

Chicheng Zhang

Curriculum Vitæ

CURRENT POSITION

Assistant Professor 2019.8 - Now
Department of Computer Science

Affiliate Member
Graduate Interdisciplinary Program in Statistics 2019.11 - Now
Graduate Interdisciplinary Program in Applied Mathematics 2020.2 - Now

University of Arizona
Tucson, AZ
Research interests: interactive machine learning, machine learning theory
Personal webpage: <https://zcc1307.github.io/>

EDUCATION

PhD, Computer Science 2012.9-2017.9
UC San Diego, La Jolla, CA
Advisor: Kamalika Chaudhuri
Thesis: Active learning and confidence-rated prediction

Master of Science, Computer Science 2012.9-2015.6
UC San Diego, La Jolla, CA

Bachelor of Science, Machine Intelligence, School of EECS 2008.9-2012.7
Peking University, Beijing, China

Second Degree Certificate, Mathematics and Applied Mathematics 2008.9-2012.7
Peking University, Beijing, China

EXPERIENCE

Postdoctoral Researcher 2017.9-2019.6
Machine Learning Group
Microsoft Research, New York City

Research Intern 2016.6-2016.9
Yahoo! Research, New York City
Supervisor: Dr. Alina Beygelzimer and Dr. Francesco Orabona

Research Intern 2015.6-2015.9
Yahoo! Labs, New York City
Supervisor: Dr. Alina Beygelzimer

Research Assistant 2012.9-2017.8
Department of Computer Science and Engineering, UC San Diego
Supervisor: Prof. Kamalika Chaudhuri

Undergraduate Research Assistant 2010.6-2012.6
Department of Machine Intelligence, Peking University
Supervisor: Prof. Liwei Wang

Software Testing Intern 2011.7-2011.8
MicroVu Co. China

PUBLICATIONS

CONFERENCE PAPERS

Chicheng Zhang and Yinan Li. Improved algorithms for efficient active learning halfspaces with massart and tsybakov noise. In *Proceedings of Thirty Fourth Conference on Learning Theory (COLT)*, volume 134 of *Proceedings of Machine Learning Research*, pages 4526–4527. PMLR, 15–19 Aug 2021.

Zhi Wang*, Chicheng Zhang*, Manish Kumar Singh, Laurel Riek, and Kamalika Chaudhuri. Multitask bandit learning through heterogeneous feedback aggregation. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 130 of *Proceedings of Machine Learning Research*, pages 1531–1539. PMLR, 13–15 Apr 2021.

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)*, volume 130 of *Proceedings of Machine Learning Research*, pages 2053–2061. PMLR, 13–15 Apr 2021.

Jie Shen and Chicheng Zhang. Attribute-efficient learning of halfspaces with malicious noise: Near-optimal label complexity and noise tolerance. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory (ALT)*, volume 132 of *Proceedings of Machine Learning Research*, pages 1072–1113. PMLR, 16–19 Mar 2021.

Chicheng Zhang, Jie Shen, and Pranjal Awasthi. Efficient active learning of sparse halfspaces with arbitrary bounded noise. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 7184–7197. Curran Associates, Inc., 2020.

Maryam Majzoubi, Chicheng Zhang, Rajan Chari, Akshay Krishnamurthy, John Langford, and Aleksandrs Slivkins. Efficient contextual bandits with continuous actions. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 349–360. Curran Associates, Inc., 2020.

Kwang-Sung Jun and Chicheng Zhang. Crush optimism with pessimism: Structured bandits beyond asymptotic optimality. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6366–6376. Curran Associates, Inc., 2020.

Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.

Akshay Krishnamurthy*, John Langford*, Aleksandrs Slivkins*, and Chicheng Zhang*. Contextual bandits with continuous actions: Smoothing, zooming, and adapting. In *Proceedings of the Thirty-Second Conference on Learning Theory (COLT)*, volume 99 of *Proceedings of Machine Learning Research*, pages 2025–2027. PMLR, 2019.

Alina Beygelzimer*, David Pal*, Balazs Szorenyi*, Devanathan Thiruvengatathari*, Chen-Yu Wei*, and Chicheng Zhang*. Bandit multiclass linear classification: Efficient algorithms for the separable case. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 624–633. PMLR, 2019.

Chicheng Zhang, Alekh Agarwal, Hal Daumé III, John Langford, and Sahand Negahban. Warm-starting contextual bandits: Robustly combining supervised and bandit feedback. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 7335–7344. PMLR, 2019.

Chicheng Zhang. Efficient active learning of sparse halfspaces. In *Proceedings of the 31st Conference on Learning Theory (COLT)*, volume 75 of *Proceedings of Machine Learning Research*, pages 1856–1880. PMLR, 06–09 Jul 2018.

Songbai Yan and Chicheng Zhang. Revisiting perceptron: Efficient and label-optimal learning of halfspaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pages 1056–1066. Curran Associates, Inc., 2017.

Alina Beygelzimer*, Francesco Orabona*, and Chicheng Zhang*. Efficient online bandit multiclass learning with $\tilde{O}(\sqrt{T})$ regret. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pages 488–497. PMLR, 06–11 Aug 2017.

Alina Beygelzimer*, Daniel J Hsu*, John Langford*, and Chicheng Zhang*. Search improves label for active learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, pages 3342–3350. Curran Associates, Inc., 2016.

Chicheng Zhang and Kamalika Chaudhuri. The extended littlestone’s dimension for learning with mistakes and abstentions. In *29th Annual Conference on Learning Theory (COLT)*, volume 49 of *Proceedings of Machine Learning Research*, pages 1584–1616. PMLR, 23–26 Jun 2016.

Chicheng Zhang and Kamalika Chaudhuri. Active learning from weak and strong labelers. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, pages 703–711. Curran Associates, Inc., 2015.

Chicheng Zhang, Jimin Song, Kamalika Chaudhuri, and Kevin Chen. Spectral learning of large structured hmms for comparative epigenomics. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, pages 469–477. Curran Associates, Inc., 2015.

Chicheng Zhang and Kamalika Chaudhuri. Beyond disagreement-based agnostic active learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27, pages 442–450. Curran Associates, Inc., 2014.

JOURNAL PAPERS Akshay Krishnamurthy*, John Langford*, Aleksandrs Slivkins*, and Chicheng Zhang*. Contextual bandits with continuous actions: Smoothing, zooming, and adapting. *Journal of Machine Learning Research (JMLR)*, 21(137):1–45, 2020.

PREPRINTS Chicheng Zhang and Zhi Wang. Provably efficient multi-task reinforcement learning with model transfer. *arXiv preprint arXiv:2107.08622*, 2021.

Chicheng Zhang, Eran A. Mukamel, and Kamalika Chaudhuri. Spectral learning of binomial hmms for DNA methylation data. *CoRR*, abs/1802.02498, 2018.

SUBMITTED Tom Yan and Chicheng Zhang. Margin-distancing for safe model explanation. *Submitted*, 2021.

OPEN SOURCE SOFTWARE CONTRIBUTIONS

Vowpal Wabbit - Warm-starting contextual bandits

Code:

https://github.com/VowpalWabbit/vowpal_wabbit/blob/master/vowpalwabbit/warm.cb.cc

Documentation:

https://github.com/VowpalWabbit/vowpal_wabbit/wiki/Warm-starting-contextual-bandits

Scripts for experimental evaluation:

https://github.com/zcc1307/warmcb_scripts

SELECTED TALKS

- Efficient active learning of halfspaces: noise tolerance and exploiting sparsity**
Arizona State University Machine Learning Day, Virtual Online Apr 2021
<https://asuevents.asu.edu/content/machine-learning-day-1>
- Efficient active learning of sparse halfspaces with arbitrary bounded noise**
NeurIPS 2020 oral presentation, Virtual Online Dec 2020
https://neurips.cc/virtual/2020/protected/session_oral_21088.html
- Los Alamos - Arizona Days, Virtual Online May 2020
<https://appliedmath.arizona.edu/events/los-alamos-arizona-days>
- Efficient contextual bandits with continuous actions via extreme classification**
ICML 2020 workshop on extreme classification, Virtual Online July 2020
<http://manikvarma.org/events/XC20/index.html>
- New directions in contextual bandits learning: continuous actions and linear separability**
TRIPODS Seminar, University of Arizona, Tucson, AZ Sep 2019
<https://tripods.math.arizona.edu/events>
- Contextual bandits with continuous actions: smoothing, zooming, and adapting**
COLT 2019, Phoenix, AZ June 2019
<https://learningtheory.org/colt2019/program.html>
- Efficient and robust interactive learning** February - April 2019
Illinois Institute of Technology, Chicago, IL Host: Gady Agam and Zhiling Lan
University of Arizona, Tucson, AZ Host: John Kececioglu
<https://www.cs.arizona.edu/news-calendars/colloquia/colloquia-2018-2019>
Stevens Institute of Technology, Hoboken, NJ Host: Jie Shen
University of Minnesota Twin Cities, Minneapolis, MN Host: Dan Boley
Pennsylvania State University, State College, PA Host: David Miller
University of Connecticut, Storrs, CT Host: Alexander Russell
<https://www.cse.uconn.edu/events-calendar/category/colloquia/list/>
Rensselaer Polytechnic Institute, Troy, NY Host: Alex Gittens
<https://science.rpi.edu/computer-science/events/computer-science-colloquium>
University of Illinois at Chicago, Chicago, IL Host: Brian Ziebart
- Efficient active learning of sparse halfspaces**
TRIPODS RWG6 Seminar, University of Arizona, Tucson, AZ Sep 2019
<https://sites.google.com/math.arizona.edu/tripods-rwg6/group-meeting>
- ALT 2019 Workshop on “When Smaller Sample Sizes Suffice for Learning”
Chicago, IL March 2019
<http://algorithmiclearningtheory.org/alt2019/workshop/>
- Interactive learning with data-efficiency and robustness guarantees**
Peking University EECS Youth Forum, Beijing, China December 2018
Microsoft Research Asia, Beijing, China January 2019
Baidu Research, Beijing, China January 2019

Efficient online bandit multiclass learning with $\tilde{O}(\sqrt{T})$ regret
 ICML 2017, Sydney, Australia August 2017
<https://vimeo.com/channels/1301905/237240526>

International Chinese Statistical Association, New Brunswick, NJ June 2018
<https://www.icsa.org/june-14-17-2018-icsa-2018-applied-statistics-symposium/>

Computationally and statistically efficient active learning of linear separators
 NYU Machine Learning PhD Seminar, New York, NY March 2018
<http://people.stern.nyu.edu/nzhang/mlseminar18.html>

Tutorial on statistical foundations of interactive learning June 2017
 ISIT 2017, Aachen, Germany (co-presented with Kamalika Chaudhuri and Tara Javidi)
<https://isit2017.org/tutorials.html>

New directions in active learning
 Microsoft Research, New York, NY March 2017
 Google Research, New York, NY March 2017

Confidence-based active learning
 Yahoo! Research, New York, NY February 2017
 UCSD Computational Statistics and Machine Learning Seminar, La Jolla, CA May 2015

The extended Littlestone’s dimension for learning with mistakes and abstentions
 COLT 2016, New York, NY June 2016
<https://www.youtube.com/watch?v=wSu9063-dko>

TEACHING

INSTRUCTOR CSC 696H – Topics in Reinforcement Learning Theory, University of Arizona. Fall 2021
<https://zcc1307.github.io/courses/csc696fa21/>

CSC 588 – Machine Learning Theory, University of Arizona. Spring 2021
<https://zcc1307.github.io/courses/csc588sp21/>

CSC 665 – Machine Learning Theory, University of Arizona. Fall 2019
<https://zcc1307.github.io/courses/csc665fa19/>

TEACHING ASSISTANT CSE 151 – Introduction to Machine Learning, UC San Diego. Spring 2015, Winter 2017
 CSE 202 – Graduate Algorithms, UC San Diego. Spring 2016
 CSE 250C – Machine Learning Theory, UC San Diego. Spring 2017

AWARDS

ICML student / postdoc travel grant 2017, 2019
 NeurIPS student travel grant 2016
 4th place in ACM Southern California Regional Programming Contest 2015
 8th place in ACM Southern California Regional Programming Contest 2013
 2nd place in UCSD Programming Contest 2013
 UCSD Graduate Student Fellowship 2012
 Li Huirong Scholarship 2011
 3rd Prize in Beijing Collegiate Mathematical Contest 2011
 Starlight International Media Scholarship 2010

Merit Student Award	2009
3rd Prize in National Mathematics Olympiad in Province	2007

SERVICES

WORKSHOP CO-ORGANIZER	ICML 2017 Workshop on Picky Learners: Choosing Alternative Ways to Process Data. Sydney, Australia (with Kamalika Chaudhuri, Corinna Cortes, Giulia DeSalvo, and Ning- shan Zhang) August 2017 https://sites.google.com/site/rejectionactiveicml/
AREA CHAIR / SENIOR PROGRAM COMMITTEE	Conference on Neural Information Processing Systems (NeurIPS): 2019, 2021 International Conference on Machine Learning (ICML): 2020, 2021 International Conference on Artificial Intelligence and Statistics (AISTATS): 2019 International Conference on Learning Representations (ICLR): 2021, 2022 (invited)
CONFERENCE REVIEWER	Conference on Neural Information Processing Systems (NeurIPS): 2015, 2016, 2017, 2018, 2020 International Conference on Machine Learning (ICML): 2016, 2017, 2018, 2019 Conference on Learning Theory (COLT): 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021 International Conference on Artificial Intelligence and Statistics (AISTATS): 2016, 2017, 2018, 2020, 2021 International Conference on Learning Representations (ICLR): 2018 Conference on Uncertainty in Artificial Intelligence (UAI): 2015, 2016 Conference on Algorithmic Learning Theory (ALT): 2015, 2021 AAAI Conference on Artificial Intelligence (AAAI): 2019 IEEE Symposium on Foundations of Computer Science (FOCS): 2020, 2021
JOURNAL REVIEWER	IEEE Transactions on Pattern Analysis and Machine Intelligence IEEE Transactions on Information Theory IEEE Transactions on Signal Processing IEEE Journal on Selected Areas in Information Theory Journal of Artificial Intelligence Research Journal of the ACM Journal of Machine Learning Research Theoretical Computer Science
GRANT PROPOSAL REVIEWER	US National Science Foundation Core Program (served as a panelist) US National Science Foundation SBIR/STTR Program The Israel Science Foundation