Computer Science (2017)
- External examiner for the PhD viva of Naoise Holohan at Trinity College Dublin (January 2017)
- Member of the PhD examining board for Guillaume Rabusseau at Université d’Aix-Marseille (October 2016)