

Haipeng Luo

University of Southern California
Computer Science Department
941 Bloom Walk, Los Angeles, CA 90089
E-mail: haipengl@usc.edu
Homepage: <https://haipeng-luo.net>

RESEARCH INTEREST

Developing practical machine learning algorithms with strong theoretical guarantees, with a focus on online decision making, bandit problems, reinforcement learning, learning in games, and so on.

EXPERIENCE

07/2017~present Assistant Professor, Computer Science Department, University of Southern California
07/2016~06/2017 Postdoctoral Researcher, Microsoft Research, NYC
03/2016~04/2016 Visiting student, CWI (Centrum Wiskunde & Informatica), Amsterdam
06/2015~08/2015 Research intern, Microsoft Research, NYC
06/2014~08/2014 Research intern, Yahoo Labs, NYC
06/2013~08/2013 Research intern, AT&T Labs, San Francisco

EDUCATION

2011~2016 PhD, Department of Computer Science, Princeton University
Advisor: Robert Schapire Thesis: Optimal and Adaptive Online Learning
M.A. received in 09/2013
2007~2011 B.S., Department of Computer Science, Peking University (double major in Mathematics)

AWARDS

Best paper award, COLT 2021 (with Chen-Yu Wei).
NSF CAREER Award, 2020.
Google Faculty Research Award, 2020.
Best student paper award, COLT 2018 (with Dylan Foster, Satyen Kale, Mehryar Mohri, Karthik Sridharan).
NSF CRII Award, 2018.
Best paper award, NIPS 2015 (with Vasilis Syrgkanis, Alekh Agarwal and Robert E. Schapire).
Best paper award, ICML 2015 (with Alina Beygelzimer and Satyen Kale).
Wu Prize for Excellence, 2015 (awarded to engineering students who have performed at the highest level).
Google Spotlight Presentation Awards, Annual Machine Learning Symposium, 2014, 2015, 2017.
ACM Doctoral Dissertation Award Nominee (one per institution).
Symantec Fellowship finalist, 2015.
Tung OOCL Scholarship, 2010.
Suzhou Industrial Park Scholarship, 2009.
May Fourth Scholarship, 2008.
Silver Medal, 33rd ACM International Collegiate Programming Contest, Asia Regional, 2008.
Bronze Medal, ACM International Collegiate Programming Contest, Beijing Sub-regional, 2008.

GRANTS

NSF CAREER Award, 2020. Total amount: \$499,863.
Google Faculty Research Award, 2020. Total amount: \$50,460.
NSF CISE Research Initiation Initiative (CRII), 2018. Total amount: \$175,000.

PUBLICATIONS

[51] Tiancheng Jin, Longbo Huang, and Haipeng Luo.

The Best of Both Worlds: Stochastic and Adversarial Episodic MDPs with Unknown Transition.

To appear in Neural Information Processing Systems 34 (NeurIPS 2021, oral, <1% acceptance rate).

[50] Haipeng Luo, Chen-Yu Wei, and Chung-Wei Lee.

Policy Optimization in Adversarial MDPs: Improved Exploration via Dilated Bonuses.

To appear in Neural Information Processing Systems 34 (NeurIPS 2021).

[49] Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, and Haipeng Luo.

Implicit Finite-Horizon Approximation and Efficient Optimal Algorithms for Stochastic Shortest Path.

To appear in Neural Information Processing Systems 34 (NeurIPS 2021).

[48] Chung-Wei Lee, Christian Kroer, and Haipeng Luo.

Last-iterate Convergence in Extensive-Form Games.

To appear in Neural Information Processing Systems 34 (NeurIPS 2021).

[47] Chen-Yu Wei and Haipeng Luo.

Non-stationary Reinforcement Learning without Prior Knowledge: An Optimal Black-box Approach.

In Proceedings of the 34th Conference on Learning Theory (COLT 2021). ***Best Paper Award.***

[46] Chen-Yu Wei, Haipeng Luo, and Liyu Chen.

Impossible Tuning Made Possible: A New Expert Algorithm and Its Applications.

In Proceedings of the 34th Conference on Learning Theory (COLT 2021).

[45] Liyu Chen, Haipeng Luo, and Chen-Yu Wei.

Minimax Regret for Stochastic Shortest Path with Adversarial Costs and Known Transition.

In Proceedings of the 34th Conference on Learning Theory (COLT 2021).

[44] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo.

Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games.

In Proceedings of the 34th Conference on Learning Theory (COLT 2021).

[43] Liyu Chen and Haipeng Luo.

Finding the Stochastic Shortest Path with Low Regret: The Adversarial Cost and Unknown Transition Case.

In Proceedings of the 38th International Conference on Machine Learning (ICML 2021).

[42] Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, Mengxiao Zhang, and Xiaojin Zhang

Achieving Near Instance-Optimality and Minimax-Optimality in Stochastic and Adversarial Linear Bandits Simultaneously.

In Proceedings of the 38th International Conference on Machine Learning (ICML 2021).

[41] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo.

Linear Last-iterate Convergence in Constrained Saddled-point Optimization.

In Proceedings of the 9th International Conference on Learning Representations (ICLR 2021)

[40] Chen-Yu Wei, Mehdi Jafarnia-Jahromi, Haipeng Luo, and Rahul Jain.

Learning Infinite-horizon Average-reward MDPs with Linear Function Approximation.

In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)

[39] Yining Chen, Haipeng Luo, Tengyu Ma, and Chicheng Zhang.

Active Online Learning with Hidden Shifting Domains.

In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)

[38] Chen-Yu Wei, Ehsan Emamjomeh-Zadeh, Haipeng Luo, and David Kempe.

Adversarial Online Learning with Changing Action Sets: Efficient Algorithms with Approximate Regret Bounds.

In Proceedings of the 32nd International Conference on Algorithmic Learning Theory (ALT 2021)

[37] Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang.

Bias no more: high-probability data-dependent regret bounds for adversarial bandits and MDPs.

In Advances in Neural Information Processing Systems 33 (NeurIPS 2020, oral, 1.11% acceptance rate).

[36] Tiancheng Jin and Haipeng Luo.

Simultaneously Learning Stochastic and Adversarial Episodic MDPs with Known Transition.

In Advances in Neural Information Processing Systems 33 (NeurIPS 2020, spotlight, 3% acceptance rate).

[35] Dirk van der Hoeven, Ashok Cutkosky, and Haipeng Luo.

Comparator-Adaptive Convex Bandits.

In Advances in Neural Information Processing Systems 33 (NeurIPS 2020).

[34] Miroslav Dudik, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan.

Oracle-Efficient Online Learning and Auction Design.

Journal of ACM (JACM 2020).

[33] Chen-Yu Wei, Haipeng Luo, and Alekh Agarwal.

Taking a Hint: How to Leverage Loss Predictors in Contextual Bandits?

In Proceedings of the 33rd Conference on Learning Theory (COLT 2020).

[32] Chung-Wei Lee, Haipeng Luo, and Mengxiao Zhang.

A Closer Look at Small-loss Bounds for Bandits with Graph Feedback.

In Proceedings of the 33rd Conference on Learning Theory (COLT 2020).

[31] Dylan Foster, Akshay Krishnamurthy, and Haipeng Luo.

Open Problem: Model Selection for Contextual Bandits.

In Proceedings of the 33rd Conference on Learning Theory (COLT 2020).

[30] Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu.

Learning Adversarial Markov Decision Processes with Bandit Feedback and Unknown Transition.

In Proceedings of the 37th International Conference on Machine Learning (ICML 2020).

[29] Chen-Yu Wei, Mehdi Jafarnia-Jahromi, Haipeng Luo, Hiteshi Sharma, and Rahul Jain.

Model-free Reinforcement Learning in Infinite-horizon Average-reward Markov Decision Processes.

In Proceedings of the 37th International Conference on Machine Learning (ICML 2020).

[28] Yifang Chen, Alex Cuellar, Haipeng Luo, Jignesh Modi, Heramb Nemlekar, and Stefanos Nikolaidis.

Fair Contextual Multi-Armed Bandits: Theory and Experiments.

In Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI 2020).

[27] Dylan J. Foster, Akshay Krishnamurthy, and Haipeng Luo.

Model Selection for Contextual Bandits.

In Advances in Neural Information Processing Systems 32 (NeurIPS 2019, Spotlight).

[26] Kai Zheng, Haipeng Luo, Ilias Diakonikolas, and Liwei Wang.

Equipping Experts/Bandits with Long-term Memory.

In Advances in Neural Information Processing Systems 32 (NeurIPS 2019).

[25] Dylan J. Foster, Spencer Greenberg, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan.

Hypothesis Set Stability and Generalization.

In Advances in Neural Information Processing Systems 32 (NeurIPS 2019).

[24] Sébastien Bubeck, Yuanzhi Li, Haipeng Luo, and Chen-Yu Wei.

Improved Path-length Regret bounds for Bandits.

In Proceedings of the 32nd Conference on Learning Theory (COLT 2019).

[23] Yifang Chen, Chung-Wei Lee, Haipeng Luo, and Chen-Yu Wei.

A New Algorithm for Non-stationary Contextual Bandits: Efficient, Optimal, and Parameter-free.

In Proceedings of the 32nd Conference on Learning Theory (COLT 2019).

[22] Peter Auer, Yifang Chen, Pratik Gajane, Chung-Wei Lee, Haipeng Luo, Ronald Ortner, and Chen-Yu Wei.

Achieving Optimal Dynamic Regret for Non-stationary Bandits without Prior Information.

In Proceedings of the 32nd Conference on Learning Theory (COLT 2019), extended abstract.

[21] Julian Zimmert, Haipeng Luo, and Chen-Yu Wei.

Beating Stochastic and Adversarial Semi-Bandits Optimally and Simultaneously.

In Proceedings of the 36th International Conference on Machine Learning (ICML 2019, long talk).

[20] Haipeng Luo, Chen-Yu Wei, and Kai Zheng.

Efficient Online Portfolio with Logarithmic Regret.

In Advances in Neural Information Processing Systems 31 (NeurIPS 2018, spotlight).

[19] Chen-Yu Wei and Haipeng Luo.

More Adaptive Algorithms for Adversarial Bandits.

In Proceedings of the 31st Conference on Learning Theory (COLT 2018).

- [18] Haipeng Luo, Chen-Yu Wei, Alekh Agarwal, and John Langford.
Efficient Contextual Bandits in Non-stationary Worlds.
In Proceedings of the 31st Conference on Learning Theory (COLT 2018).
- [17] Dylan J. Foster, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan
Logistic Regression: The Importance of Being Improper.
In Proceedings of the 31st Conference on Learning Theory (COLT 2018). ***Best Student Paper Award.***
- [16] Dylan J. Foster, Alekh Agarwal, Miroslav Dudik, Haipeng Luo, and Robert E. Schapire
Practical Contextual Bandits with Regression Oracles.
In Proceedings of the 35th International Conference on Machine Learning (ICML 2018).
- [15] Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E. Schapire.
Corralling a Band of Bandit Algorithms.
In Proceedings of the 30th Conference on Learning Theory (COLT 2017).
- [14] Miroslav Dudik, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan.
Oracle-Efficient Online Learning and Auction Design.
In the 58th Annual Symposium on Foundations of Computer Science (FOCS 2017).
- [13] Alekh Agarwal, Akshay Krishnamurthy, John Langford, Haipeng Luo, and Robert E. Schapire.
Open Problem: First-Order Regret Bounds for Contextual Bandits.
In Proceedings of the 30th Conference on Learning Theory (COLT 2017).
- [12] Haipeng Luo, Alekh Agarwal, Nicolò Cesa-Bianchi, and John Langford.
Efficient Second Order Online Learning via Sketching.
In Advances in Neural Information Processing Systems 29 (NIPS 2016).
- [11] Vasilis Syrgkanis, Haipeng Luo, Akshay Krishnamurthy, and Robert E. Schapire.
Improved Regret Bounds for Oracle-Based Adversarial Contextual Bandits.
In Advances in Neural Information Processing Systems 29 (NIPS 2016).
- [10] Elad Hazan and Haipeng Luo.
Variance-Reduced and Projection-Free Stochastic Optimization.
In Proceedings of the 33rd International Conference on Machine Learning (ICML 2016).
- [9] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire.
Fast Convergence of Regularized Learning in Games.
In Advances in Neural Information Processing Systems 28 (NIPS 2015). ***Best Paper Award.***
- [8] Alina Beygelzimer, Satyen Kale, and Haipeng Luo
Optimal and Adaptive Algorithms for Online Boosting.
In Proceedings of the 32nd International Conference on Machine Learning (ICML 2015). ***Best Paper Award.***
- [7] Alina Beygelzimer, Elad Hazan, Satyen Kale, and Haipeng Luo.

Online Gradient Boosting.

In Advances in Neural Information Processing Systems 28 (NIPS 2015).

[6] Haipeng Luo and Robert E. Schapire.

Achieving All with No Parameters: AdaNormalHedge.

In Proceedings of the 28th Conference on Learning Theory (COLT 2015).

[5] Haipeng Luo and Robert E. Schapire.

A Drifting-Games Analysis for Online Learning and Applications to Boosting.

In Advances in Neural Information Processing Systems 27 (NIPS 2014).

[4] Haipeng Luo, Patrick Haffner, and Jean-Francois Paiement.

Accelerated Parallel Optimization Methods for Large Scale Machine Learning.

In the 7th NIPS Workshop on Optimization for Machine Learning, 2014.

[3] Haipeng Luo and Robert E. Schapire.

Towards Minimax Online Learning with Unknown Time Horizon.

In Proceedings of the 31st International Conference on Machine Learning (ICML 2014).

[2] Weijia Song, Zhen Xiao, Qi Chen, and Haipeng Luo.

Adaptive Resource Provisioning for the Cloud Using Online Bin Packing.

IEEE Transactions on Computers, 63:2647-2660, 2013.

[1] Zhen Xiao, Qi Chen, and Haipeng Luo.

Automatic Scaling of Internet Applications for Cloud Computing Services.

IEEE Transactions on Computers, 63:1111-1123, 2012.

SERVICES

Editorial board member: Journal of Machine Learning Research (JMLR)

Program Committee Member: COLT 2021, 2020, 2019, ICML 2020, 2019, NeurIPS 2021, 2020, 2019, 2018, ALT 2022, 2021, 2020, 2019, 2017

Workshop Proposal Reviewer: ICML 2021

Conference Reviewer: FOCS 2020, ICML 2018, 2017, 2016, NeurIPS 2021, 2017, 2016, AISTATS 2019, 2017, 2016, COLT 2018, 2017, 2016, SODA 2022, 2017, 2016, IJCAI 2015, 2013

Journal Reviewer: JACM, JMLR, Machine Learning, Mathematical Programming, Theoretical Computer Science, Operations Research, Mathematics of Operations Research, IEEE Transactions on Information Theory, IEEE Transactions on Signal Processing, IEEE Transactions on Information Forensics & Security, IEEE Transactions on Automatic Control, Management Science

Proposal Panelist: NSF panels

Workshop Organizer: KDD 2021, “Multi-Armed Bandits and Reinforcement Learning: Advancing Decision Making in E-Commerce and Beyond”.

MENTORING

Current PhD students: Liyu Chen, Tiancheng Jin, Chung-Wei Lee, Chen-Yu Wei, and Mengxiao Zhang

Past undergraduate/master students: Yifang Chen (now graduate student at University of Washington)

Past visiting students: Dirk van der Hoeven (now postdoc at University of Milan), Kai Zheng (now researcher at Kuaishou), Julian Ulf Zimmert (now research scientist at Google).

TEACHING

Instructor, Machine Learning (CSCI 567), USC, Fall 2021.

Instructor, Machine Learning (CSCI 567), USC, Fall 2020.

Instructor, Theoretical Machine Learning (CSCI 699), USC, Fall 2019.

Instructor, Machine Learning (CSCI 567), USC, Fall 2018.

Instructor, Introduction to Online Optimization/Learning, Peking University, Summer 2018.

Instructor, Introduction to Online Learning (CSCI 699), USC, Fall 2017.

Teaching assistant, Artificial Intelligence (COS 402), Princeton, Fall 2012, 2013.

SELECTED INVITED TALKS

The best of both worlds: stochastic and adversarial episodic MDPs with unknown transition. DeLTA learning theory seminar, University of Copenhagen (online). Aug 2021.

From bandits to RL: recent advances in theory and algorithms for reinforcement learning. USC-THU symposium (online). May 2021.

From bandits to MDPs: optimally and adaptively learning episodic MDPs with adversarial losses. Caltech (online), Jun 2020; Berkeley (online), Oct 2020; John Hopkins University (online), Oct 2020.

Leveraging Loss Predictors in Contextual Bandits. Princeton (online), Apr 2020.

Beating Stochastic and Adversarial Semi-bandits Optimally and Simultaneously. Microsoft Research, Redmond, Aug 2019; Google Machine Learning Theory Workshop, Sep 2019.

Logistic Regression: The Importance of Being Improper. TTIC Recent Trends in Clustering and Classification Workshop, Sep 2019.

Recent Advances in (Contextual) Multi-armed Bandits. Alibaba, Aug 2019.

More Adaptive Algorithms for Adversarial Bandits. Microsoft Research, Redmond, May 2018.

Efficient Second Order Online Learning via Sketching. SIAM Conference on Optimization, May 2017.

Corralling a Band of Bandit Algorithms. NYAS 11th Annual Machine Learning Symposium, Mar 2017; Theoretical Foundations for Learning from Easy Data, Oct 2016.

Optimal and Adaptive Algorithms for Online Boosting. IJCAI sister conference best paper track, Jul 2016.

Achieving All with No Parameters: AdaNormalHedge. NYAS 9th Annual Machine Learning Symposium, Mar 2014.

Accelerated Parallel Optimization Methods for Large Scale Machine Learning. NYAS 8th Annual Machine Learning Symposium, Mar 2014.

Towards Minimax Online Learning with Unknown Time Horizon. Microsoft Research, Redmond, Feb 2014.