

Research Area

Machine learning and optimization with a focus on GPU acceleration, recommendation system, SDP, and differentiable optimization layer.

Experience

Applied Scientist at Pinterest

Sep 2021 – Now

Ph.D. at Machine Learning Dept., Carnegie Mellon University

Aug. 2015 – Aug 2021

- Advised by Prof. J. Zico Kolter
- My thesis work is on Learning and Reasoning using fast scalable semidefinite programming (SDP)
 - We scales SDPs to millions of variables, unleashing their power to ML and AI.
 - The proposed method is applied to probabilistic inference [2], parameter learning [4], and clustering [1], achieving better performance than state-of-the-art methods in shorter time.
 - We are applying the method to graph representation learning and language modeling.

B.S. at Computer Science Dept., National Taiwan University

Sep. 2008 – June 2012

- Advised by Prof. Chih-Jen Lin
- Worked on optimization of SVMs and its convergence properties

Research Topics

- Differentiable satisfiability solver as a layer** [4]
 - Logical reasoning within deep learning using a differentiable SAT solver.
 - Approximates MAXSAT with an implicitly differentiable SDP optimization layer.
 - Wrap-level GPU optimization with CUDA C. **Best paper honorable mention** at ICML '19.
- Low-rank semidefinite solvers for MAXCUT and MAXSAT** [11, 10, 5]
 - The first proof of optimal convergence using Lyapunov analysis and the stable-manifold theorem.
 - Breakthrough in scaling semidefinite program to millions of variables.
 - Orders of magnitude faster than other existing methods in experiments.
- The distributed common-direction solver for linear classification** [9, 6]
 - Reuse gradient information to reduce communication cost in distributed optimization.
 - Outperforms the state-of-the-art first- and second-order methods in experiments.
 - Converges linearly in optimal rate and enjoys local quadratic convergence.
- Global linear convergence for non-strongly convex problems (i.e., RSC)** [14]
 - The first global linear rate for first-order methods on non-strongly convex problems.
 - Providing theoretical foundation for libLinear package for SVM, published in JMLR.
 - Solves the open problem of convergence rate of the Gauss-Seidel method on PSD matrices.

Publications

- Community detection using fast low-cardinality semidefinite programming.** NeurIPS '20
Po-Wei Wang and J. Zico Kolter. *NeurIPS, 2020*
- Efficient semidefinite-programming-based inference for binary and multi-class MRFs.** NeurIPS '20
Chirag Pabbaraju, **Po-Wei Wang**, and J. Zico Kolter. *NeurIPS, 2020*
- Differentiable learning of numerical rules in knowledge graphs.** ICLR '20
Po-Wei Wang, Daria Stepanova, Csaba Domokos, J. Zico Kolter. *ICLR, 2020*
- SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver.** ICML '19
Po-Wei Wang, Priya L. Donti, Bryan Wilder, Zico Kolter. *ICML, 2019 (best paper honorable mention)*

5. **Low-rank semidefinite programming for the MAX2SAT problem.** AAAI '19
Po-Wei Wang and J. Zico Kolter. *AAAI, 2019*
6. **The Common-directions Method for Regularized Empirical Risk Minimization.** JMLR '19
Po-Wei Wang, Ching-pei Lee, and Chih-Jen Lin. *JMLR, 2019*
7. **Realtime query completion via deep language models.** ECOM '18
Po-Wei Wang, Huan Zhang, Vijai Mohan, Inderjit S. Dhillon, and J. Zico Kolter. *SIGIR eCom, 2018*
8. **Polynomial optimization methods for matrix factorization.** AAAI '17
Po-Wei Wang, Chun-Liang Li, and J. Zico Kolter. *AAAI, 2017*
9. **Limited-memory common-directions method for distributed optimization and its application on empirical risk minimization.** SDM '17
 Ching-pei Lee, **Po-Wei Wang**, Weizhu Chen, and Chih-Jen Lin. *SDM, 2017*.
10. **The Mixing method: low-rank coordinate descent for semidefinite programming with diagonal constraints.** ArXiv '17
Po-Wei Wang, Wei-Cheng Chang, and J. Zico Kolter. *Tech report, 2017*
11. **The Mixing method for Maxcut-SDP problem.** LHDS '16
Po-Wei Wang and J. Zico Kolter. *NIPS LHDS Workshop, 2016*
12. **Epigraph Projections for Fast General Convex Programming.** ICML '16
Po-Wei Wang, Matt Wytock, and J. Zico Kolter. *ICML, 2016*
13. **Epigraph Proximal Algorithms for General Convex Programming.** OPT '15
 Matt Wytock, **Po-Wei Wang**, and J. Zico Kolter. *Proceedings of the NIPS 2015 Optimization Workshop*
14. **Iteration Complexity of Feasible Descent Methods for Convex Optimization.** JMLR '15
Po-Wei Wang and Chih-Jen Lin. *The Journal of Machine Learning Research 2015, 1523-1548*
15. **Support Vector Machines.** BOOK '14
Po-Wei Wang and Chih-Jen Lin. *Data Classification: Algorithms and Applications. CRC Press. 2014*
16. **On Convergence Rate of Concave-Convex Procedure.** OPT '12
 Ian E.H. Yen, Nanyun Peng, **Po-Wei Wang**, Shou-De Lin. *Proceedings of the NIPS 2012 Optimization Workshop*

Selected Awards

1. **Best paper honorable mention from ICML 2019**
2. **Double Championships in KDD Cup competition on Music Rating Prediction (2011)**
3. **Third Place in Robocup Standard Platform League World Final (2011)**

Professional Services

1. Program committees: NeurIPS (16-now), ICML (18), AAAI (20-21), ICLR (21), TPAMI (19), TKDE (17-18), JMLR (14,17), CDC (14), Neurocomputing (15), and DAMI (13-14).
2. Teaching: Head TA in convex optimization 2018 and a normal TA in 2019.
3. Organizer of NTU Machine Learning Symposium 2014

Intern Experience

1. Intern at Bosch (Pittsburgh and Renningen) for differentiable logic. May 2019 – Aug 2019
 - Filed two patents, published at ICLR.
2. Intern at A9 (Amazon) for Real-time Deep Query Completion. May 2017 – Aug 2017
 - Filed a patent for A9 and was gifted a scholarship.