ALEC KOPPEL

U.S. Army Research Laboratory
Computational and Information Sciences Directorate
2800 Powder Mill Road
Adelphi, MD 20783

Tel: (301)-394-1126 (Office)
Tel: (314)-303-2399 (Mobile)
Email: akoppel@seas.upenn.edu
http://koppel.netlify.com/

RESEARCH INTERESTS

Applications of stochastic optimization to develop new learning methods, both for statistical inference and autonomous control. Topics of interest include: centralized and decentralized convex optimization, memory-efficient kernel methods, supervised and reinforcement learning, collaboration in robotic teams, networked control systems, adaptive signal processing, online machine learning, and mathematical statistics.

EMPLOYMENT

U.S. Army Research Laboratory — Adelphi, MD
Research Scientist — August 2017 - Present
Computational and Information Sciences Directorate

University of Pennsylvania — Philadelphia, PA
Research Assistant — August 2012 - August 2017
Electrical and Systems Engineering

U.S. Army Research Laboratory — Adelphi, MD
SMART Fellow — July 2013 - July 2017
Computational and Information Sciences Directorate

Washington University in St. Louis — St. Louis, MO
Research Assistant — July 2010 - July 2012
Department of Mathematics; Division of Biostatistics

EDUCATION

University of Pennsylvania — Philadelphia, PA
Ph. D., Electrical & Systems Engineering — August 2017
Thesis: "Stochastic Optimization for Multi-Agent Statistical Learning and Control"
GPA: 3.72; *Advisor*: Prof. Alejandro Ribeiro

The Wharton School, University of Pennsylvania — Philadelphia, PA
M. Sc. Statistics — August 2017
Thesis: "Parameter Estimation in High-Dimensions using Doubly Stochastic Approximation"
GPA: 3.72; *Advisor*: Prof. Dylan Small

Washington University — St. Louis, MO
M. Sc. Systems Science & Mathematics — May 2012
GPA: 3.63; *Advisor*: Prof. Hiro Mukai

Washington University — St. Louis, MO
B.A. Mathematics, Magna Cum Laude — May 2011
Thesis: "Stochastic Methods for the Lotka-Volterra Model with Migration"
GPA: 3.66; *Advisor*: Prof. Renato Feres

ACADEMIC HONORS

- IEEE Asilomar Signals, Systems, & Computers Best Paper Finalist — Nov. 2017
- University of Pennsylvania Award for Exceptional Service to ESE — Fall 2016
- SMART National Fellowship, sponsored by U.S. Defense Dept. — Summer 2013
- University of Pennsylvania Fellowship Award for graduate studies — Fall 2012
- George Washington University Fellowship Award for graduate studies — Fall 2012
- Latin Honors: Magna Cum Laude, WashU Dept. of Mathematics — Spring 2011

PROFESSIONAL AFFILIATIONS AND SERVICES

- Member of: IEEE Signal Processing Society. Mathematical Optimization Society, INFORMS Optimization Society
- Reviewer the following publications: SIAM Journal on Optimization, IEEE Transactions on Signal Processing, IEEE Transactions on Automatic Control, IEEE Transactions on Control of Network Systems, IEEE Int. Conf. Acoustics, Speech, Signal Process., IEEE Workshop on Signal Process. Advances in Wireless Com., IEEE Conference on Decision and Control (CDC), Automatica, European Journal on Signal Processing, Elsevier Signal Processing, Neural Information Processing Systems
- Conference sessions chaired: SIAM ICCOPT, Aug. 2019; IEEE Asilomar Conference, Nov. 2019; INFORMS Annual Meeting, Nov. 2018; INFORMS Optimization Society Conference, Mar. 2018; IEEE Asilomar Conference, Nov. 2015; INFORMS Optimization Society Conference, Mar. 2016
- UPenn ESE Phd Student Colloquium Coordinator Fall 2013-Summer 2017
- UPenn Graduate Students Engineering Group (GSEG) Rep. Fall 2015-Summer 2017

PUBLICATIONS

Journal Papers

1. A. S. Bedi, A. Koppel, K. Rajawat, and B.M. Sadler. "Nonstationary Nonparametric Online Learning: Balancing Dynamic Regret and Model Parsimony," in *IEEE Trans. Signal Process* (submitted), Sept. 2019.

2. R. Pradhan, A. S. Bedi, A. Koppel, and K. Rajawat. "Adaptive Kernel Learning in Heterogeneous Networks," in *IEEE Trans. Signal Process* (submitted), Aug. 2019.

3. A. S. Bedi, A. Koppel, and K. Rajawat. "Nonparametric Compositional Stochastic Optimization: Algorithms for Robust Online Learning with Kernels," in *IEEE Trans. Signal Process* (submitted), Feb. 2019.

4. K. Zhang, A. Koppel, H. Zhu, and T. M. Baser. "Global Convergence of Policy Gradient Methods: A Nonconvex Optimization Perspective" in *SIAM Journal on Control and Optimization* (submitted), Jan. 2019.

5. A. S. Bedi, A. Koppel, and K. Rajawat. "Asynchronous Online Learning in Multi-Agent Systems with Proximity Constraints" in *IEEE Transactions on Signal and Information Processing over Networks*, 2019

6. A. Koppel, K. Zhang, H. Zhu, and T. M. Baser. "Projected Stochastic Primal-Dual Method for Constrained Online Learning with Kernels" in *IEEE Trans. Signal Process.*, May. 2019.

7. A. Koppel, E. Tolstaya, E. Stump, and A. Ribeiro. "Nonparametric Stochastic Compositional Gradient Descent for Q-Learning in Continuous Markov Decision Problems" in *IEEE Trans. Automatic Control* (to appear), 2019.

8. A. S. Bedi, A. Koppel, and K. Rejawat, "Asynchronous Decentralized Stochastic Optimization in Heterogeneous Networks" in *IEEE Trans. Signal Process.*, Jan. 2019.

9. A. Koppel, G. Warnell, E. Stump, P. Stone, and A. Ribeiro. "Policy Evaluation in Continuous MDPs with Efficient Kernelized Gradient Temporal Difference," in *IEEE Trans. Automatic Control* (submitted), Dec. 2017.

10. M. Fazlyab, A. Koppel, V. Preciado, and A. Ribeiro, "Synthesis of Accelerated Optimization Methods: From Continuous-time Systems to Discrete-time Algorithms," in *IEEE Trans. Automatic Control* (under review), Nov. 2017.

11. A. Koppel, S. Paternain, C. Richard, and A. Ribeiro, "Decentralized Online Learning with Kernels", in *IEEE Trans. Signal Process*, Apr. 2018.

12. A. Koppel, G. Warnell, E. Stump, and A. Ribeiro, "Parsimonious Online Learning with Kernels via Sparse Projections in Function Space," in Journal of Machine Learning Research, Jan. 2019

13. A. Mokhtari, A. Koppel, and A. Ribeiro, "A Class of Doubly Random Parallel Stochastic Methods for Large Scale Learning," in *Journal of Machine Learning Research* (under review), June 2016

14. A. Koppel, B. Sadler, and A. Ribeiro, "Proximity without Consensus in Online Multi-Agent Optimization," in *IEEE Trans. Signal Proc*, Volume: 65 Issue: 12 , Page 3062-3077, June 15, 2017.

15. A. Koppel, G. Warnell, E. Stump, and A. Ribeiro, "D4L: Decentralized Dynamic Discriminative Dictionary Learning," *IEEE Trans. Signal Info. Process. over Networks*, Mar. 2017.

16. A. Simonetto, A. Koppel, A. Mokhtari, G. Leeus, and A. Ribeiro, "Decentralized Prediction-Correction Methods for Networked Time-Varying Convex Optimization," *IEEE Trans. Automatic Control*, Volume 62, Issue 11. Nov, 2017.

17. A. Simonetto, A. Mokhtari, A. Koppel, G. Leeus, and A. Ribeiro, "A Class of Prediction-Correction Methods for Time-Varying Convex Optimization," *IEEE Trans. Signal Process*, Sept. 2015'.

18. A. Koppel, F. Jakubeic, and A. Ribeiro, "A saddle point algorithm for networked online convex optimization," *IEEE Trans. Signal Process.*, Oct. 2015.

Conference Papers

1. S. Bhatt, A. Koppel, V. Krishnamurthy, "Policy Gradient using Weak Derivatives for Reinforcement Learning," in IEEE Conference on Decision and Control (CDC) (to appear), Nice, France, Dec. 11-13, 2019.

2. K. Zhang, A. Koppel, H. Zhu, T. Basar, "Convergence and Iteration Complexity of Policy Gradient Methods for Infinite-horizon Reinforcement Learning," in IEEE Conference on Decision and Control (CDC) (to appear), Nice, France, Dec. 11-13, 2019.

3. S. Bhatt, A. Koppel, V. Krishnamurthy, "Policy Search using Jordan Decomposition for Reinforcement Learning," in IEEE Conference on Information Sciences and Systems (CISS), Baltimore, MD, Mar. 20-22, 2019.

4. K. Zhang, A. Koppel, H. Zhu, T. Basar, "Policy Search in Infinite-Horizon Discounted Reinforcement Learning: Advances through Connections to Non-Convex Optimization," in IEEE Conference on Information Sciences and Systems (CISS), Baltimore, MD, Mar. 20-22, 2019.

5. A. Koppel, A. S. Bedi, K. Rajawat, "Controlling the the Bias-Variance Tradeoff via Coherent Risk for Robust Learning with Kernels," in *American Control Conference*, Philadelphia, PA, July 10-12, 2019.

6. A. Koppel, "Consistent Online Gaussian Process Regression Without the Sample Complexity Bottleneck," in *American Control Conference*, Philadelphia, PA, July 10-12, 2019.

7. H. Pradhan, A. S. Bedi, A. Koppel, and K. Rajawat, "Exact Decentralized Online Nonparametric Optimization," in *IEEE Global Conf. on Signal and Info. Processing* (GlobalSIP), Anaheim, CA, Nov. 26-28, 2018.

8. A. Koppel, S. Paternain, C. Richard, and A. Ribeiro,"Decentralized Online Nonparametric Learning", in Proc. Asilomar Conf. Signals, Systems, Computers, Pacific Grove, CA, Oct. 28-31, 2018.

9. A. S. Bedi, A. Koppel, and K. Rajawat, "Asynchronous Saddle Point Method: Interference Management Through Pricing," in *IEEE Conf. on Decision and Control* (CDC), Miami Beach, FL, Dec. 17-19, 2018.

10. K. Zhang, H. Zhu, T. Baser, and A. Koppel , "Projected Stochastic Primal-Dual Method for Constrained Online Learning with Kernels,"in *IEEE Conf. on Decision and Control* (CDC), Miami Beach, FL, Dec. 17-19, 2018.

11. E. Tolstaya, E. Stump, A. Koppel, and A. Ribeiro, "Composable Learning with Sparse Kernel Representations," in *International Conference on Intelligent Robots and Systems* (IROS), Madrid, Spain, Oct 1-5, 2018.

12. E. Tolstaya, A. Koppel, E. Stump, and A. Ribeiro, "Nonparametric Stochastic Compositional Gradient Descent for Q-Learning in Continuous Markov Decision Problems," in *American Control Conference*, Milwaukee, WI, June 27-29, 2018.

13. A. Koppel, A. Mokhtari, and A. Ribeiro, "Parallel Stochastic Successive Convex Approximation Method for Large-Scale Dictionary Learning," in *Proc. Int. Conf. Acoustics Speech Signal Process*, Calgary, Canada, Apr. 15-20, 2018.

14. A. Koppel, S. Paternain, C. Richard, and A. Ribeiro, "Decentralized Efficient Nonparametric Stochastic Optimization", in *IEEE Global Conference on Signal and Information Processing*, Montreal, Canada, Nov. 14-16, 2017.

15. A. S. Bedi, A. Koppel, and K. Rejawat, "Beyond Consensus and Synchrony in Decentralized Online Optimization using Saddle Point Method" in *Proc. Asilomar Conf. on Signals Systems Computers*, Pacific Grove, CA, Oct. 29-Nov. 1, 2017.

16. M. Fazylab, A. Koppel, V. Preciado, and A. Ribeiro, "A Variational Approach to Dual Methods for Constrained Convex Optimization," in *American Control Conference*, Seattle, WA, May 24-26, 2017.

17. A. Mokhtari, A. Koppel, and G. Scutari, A. Ribeiro, "Large-Scale Non-Convex Stochastic Optimization by Doubly Stochastic Successive Convex Approximation," in *Proc. Int. Conf. Acoustics Speech Signal Processing*, New Orleans, LA, USA Mar. 5-9 2017.

18. A. Koppel, G. Warnell, E. Stump, and A. Ribeiro, "Parsimonious Online Learning with Kernels via Sparse Projections in Function Space," in *Proc. Int. Conf. Acoustics Speech Signal Process*, New Orleans, LA, USA Mar. 5-9 2017.

19. A. Koppel, B. M. Sadler, and A. Ribeiro, "Decentralized Online Optimization with Heterogeneous Data Sources", *IEEE Global Conference on Signal and Information Processing* (to appear), Washington, DC, Dec. 7-9, 2016.

20. A. Koppel, A. Mokhtari, and A. Ribeiro "Doubly Random Parallel Stochastic Methods for Large Scale Optimization." in *Proc. Asilomar Conf. on Signals Systems Computers*, Pacific Grove, CA, November 6-9 2016.

21. A. Koppel, J. Fink, G. Warnell, E. Stump, and A. Ribeiro, "Online Learning for Characterizing Unknown Environments in Ground Robotic Vehicle Models," in *Proc. Int. Conf. Intelligent Robotics and Systems*, South Korea, Oct. 2016

22. A. Simonetto, A. Koppel, A. Mokhtari, G. Leus, and A. Ribeiro, "A Quasi-Newton Prediction-Correction Method for Decentralized Dynamic Convex Optimization", *European Control Conference*, Aalborg, Denmark, June 29 - July 1, 2016.

23. A. Mokhtari, A. Koppel, and A. Ribeiro, "Doubly Random Parallel Stochastic Methods for Large Scale Learning," in *American Control Conference*, Boston, MA, July 6-8 2016.

24. A. Koppel, B. M. Sadler and A. Ribeiro, "Proximity without consensus in online multi-agent optimization," in *Proc. Int. Conf. Acoustics Speech Signal Process*, Shanghai, China, Mar. 20-25 2016.

25. A. Simonetto, A. Mokhtari, A. Koppel, G. Leus, and A. Ribeiro, "A Decentralized Prediction-Correction Method for Networked Time-Varying Convex Optimization", *IEEE Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, Cancun, Mexico, Dec. 13-16, 2015.

26. A. Koppel, A. Simonetto, A. Mokhtari, G. Leus, and A. Ribeiro, "Target Tracking with Dynamic Convex Optimization", *IEEE Global Conference on Signal and Information Processing*, Orlando, FL, Dec. 14-16, 2015.

27. A. Simonetto, A. Mokhtari, A. Koppel, G. Leeus, and A. Ribeiro "Prediction-Correction Methods for Time-Varying Convex Optimization." in *Proc. Asilomar Conf. on Signals Systems Computers*, Pacific Grove, CA, November 8-11 2015.

28. A. Koppel, G. Warnell, E. Stump, and A. Ribeiro "Task-Driven Dictionary Learning in Distributed Online Settings." in *Proc. Asilomar Conf. on Signals Systems Computers*, Pacific Grove, CA, November 8-11 2015.

29. A. Koppel, G. Warnell, E. Stump, and A. Ribeiro, "D4L: Decentralized Dynamic Discrminative Dictionary Learning," in *Proc. Int. Conf. Intelligent Robotics and Systems*, Hamburg, Germany, Sep 28-Oct2 2015.

30. A. Koppel, F. Jakubeic and A. Ribeiro, "Regret Bounds of a distributed saddle point algorithm," in *Proc. Int. Conf. Acoustics Speech Signal Process.*, Brisbane Australia, Apr 19-24 2015.

31. A. Koppel, F. Y. Jakubiec, and A. Ribeiro, "A Saddle Point Algorithm for Networked Online Convex Optimization." in *39th Proc. Int. Conf. Acoust. Speech Signal Process.*, May 4-9 2014, pp. 8292 - 8296.

SKILLS

- *Programming Languages*: MATLAB & SIMULINK, Python, R, SAS
- *Applications*: LATEX, Microsoft Office, CVX
- *Operating Systems*: Microsoft Windows 7/XP/2000, Linux, Ubuntu, Mac OSX
- *Languages*: English (native), Spanish (proficient)

TEACHING EXPERIENCE

Teaching Certification, **UPenn's Center for Teaching & Learning** — Fall 2015

Teaching assistant, **University of Pennsylvania**

- "Signal and Information Processing" (Instructor: Prof. Ribeiro) — Spring 2015
- "Modern Convex Optimization" (Instructor: Prof. Ribeiro) — Spring 2014
- "Engineering Probability", (Instructor: Prof. Sarkar) — Fall 2013

Teaching assistant & ***Peer Academic Mentor***, **WashU**

- "Calculus of Several Variables" (Instructor: Prof. Thornton) — Spring 2011
- "Matrix Algebra" (Instructor: Prof. Freiwald) — Fall 2010
- "Calculus III", (Instructor: Prof. Blank) — Fall 2010
- "Calculus II", (Instructor: Prof. Feres) — Spring 2009

MENTORING EXPERIENCE

ARL Research Associates: **Graduate Fellows**

- Amrit Singh Bedi, "Risk-awareness in online nonparametric methods" Fall 2017
  Affiliation: Postdoctoral fellow, U.S. Army Research Laboratory
- Kaiqing Zhang, "Non-convex optimi zation approaches to policy search" Summer 2018
  Affiliation: Phd student, University of Illinois Urbana-Champaign
- Hrusikesha Pradhan, "Consistent compressions of Gaussian Processes" Fall 2018
  Affiliation: Phd student, EE, India Institute of Technology, Kanpur
- Sujay Bhatt, "Variance-reduced policy gradients via weak-derivatives" Fall 2018
  Affiliation: Phd, Cornell University
- Amir Daneshmand, "Multi-agent RL with Incomplete Information" Summer 2019
  Affiliation: Phd student, Purdue University

ARL Research Associates: **Postdoctoral Fellows**

- Amrit Singh Bedi Winter 2019 - present

Other Mentorship Experience: **Graduate Fellows**

- Ehsan Zobeidi, "Gaussian Processes for Occupancy Mapping" Spring 2019 - present
  Affiliation: Phd student, University of California, San Diego
- Mathew Rittonia, " Defenses to Black Box Data Poisoning " Spring 2019 - present
  Affiliation: MS student, Johns Hopkins University
- Zhan Gao, "Covarying Step and Batch-sizes in Learning" Spring 2019 - present
  Affiliation: Phd student, UPenn
- Harshat Kumar, "Sample Complexity of Actor-Critic in RL" Summer 2018 - present
  Affiliation: Phd student, UPenn
- Ekaterina Tolstaya, "Compressed Kernel Q Learning" Fall 2016 - Spring 2018
  Affiliation: Phd student, UPenn