Elad Hazan - CV

Email ehazan@princeton.edu
Web page www.ehazan.com

Academic Degrees

2006, Ph.D, Department of Computer Science, Princeton University
Advisor: Prof. Sanjeev Arora
2002, M.Sc, Department of Computer Science, Tel Aviv University
2001, B.Sc, Department of Computer Science, Tel Aviv University

Academic and Research Appointments

Jul 2016-today Professor of Computer Science, Princeton University
Jan 2018-today Co-founder and Director , Google AI Princeton
Apr 2016 - Apr 2017 Visiting Professor, Google Research, NYC (part time)
Jan 2015-June 2016 Assistant Professor of Computer Science, Princeton University
2013-2015 Associate Professor (tenured), Technion - Israel Institute of Technology
Mar 2014- Aug 2015 Visiting Researcher, Microsoft Research, Herzliya , Israel (part time)
2010-2013 Assistant Professor , Technion - Israel Institute of Technology
2006-2010 Research Staff Member "The Theory Group", IBM Almaden Research Center

Research Interests

I am interested in designing algorithms for fundamental problems in machine learning and optimization that are practically efficient and whose analysis is theoretically elegant. Most of my research can be classified under machine learning and mathematical optimization. More recently I'm also interested in control theory and reinforcement learning.

Industry relations

Co-inventor of **six US patents**, three issued and two filed. **Co-founder and CEO** of machine learning and optimization startup In8 inc. in 2017 based in Princeton, NJ. In8 was acquired by the Google corporation in 2018 to create **Google AI Princeton**. Currently co-founder and director of Google AI Princeton.

Honors

2022	**ACM Fellow**.
2018	Best Paper Award, Optimization Foundations for RL workshop, NeurIPS2019.
2018	Amazon Research Award
2017	**Bell Labs Prize**
2015	Google Research Award
2013	**European Research Council (ERC) Starting Grant Recipient**
2012	The 2012 IBM Pat Goldberg Memorial **Best Paper** Award.
2012	ICML **Best Student Paper** Runner Up
2012	Bar-Ilan Financial Models **Best Paper** Award
2011	Google Research Award
2011-2015	European Research Council Marie Curie Fellow
2009	The 2009 IBM Pat Goldberg Memorial **Best Paper** Award.
2008	Machine Learning Journal Award **best student paper** (COLT 2008).
2002–2006	Gordon Wu fellowship, Princeton University.
1999 and 2000	Dean prize from Tel Aviv University, Faculty of Exact Sciences.
1999	Dean prize from Tel Aviv University, Faculty of Engineering.

Graduate Students

Theses Completed

2023	Udaya Ghai	Ph.D	primary supervisor, Princeton
2023	Edgar Minasyan	Ph.D	primary supervisor, Princeton
2023	Daniel Suo	Ph.D	co-advised w. Kai Li, Princeton
2021	Karan Singh	Ph.D	primary supervisor, Princeton
2020	Cyril Zhang	Ph.D	primary supervisor, Princeton
2019	Brian Bullins	Ph.D	primary supervisor, Princeton
2018	Naman Agarwal	Ph.D	primary supervisor, Princeton
2016	Tomer Koren	Ph.D	primary supervisor, Technion
2016	Kfir Levi	Ph.D	primary supervisor, Technion
2016	Oren Anava	Ph.D	primary supervisor, Technion
2015	Dan Garber	Ph.D	primary supervisor, Technion

Theses in Progress

2023	Xinyi Chen	Ph.D	primary supervisor, Princeton
2023	Nataly Brukhim	Ph.D	primary supervisor, Princeton
2024	Zhou Lu	Ph.D	primary supervisor, Princeton
2024	Wenhan Xia	Ph.D	primary supervisor, Princeton
2025	Jennifer Sun	Ph.D	primary supervisor, Princeton

Postdoctoral Fellows

2015-2017	Zeyuan Allen-Zhu	Princeton and IAS
2015-2017	Roi Livni	Princeton and IAS
2016-2018	Nadav Cohen	Princeton and IAS
2017-2019	Alon Gonen	Princeton
2019-2020	Shay Moran	Google AI Princeton

Research Grants

- 2023-2027, 500,000 $, Open Philantropy , PI: , Elad Hazan "Game theoretic approaches to AI safety and alignment"
- 2021-2023, 150,000 $, National Science Foundation , PI: , Thomas Poggio, Christos Papadimitriou, Elad Hazan, Santosh Vempala, "Collaborative Research: Foundations of Deep Learning: Theory, Robustness, and the Brain"
- 2017-2020, 2,200,000 $, National Science Foundation , PI: Sanjeev Arora, Elad Hazan, Yoram Singer and Rong Ge, "Collaborative Research: Nonconvex methods and models for learning"
- 2015-2019, 500,000 $, National Science Foundation , PI: Elad Hazan, "Efficient projection-free algorithms for optimization and learning"
- 2015-2016, 80000 $, Google Research Award, PI: Elad Hazan
- 2013-2018, 1,450,000 €, European Research Council, Starting Grant (ERC-StG) SUBLRN , PI: Elad Hazan, "Information-optimal machine learning"
- 2011-2015, 858,000 NIS, Israeli Science Foundation (ISF) Grant 810/11 , PI: Elad Hazan, "Data Analysis in Sublinear Time"
- 2010-2014, 100,000 €, European Union (EU) Grant FP7-MARIE CURIE-IRG, PI: Elad Hazan, "Sublinear Optimization for Machine Learning"
- 2011-2012, 37000 $, Google Research Award, PI: Elad Hazan
- 2012-2015, 45000 $, Microsoft Electronic Commerce Center grant, PI: Elad Hazan and Sham Kakade, "Efficient algorithms for contextual bandits"
- 2012-2014, 50000 $, Intel Machine-Learning grant, PI: Elad Hazan , "Machine learning with 20/20 sight"

Public Professional Activities

- Senior area chair and/or area chair for NeurIPS 2023-2017, COLT 2023-2020, ICML 2018-2017.
- **Program chair for COLT (Computational Learning Theory) 2015**.
- Associate editor for "Mathematical Programming Series A", 2015-2019.
- On the editorial board for "Machine Learning Journal", since 2011.
- Steering committee member of COLT (Computational Learning Theory), 2013-2016
- Program committee member for the conferences NeurIPS 2017 (area chair), ICML 2018 (area chair), COLT 2018, NeurIPS 2015 (area chair), COLT 2014, NeurIPS 2013 (area chair), ICML 2013 (area chair), COLT 2013 (open problems chair), NeurIPS 2012, ICML 2012 (area chair), COLT 2012, NIPS 2011, COLT 2011, NIPS 2010, COLT 2009, AISTAT 2009, APPROX 2008
- Reviewer for *Math of Operations Research, Mathematical Programming, SIAM Journal on Computing, Journal of Machine Learning Research, IEEE Transactions on Information Theory, IEEE Transactions on Neural Networks, Algorithmica*, Machine Learning Journal, and other journals.

Invited Talks (selected)

- *The Online Convex Optimization approach to Control*
  Hariri Institute Director's Esteemed seminar at Boston University, November 2022
- *Keynote: Optimization for machine learning*
  Keynote talk at INFORMS, November 2020
- *Inaugural invited talk: The Non-Stochastic Control Problem*
  Inaugural invited talk, control meets learning seminar, Sept 2020
  EPFL IC colloquium, Nov 2019
- *YINS Distinguished Lecture: Optimization for machine learning*
  Yale University, October 2018
- *Invited tutorial: Optimization for machine learning*
  Machine Learning Summer School (MLSS 2018), Buenos Aires, June 2018
  and Simons Institute for the Theory of Computing, Jan 2017.

- *Semi-Plenary: Simulated annealing with an efficient universal barrier*
  Semi-plenary talk at The fifth International Conference on Continuous Optimization (ICCOPT 2016)
  Tokyo, Japan, Aug 10, 2016

- *Keynote: Fast convex optimization: simulated annealing with an efficient universal barrier*
  Keynote talk at Optimization for Machine Learning (OPT 2015), NIPS 2015
  Montreal, Canada, Dec 11, 2015

- *Plenary: The computational power of optimization in online learning*
  Plenary talk at Challenges in Optimization for Data Science
  Université Pierre et Marie Curie, Paris, July 1-2 2015

- *Keynote: Agnostic and Non-proper learning: a methodology for overcoming NP-hardness*
  Keynote talk at 9th Annual Machine Learning Symposium
  New York Academy of Sciences, 9 March 2015, NYC

- *Tutorial: online convex optimization*
  Invited talk at Online Algorithms and Learning 17-21 Nov, 2014, Lorentz center, Netherlands

- *Sublinear Optimization*
  Invited talk at Algorithmic Frontiers Workshop June 11-14, 2012, EPFL, Lausanne and Optimization and Statistical Learning workshop, January 6-11, 2013, Les Houches, France

- *Sublinear Optimization for Machine Learning*
  Invited talk at IPAM workshop on discrete math, Oct. 2010, Los Angeles, CA and Oberwolfach workshop Aug 2011.

- *Learning with limited information*
  Invited talk at 2nd CompView Fall School on Computational View to Machine Learning, Tokyo, Japan September 17-19 2009.

Teaching

AI Safety & Alignment, COS 598:
primary instructor
2022-2023 at Princeton University.

Computational Control Theory, COS 598:
primary instructor
2020-2021 at Princeton University.

Optimization for Machine Learning, COS 598:
primary instructor
2019-2020 at Princeton University.

Theoretical Machine Learning, COS 511:
primary instructor
2014-2016, 2019-2020, 2022-2023 at Princeton University.

Introduction to Machine Learning, COS 324:
primary co-instructor
2017-2018, 2020-2021 at Princeton University.

Artificial Intelligence, COS 402:
primary co-instructor
2016-2017 at Princeton University.

Quantitative decision analysis, 097328: (joint undergrad-grad course)
primary instructor, **new course** designed and taught in academic years
2010-2011, 2011-2012 and 2014-2015 at the Technion.

Machine learning and online optimization, 097209: (joint undergrad-grad course),
primary instructor, **new course** designed and taught in academic years
2012-2013 at the Technion.

Polynomial time algorithms for linear programming, 096328: (graduate level course),
primary instructor, **new course** designed and taught in academic years
2011-2012 at the Technion.

Stochastic Simulation, 094334:
undergrad course, compulsory in the IE&M undergraduate track
primary instructor, taught in academic years 2010-2011, 2011-2012 and
2012-2013 at the Technion.

Publications

Authors with asterisks* next to their names are student advisees.

Books and book chapters

1. Introduction to Online Convex Optimization
   Elad Hazan
   Second Edition: MIT Press, 2022
   First Edition: Foundations and Trends in Optimization 2(3-4): 157-325 (2016)

2. Introduction to Online Nonstochastic Control
   Elad Hazan and Karan Singh
   arXiv:2211.09619
   Cambridge University Press forthcoming

3. The convex optimization approach to regret minimization
   Elad Hazan
   Optimization for Machine Learning, The MIT Press, pages 287-302, 2011

4. Lecture Notes: Optimization for Machine Learning
   Elad Hazan
   arXiv:1909.03550, 2019

Refereed papers in conference proceedings

5. Partial Matrix Completion
   E Hazan, AT Kalai, V Kanade, C Mohri, YJ Sun
   Thirty-seventh Conference on Neural Information Processing Systems 2023

6. Online Control for Meta-optimization
   X Chen, E Hazan
   Thirty-seventh Conference on Neural Information Processing Systems 2023

7. Sketchy: Memory-efficient Adaptive Regularization with Frequent Directions
   V Feinberg, X Chen, YJ Sun, R Anil, E Hazan
   Thirty-seventh Conference on Neural Information Processing Systems 2023

8. Online Nonstochastic Model-Free Reinforcement Learning
   U Ghai, A Gupta, W Xia, K Singh, E Hazan
   Thirty-seventh Conference on Neural Information Processing Systems 2023

9. Optimal Rates for Bandit Nonstochastic Control
   YJ Sun, S Newman, E Hazan
   Thirty-seventh Conference on Neural Information Processing Systems 2023

10. Regret Guarantees for Online Deep Control
X Chen, E Minasyan, JD Lee, E Hazan
Learning for Dynamics and Control Conference, 1032-1045 4 2023

11. Best of Both Worlds in Online Control: Competitive Ratio and Policy Regret
G Goel, N Agarwal, K Singh, E Hazan
Learning for Dynamics and Control Conference, 1345-1356 3 2023

12. Adaptive regret for control of time-varying dynamics
P Gradu, E Hazan, E Minasyan
Learning for Dynamics and Control Conference, 560-572 41 2023

13. Projection-free adaptive regret with membership oracles
Z Lu, N Brukhim, P Gradu, E Hazan
International Conference on Algorithmic Learning Theory, 1055-1073 7 2023

14. A Regret Minimization Approach to Multi-Agent Control
Udaya Ghai, Udari Madhuhshani, Naomi Leonard, Elad Hazan
Thirty-ninth International Conference on Machine Learning (ICML 2022)

15. Non-convex online learning via algorithmic equivalence
U Ghai, Z Lu, E Hazan
Advances in Neural Information Processing Systems 35 (NeurIPS 2022)

16. A boosting approach to reinforcement learning
N Brukhim, E Hazan, K Singh
Advances in Neural Information Processing Systems 35 (NeurIPS 2022)

17. Robust Online Control with Model Misspecification
U Ghai, X Chen, E Hazan, A Megretski
Learning for Dynamics and Control Conference 2022, 1163-1175

18. Machine Learning for Mechanical Ventilation Control
Daniel Suo , Cyril Zhang , Paula Gradu , Udaya Ghai , Xinyi Chen , Edgar Minasyan , Naman Agarwal , Karan Singh , Julienne LaChance , Tom Zajdel , Manuel Schottdorf , Daniel Cohen , Elad Hazan
Machine Learning for Health (ML4H 2021)

19. Online Control of Unknown Time-Varying Dynamical Systems
E Minasyan, P Gradu, M Simchowitz, E Hazan
Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

20. Multiclass Boosting and the Cost of Weak Learning
N Brukhim, E Hazan, S Moran, I Mukherjee, RE Schapire
Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

21. Black-box control for linear dynamical systems
X Chen, E Hazan
Conference on Learning Theory, 1114-1143 (COLT 2021)

22. Online Boosting with Bandit Feedback
N Brukhim, E Hazan
Algorithmic Learning Theory, 397-420 (ALT 2021)

23. Generating Adversarial Disturbances for Controller Verification
U Ghai, D Snyder, A Majumdar, E Hazan
Learning for Dynamics and Control, 1192-1204 (L4DC 2021)

24. A Regret Minimization Approach to Iterative Learning Control
N Agarwal, E Hazan, A Majumdar, K Singh
Learning for Dynamics and Control, (L4DC 2021)

25. Boosting for control of dynamical systems
N Agarwal, N Brukhim, E Hazan, Z Lu
International Conference on Machine Learning (ICML 2020), 96-103

26. The gradient complexity of linear regression
M Braverman, E Hazan, M Simchowitz, B Woodworth
Conference on Learning Theory (COLT 2020), 627-647

27. Faster Projection-free Online Learning
E Hazan, E Minasyan*
Conference on Learning Theory (COLT 2020), 627-647

28. Improper learning for non-stochastic control M Simchowitz, K Singh*, E Hazan
Conference on Learning Theory (COLT 2020)

29. The nonstochastic control problem
E Hazan, S Kakade, K Singh*
Algorithmic Learning Theory, 408-421 14 (ALT 2020)

30. Exponentiated gradient meets gradient descent
U Ghai*, E Hazan, Y Singer
Algorithmic Learning Theory, 386-407 (ALT 2020)

31. Geometric Exploration for Online Control
O Plevrakis, E Hazan
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)

32. Non-stochastic control with bandit feedback
P Gradu*, J Hallman*, E Hazan
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)

33. Online Agnostic Boosting via Regret Minimization
N Brukhim*, X Chen*, E Hazan, S Moran
Advances in Neural Information Processing Systems 33 (NeurIPS 2020)

34. Logarithmic Regret for Online Control
Naman Agarwal, Elad Hazan, Karan Singh*
Thirty-third Conference on Neural Information Processing Systems (NeurIPS 2019)
Best Paper Award: Optimization Foundations for RL workshop, NeurIPS2019

35. Private Learning Implies Online Learning: An Efficient Reduction
Alon Gonen, Elad Hazan, Shay Moran
Thirty-third Conference on Neural Information Processing Systems (NeurIPS 2019)

36. Provably efficient maximum entropy exploration
Elad Hazan, Sham Kakade, Karan Singh*, Abby Van Soest*
Proceedings of the 36th International Conference on Machine Learning (ICML 2019)

37. Efficient Full-Matrix Adaptive Regularization
Naman Agarwal, Brian Bullins*, Xinyi Chen*, Elad Hazan, Karan Singh*, Cyril Zhang*, Yi Zhang
Proceedings of the 36th International Conference on Machine Learning (ICML 2019)

38. Online Control with Adversarial Disturbances
Naman Agarwal, Brian Bullins*, Elad Hazan, Sham Kakade, Karan Singh*
Proceedings of the 36th International Conference on Machine Learning (ICML 2019)

39. Learning in Non-convex Games with an Optimization Oracle
Naman Agarwal, Alon Gonen, Elad Hazan
COLT 2019

40. Generalize Across Tasks: Efficient Algorithms for Linear Representation Learning
Brian Bullins*, Elad Hazan, Adam Kalai, Roi Livni
Algorithmic Learning Theory (ALT 2019), 235-246

41. Spectral Filtering for General Linear Dynamical Systems.
Elad Hazan, Holden Lee, Karan Singh*, Cyril Zhang*, Yi Zhang
Thirty-second Conference on Neural Information Processing Systems (NIPS 2018)

42. Online Learning of Quantum States.
Scott Aaronson, Xinyi Chen*, Elad Hazan, Satyen Kale, Ashwin Nayak
Thirty-second Conference on Neural Information Processing Systems (NIPS 2018)

43. Online Improper Learning with an Approximation Oracle.
Elad Hazan, Wei Hu, Yuanzhi Li, Zhiyuan Li
Thirty-second Conference on Neural Information Processing Systems (NIPS 2018)

44. Lower Bounds for Higher-Order Convex Optimization.
Naman Agarwal, Elad Hazan
Proceedings of the 31st Conference on Learning Theory (COLT 2018)

45. On the Optimization of Deep Networks: Implicit Acceleration by Overparameterization.
Sanjeev Arora, Nadav Cohen, Elad Hazan
Proceedings of the 35th International Conference on Machine Learning (ICML 2018)

46. Linear Convergence of a Frank-Wolfe Type Algorithm over Trace-Norm Balls
Zeyuan Allen-Zhu, Elad Hazan, Wei Hu*, Yuanzhi Li*
Proceedings of 31st Annual Conference on Neural Information Processing Systems (NIPS 2017) 6192-6201

47. Learning Linear Dynamical Systems via Spectral Filtering
Elad Hazan, Karan Singh*, Cyril Zhang*
Proceedings of 31st Annual Conference on Neural Information Processing Systems (NIPS 2017) 6705-6715

48. Efficient Regret Minimization in Non-Convex Games.
Elad Hazan, Karan Singh*, Cyril Zhang*
Proceedings of the 34th International Conference on Machine Learning (ICML 2017), 1433-1441

49. Finding Approximate Local Minima Faster than Gradient Descent
Naman Agarwal*, Zeyuan Allen-Zhu, Brian Bullins*, Elad Hazan, Tengyu Ma*
49th Annual Symposium on the Theory of Computing (STOC 2017)

50. A non-generative framework and convex relaxations for unsupervised learning
Elad Hazan and Tengyu Ma
Proceedings of 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)

51. Optimal Black-Box Reductions Between Optimization Objectives
Zeyuan Allen-Zhu and Elad Hazan
Proceedings of 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)

52. The Limits of Learning with Missing Data
Brian Bullins* , Elad Hazan and Tomer Koren
Proceedings of 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)

53. Faster Convex Optimization: Simulated Annealing with an Efficient Universal Barrier
Jacob Abernethy and Elad Hazan
Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)

54. Faster Eigenvector Computation via Shift-and-Invert Preconditioning
Dan Garber*, Elad Hazan, Chi Jin, Sham Kadade, Cameron Musco, Praneeth Netrapalli and Aaron Sidford
Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)

55. On Graduated Optimization for Stochastic Non-Convex Problems
Elad Hazan, Kfir Levi* and Shai Shalev-Shwartz
Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)

56. Variance Reduction for Faster Non-Convex Optimization
Zeyuan Allen-Zhu and Elad Hazan
Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)

57. Variance-Reduced and Projection-Free Stochastic Optimization
Elad Hazan and Haipeng Luo
Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)

58. Online Learning with Low Rank Experts
Elad Hazan, Tomer Koren*, Roi Livni and Yishay Mansour
Proceedings of the 29th Conference on Learning Theory (COLT 2016)

59. The Computational Power of Optimization in Online Learning
Elad Hazan* and Tomer Koren
48th Annual Symposium on the Theory of Computing (STOC 2016)

60. Online Learning for Adversaries with Memory: Price of Past Mistakes
Oren Anava*, Elad Hazan and Shie Mannor
Proceedings of 29th Annual Conference on Neural Information Processing Systems (NIPS 2015)

61. Online Gradient Boosting
Alina Beygelzimer, Elad Hazan, Satyen Kale and Haipeng Luo
Proceedings of 29th Annual Conference on Neural Information Processing Systems (NIPS 2015)

62. Beyond Convexity: Stochastic Quasi-Convex Optimization
Elad Hazan, Kfir Levi* and Shai Shalev-Shwartz
Proceedings of 29th Annual Conference on Neural Information Processing Systems (NIPS 2015)

63. Classification with Low Rank and Missing Data
Elad Hazan, Roi Livni and Yishay Mansour
to appear 32nd International Conference on Machine Learning (ICML 2015)

64. Faster Rates for the Frank-Wolfe Method over Strongly-Convex Sets
Dan Garber* and Elad Hazan
to appear 32nd International Conference on Machine Learning (ICML 2015)

65. Online Learning of Eigenvectors
Dan Garber*, Elad Hazan and Tengyu Ma
to appear 32nd International Conference on Machine Learning (ICML 2015)

66. Online Time Series Prediction with Missing Data
Oren Anava*, Elad Hazan and Assaf Zeevi
to appear 32nd International Conference on Machine Learning (ICML 2015)

67. Bandit Convex Optimization: Towards Tight Bounds
Elad Hazan and Kfir Y. Levy*
28th Annual Conference on Neural Information Processing Systems (NIPS 2014)

68. The Blinded Bandit: Learning with Adaptive Feedback
Ofer Dekel, Elad Hazan and Tomer Koren*
28th Annual Conference on Neural Information Processing Systems (NIPS 2014)

69. Logistic Regression: Tight Bounds for Stochastic and Online Optimization
Elad Hazan, Tomer Koren* and Kfir Y. Levy*
The 27th Conference on Learning Theory (COLT 2014)

70. Volumetric Spanners: an Efficient Exploration Basis for Learning
Elad Hazan, Zohar Shay Karnin and Raghu Meka
The 27th Conference on Learning Theory (COLT 2014)

71. Hard-Margin Active Linear Regression
Elad Hazan, Zohar Shay Karnin
The 31st International Conference on Machine Learning (ICML 2014)

72. Online Learning for Time Series Prediction
Oren Anava*, Elad Hazan, Shie Mannor and Ohad Shamir
The 26th Conference on Learning Theory (COLT 2013)

73. Better Rates for Any Adversarial Deterministic MDP
Ofer Dekel and Elad Hazan
The 30th International Conference on Machine Learning (ICML 2013)

74. Playing Non-linear Games with Linear Oracles
Dan Garber* and Elad Hazan
54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)

75. Projection-free Online Learning.
E. Hazan and S. Kale
The 29th International Conference on Machine Learning (ICML 2012)

76. (weak) Calibration is Computationally Hard.
E. Hazan and S. Kakade
The 25th conference on learning theory (COLT 2012)

77. Near-Optimal Algorithms for Online Matrix Prediction.
E. Hazan, S. Kale and S. Shalev-Shwartz
The 25th conference on learning theory (COLT 2012)

78. Linear Regression with Limited Observation.
E. Hazan and T. Koren*
Proceedings of the 29th International Conference on Machine Learning (ICML), 2012, pages 807–814
ICML 2012 Best Student Paper Runner Up

79. A Polylog Pivot Steps Simplex Algorithm for Classification
E. Hazan and Z. Karnin
Twenty-Sixth Annual Conference on Neural Information Processing Systems (NIPS 2012)

80. Blackwell Approachability and No-Regret Learning are Equivalent.
J. Abernethy, P. Bartlet and Elad Hazan
The 24th Annual Conference on Learning Theory (COLT 2011)

81. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization.
Elad Hazan and Satyen Kale
The 24th Annual Conference on Learning Theory (COLT 2011)

82. Newtron: an Efficient Bandit algorithm for Online Multiclass Prediction
E. Hazan and S. Kale
Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS 2011)

83. Beating SGD: Learning SVMs in Sublinear Time
E. Hazan, T. Koren* and N. Srebro
Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS 2011)

84. Approximating Semidefinite Programs in Sublinear Time
D. Garber* and E. Hazan
Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS 2011)

85. Sublinear Optimization for Machine Learning.
Ken Clarkson, Elad Hazan, and David Woodruff
The 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010).

86. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization.
John Duchi, Elad Hazan, and Yoram Singer
The 23rd Annual Conference on Learning Theory, (COLT 2010).

87. Learning rotations with little regret.
Elad Hazan, Satyen Kale and Manfred Warmuth
The 23rd Annual Conference on Learning Theory, (COLT 2010).

88. Better algorithms for benign bandits.
Elad Hazan and Satyen Kale
Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2009, pages 38-47.

89. How hard is it to approximate the best Nash equilibrium?
Elad Hazan and Robert Krauthgamer
Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2009, pages 720-727.

90. On Stochastic and Worst-case Models for Investing.
Elad Hazan and Satyen Kale
The 23rd Twenty-Second Annual Conference on Neural Information Processing Systems (NIPS 2009). **accepted for full oral presentation** (22 papers out of 1100 submissions)

91. Beyond Convexity: Online Submodular Minimization.
Elad Hazan and Satyen Kale
The 23rd Twenty-Second Annual Conference on Neural Information Processing Systems (NIPS 2009).

92. Efficient learning algorithms for changing environments.
Elad Hazan and C. Seshadhri
The 26th International Conference on Machine Learning (ICML 2009).

93. Sparse Approximate Solutions to Semidefinite Programs.
Elad Hazan
LATIN 2008: Theoretical Informatics, 8th Latin American Symposium, pages 306-316.

94. Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs.
Elad Hazan and Satyen Kale
Proceedings of the 21st Annual Conference on Learning Theory, (COLT) 2008, pages 57-68.

95. Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization.
Jacob Abernethy, Elad Hazan and Alexander Rakhlin
Proceedings of the 21st Annual Conference on Learning Theory, (COLT) 2008, pages 263-274.
IBM Pat Goldberg Memorial Best Paper Award
Machine Learning Journal award for best student paper

96. Adaptive Online Gradient Descent.
Peter Bartlett, Elad Hazan and Alexander Rakhlin
Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, (NIPS) 2007. **accepted for full oral presentation** (25 papers out of 1000 submissions)

97. Computational Equivalence of Fixed Points and No Regret Algorithms, and Convergence to Equilibria.
Elad Hazan and Satyen Kale
Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, (NIPS) 2007.

98. Online Learning with Prior Information
Elad Hazan and Nimrod Megiddo
Proceedings of 20th Annual Conference on Learning Theory, (COLT) 2007, pages 499-513.

99. A Fast Random Sampling Algorithm for Sparsifying Matrices.
Sanjeev Arora, Elad Hazan, and Satyen Kale
9th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems - APPROX 2006

100. Logarithmic Regret Algorithms for Online Convex Optimization
Elad Hazan, Adam Kalai, Satyen Kale and Amit Agarwal
Proceedings of 19'th Annual Conference on Learning Theory, (COLT) 2006.

101. Algorithms for Portfolio Management based on the Newton Method
Amit Agarwal, Elad Hazan, Satyen Kale and Robert E. Schapire
23rd International Conference of Machine Learning, (ICML) 2006.

102. HAPLOFREQ - Estimating Haplotype Frequencies Efficiently
Eran Halperin and Elad Hazan
Ninth Annual International Conference on Research in Computational Molecular Biology (RECOMB 2005)

103. Analysis and Algorithms for Content-based Event Matching
Elad Hazan, Satyen Kale, Fengyun Cao and Jaswinder Pal Singh
4th International Workshop on Distributed Event-Based Systems (DEBS 2005)

104. Fast Algorithms for Approximate Semidefinite Programming using the Multiplicative Weights Update Method
Sanjeev Arora, Elad Hazan and Satyen Kale
46th Symposium on Foundations of Computer Science (FOCS 2005)

105. On Non-Approximability for Quadratic Programs
Sanjeev Arora, Eli Berger, Elad Hazan , Guy Kindler, Muli Safra
46th Symposium on Foundations of Computer Science (FOCS 2005)

106. $O(\sqrt{\log n})$ approximation to SPARSEST CUT in $\tilde{O}(n^2)$ time
Sanjeev Arora, Elad Hazan and Satyen Kale
appeared in 45th Symposium on Foundations of Computer Science (FOCS 2004)
To appear in SIAM Journal on Computing

107. On the Complexity of Approximating k-Set Packing
Elad Hazan, Muli Safra and Oded Schwartz
6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2003)

Refereed papers in professional journals

108. Second-Order Stochastic Optimization for Machine Learning in Linear Time
Naman Agarwal*, Brian Bullins*, Elad Hazan
Journal of Machine Learning Research 18: 116:1-116:40 (2017)

109. Near-Optimal Algorithms for Online Matrix Prediction
Elad Hazan, Satyen Kale and Shai Shalev-Shwartz
SIAM Journal on Computing (SICOMP) 46(2): 744-773 (2017)

110. Volumetric Spanners: an Efficient Exploration Basis for Learning
Elad Hazan and Zohar Shay Karnin
Journal of Machine Learning Research (JMLR) 17(119):1–34, (2016)

111. A Linearly Convergent Variant of the Conditional Gradient Algorithm under Strong Convexity, with Applications to Online and Stochastic Optimization
Dan Garber and Elad Hazan
SIAM Journal on Optimization 26(3): 1493-1528 (2016)

112. Learning rotations with little regret.
Elad Hazan, Satyen Kale and Manfred Warmuth
Machine Learning 104(1): 129-148 (2016)

113. Sublinear time algorithms for approximate semidefinite programming
Dan Garber* and Elad Hazan
Mathematical Programming (MP) 158(1-2): 329-361 (2016)

114. A Linear-Time Algorithm for Trust Region Problems
Elad Hazan and Tomer Koren*
Mathematical Programming (MP) 158(1-2): 363-381 (2016)

115. Oracle-Based Robust Optimization via Online Learning
Aharon Ben-Tal, Elad Hazan, Tomer Koren* and Shie Mannor
Operations Research 63(3): 628-638 (2015)

116. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization.
Elad Hazan and Satyen Kale
Journal of Machine Learning Research (JMLR), v. 15(1) p. 2489-2512 (2014)

117. Universal MMSE Filtering With Logarithmic Adaptive Regret
Dan Garber* and Elad Hazan
IEEE Transactions on Signal Processing, volume 61(7): 1595-1604, 2013

118. An Online Portfolio Selection Algorithm with Regret Logarithmic in Price Variation.
Elad Hazan and Satyen Kale
Mathematical Finance, published online 2 Nov 2012

119. Sublinear Optimization for Machine Learning.
Ken Clarkson, Elad Hazan and David Woodruff
Journal of the ACM (JACM) Volume 59 Issue 5, October 2012
IBM Pat Goldberg Memorial Best Paper Award

120. Interior-Point Methods for Full-Information and Bandit Online Learning.
Jacob Abernethy, Elad Hazan and Alexander Rakhlin
IEEE Transactions on Information Theory (IEEE-IT) 58(7): 4164-4175 (2012)

121. Online Submodular Minimization.
Elad Hazan and Satyen Kale
Journal of Machine Learning Research (JMLR), 13 (2012) 2903-2922

122. The Multiplicative Weights Update Method: A Meta-Algorithm and Applications.
Sanjeev Arora, Elad Hazan and Satyen Kale
Theory of Computing (ToC) Volume 8 p. 121-164 (2012)

123. How hard is it to approximate the best Nash equilibrium?
Elad Hazan and Robert Krauthgamer
SIAM Journal on Computing (SICOMP), v. 40(1): 79-91 (2011)

124. Better algorithms for benign bandits.
Elad Hazan and Satyen Kale
Journal of Machine Learning Research (JMLR), v. 12(Apr):1287-1311, 2011.

125. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization
John Duchi, Elad Hazan and Yoram Singer
Journal of Machine Learning Research (JMLR), v. 12(Jul):2121-2159, 2011.

126. Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs
Elad Hazan and Satyen Kale
Machine Learning Journal: Volume 80, Issue 2 (2010), Page 165-188

127. $O(\sqrt{\log n})$ approximation to SPARSEST CUT in $\tilde{O}(n^2)$ time
Sanjeev Arora, Elad Hazan and Satyen Kale
SIAM Journal on Computing Vol.39, No.5, pages 1748-1771 (2010)

128. Logarithmic Regret Algorithms for Online Convex Optimization
Elad Hazan, Amit Agarwal and Satyen Kale
Machine Learning Journal Volume 69 , Issue 2-3 Pages: 169 - 192, December 2007
(invited to special issue)

129. HAPLOFREQ - Estimating Haplotype Frequencies Efficiently
Eran Halperin and Elad Hazan
Journal of Computational Biology, March 2006, Vol. 13, No. 2

130. On the Complexity of Approximating k-Set Packing
Elad Hazan, Muli Safra and Oded Schwartz
Computational Complexity, v.15 n.1, p.20-39 May 2006.

Thesis

- Efficient Algorithms for Online Convex Optimization and Their Applications
Elad Hazan
Ph.D thesis, Department of Computer Science, Princeton University (2006).

- On the Hardness of Approximating k-dimensional matching
Elad Hazan
M.Sc thesis, Department of Computer Science, Tel Aviv University (2002).

Patents

- Patent pending, 2023
  Spectral state space models
  Inventors: Naman Agarwal, Daniel Suo, Xinyi Chen and Elad Hazan

- Patent US 62/518,682, pending, 2017
  Method for online learning of linear dynamical systems
  Inventors: Elad Hazan, Karan Singh and Cyril Zhang

- Patent US7730000, issued, issue date Jun 1, 2010
  Method for machine learning with state information
  Inventors: Elad Hazan and Nimrod Megiddo

- Patent 7870082, issued, issue date Jan 11, 2011
  Method for machine learning using online convex optimization problem solving with minimum regret
  Inventors: Elad Hazan and Nimrod Megiddo

- Patent 8494994, issued, issue date July 23, 2013
  Fast Adaptation in Real Time Systems
  Inventors: Elad Hazan and Nimrod Megiddo

- Filed by IBM Reseach
  Sublinear Time Algorithms for Classification
  Inventors: Kenneth Clarkson, Elad Hazan and David Woodruff

Technical Reports

- Fast and Simple PCA via Convex Optimization.
  Dan Garber* and Elad Hazan
  arXiv 1509.05647

- A Linearly Convergent Conditional Gradient Algorithm with Applications to Online and Stochastic Optimization
  Dan Garber* and Elad Hazan
  CoRR abs/1301.4666, 2013

- Approximating Quadratic Programs with Positive Semidefinite Constraints
  Elad Hazan and Satyen Kale
  Princeton University Technical Report TR-804-07, 2007

- The "Arrangement Method" for Linear Programming is Equivalent to the Phase-One Method.
  Elad Hazan and Nimrod Megiddo
  IBM Research Report, February 2007.

- Approximate Convex Optimization by Online Game Playing.
  Elad Hazan
  arXiv.org Cornel Library report, 0610119v1, 2006

- Efficient Algorithms for Online Game Playing and Universal Portfolio Management
  Amit Agarwal and Elad Hazan
  Electronic Colloquium on Computational Complexity TR06-033, 2006