

ALEC KOPPEL

JP Morgan Chase & Co.
Artificial Intelligence Research
383 Madison Ave.
New York, NY 10017

Tel: (314)-303-2399 (Mobile)
Email: alec.koppel@jpmchase.com
Alternate: akoppel@seas.upenn.edu
<http://koppel.netlify.app/>

RESEARCH INTERESTS	Stochastic optimization for developing learning methods, both for statistical inference and autonomous control. Topics include: centralized and decentralized convex optimization, large-scale and online learning, reinforcement learning and dynamic programming, and Bayesian/nonparametric statistics. Application areas of interest include robotics and autonomy, sourcing and vendor selection, and financial markets.
EMPLOYMENT	<p>JP Morgan Chase & Co. New York, NY Vice President/AI Research Lead June 2022 - present <i>AI Research - Multi-agent Learning and Simulation team</i></p> <p>Amazon Bellevue, WA Research Scientist II Sept. 2021 - June 2022 <i>Supply Chain Optimization Technologies (SCOT)</i></p> <p>U.S. Army Research Laboratory Adelphi, MD Research Scientist Sept. 2017 - Sept. 2021 <i>Computational and Information Sciences Directorate</i></p> <p>University of Pennsylvania Philadelphia, PA Research Assistant August 2012 - August 2017 <i>Electrical and Systems Engineering</i></p> <p>U.S. Army Research Laboratory Adelphi, MD SMART Fellow July 2013 - July 2017 <i>Computational and Information Sciences Directorate</i></p> <p>Washington University in St. Louis St. Louis, MO Research Assistant July 2010 - July 2012 <i>Department of Mathematics; Division of Biostatistics</i></p>
EDUCATION	<p>University of Pennsylvania Philadelphia, PA Ph. D., Electrical & Systems Engineering August 2017 <i>Thesis: "Stochastic Optimization for Multi-Agent Statistical Learning and Control"</i> GPA: 3.72; <i>Advisor: Prof. Alejandro Ribeiro</i></p> <p>The Wharton School, University of Pennsylvania Philadelphia, PA M. Sc. Statistics August 2017 <i>Thesis: "Parameter Estimation in High-Dimensions using Doubly Stochastic Approximation"</i> GPA: 3.72; <i>Advisor: Prof. Dylan Small</i></p> <p>Washington University St. Louis, MO M. Sc. Systems Science & Mathematics May 2012 GPA: 3.63; <i>Advisor: Prof. Hiro Mukai</i></p> <p>Washington University St. Louis, MO B.A. Mathematics, Magna Cum Laude May 2011 <i>Thesis: "Stochastic Methods for the Lotka-Volterra Model with Migration"</i> GPA: 3.66; <i>Advisor: Prof. Renato Feres</i></p>

ACADEMIC
HONORS

- Co-author on paper selected for Spotlight at ICML July 2022
- Mentor to ARL Summer Symposium Best Project Awardee August 2021
- IEEE Robotics & Automation Society Honorable Mention, Best Paper Award May 2021
- Co-author on paper selected for Spotlight at NeurIPS Dec. 2020
- ARL Director's Research Award TRC May 2020
- IEEE SPCOM Best Graduation Day Talk Award (co-author) Jul. 2018
- 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
- **Area Chair for: 2024 ICLR, 2024 NeurIPS, 2023 ICLR, 2023 ICML**
- Program Committee for: 2020-2023 Learning for Dynamics and Control (L4DC), 2019-2023 NeurIPS, 2020-2023 ICML, 2018-2022 Technical Program Affiliate for IEEE Signal Processing Society
- Reviewer for the IEEE publications: Trans. Information Theory, Trans. Signal Processing, Trans. Control of Network Systems, Trans. Automatic Control, ICASSP, SPAWC, CDC
- Reviewer for the Springer publications: Journal of Optimization Theory and Applications (JOTA), Set-Valued and Variational Analysis
- Reviewer for the Machine Learning publications: The Journal of Machine Learning Research, NeurIPS (2018-2022), ICML (2019-2022), ICLR (2021-2023), AAAI (2021-2023)
- Reviewer for SIAM Journal on Optimization, Automatica, European Journal on Signal Processing, Elsevier Signal Process.
- Conference sessions chaired: 2023 IEEE CISS, 2020 INFORMS Annual Meeting, 2019 SIAM ICCOPT, 2019 IEEE Asilomar Conf; 2018 INFORMS Annual Meeting; 2018 INFORMS Optimization Soc. Conf.; 2016 INFORMS Optimization Soc. Conf.; 2015 IEEE Asilomar Conf.;
- Mentor to the Office of Secretary of Defense (OSD) Historically Black Colleges and Universities (HBCU) Summer Internship Program, and ARO Technical Advisor Program (2019-2021)
- UPenn ESE Phd Student Colloquium Coordinator Fall 2013-Summer 2017
- UPenn Graduate Students Engineering Group (GSEG) Rep. Fall 2015-Summer 2017

PUBLICATIONS

Working papers (Under Review)

1. W. A. Suttle, **A. Koppel**, J. Liu, "Occupancy Information Ratio: Infinite-Horizon, Information-Directed, Parameterized Policy Search." arXiv preprint arXiv:2201.08832. SIAM Journal on Control and Optimization (Major Revision April 2023)
2. A. S. Bedi, C. Fan, **A. Koppel**, A. K. Sahu, B. M. Sadler, F. Huang, D. Manocha, (2022). "FedBC: Calibrating Global and Local Models via Federated Learning Beyond Consensus." preprint available as arXiv:2206.10815.

Journal Papers: Refereed

3. A. Bedi, D. Peddireddy, V. Aggarwal, and **A. Koppel**, “Sublinear Regret and Belief Complexity in Gaussian Process Bandits via Information Thresholding,” in *2024 IEEE Transactions on Artificial Intelligence* Aug. pp. 1-10, vol. 1 DOI Bookmark: 10.1109/TAI.2023.3332023, preprint available as arXiv 2003.10550.
4. A. S. Bedi, A. Parayil, J. Zhang, M. Wang, **A. Koppel**, “On the Sample Complexity and Metastability of Heavy-tailed Policy Search in Continuous Control,” in *Journal of Machine Learning Research*, Jan. 2024.
5. **A. Koppel**, J. Eappen, S. Bhatt, C. Hawkins, S. Ganesh. “Online MCMC Thinning with Kernelized Stein Discrepancy” in *SIAM Journal on Mathematics of Data Science (SIMODS)*, Jan 2024. Available as ArXiv 2201.07130
6. Q. Bai, A. S. Bedi, M. Agarwal, **A. Koppel**, V. Aggarwal, “Achieving Zero Constraint Violation for Constrained Reinforcement Learning via Primal-Dual Approach,” in *Journal of Artificial Intelligence Research (JAIR)*, Dec. 2023.
7. H. Kumar, **A. Koppel**, and A. Ribeiro. “On the Sample Complexity of Actor-Critic Method for Reinforcement Learning with Function Approximation” in *Machine Learning*, p. 1-35 Springer. Feb 2023.
8. A. Chakraborty, K. Rajawat and **A. Koppel**, “Sparse Representations of Positive Functions via First- and Second-Order Pseudo-Mirror Descent,” in *IEEE Transactions on Signal Processing*, vol. 70, pp. 3148-3164, 10.1109/TSP.2022.3173146. May 2022
9. E. Noorani, Y. Savas, **A. Koppel**, J. S. Baras, U. Topcu, B. M. Sadler, “Collaborative Beamforming Under Localization Errors: A Discrete Optimization Approach,” in *Elsevier Signal Processing*, Volume 200, Nov. 2022, 108647.
10. E. Zobeidi, **A. Koppel**, and N. Atanasov, “Dense Incremental Metric-Semantic Mapping via Sparse Gaussian Process Regression,” in *IEEE Transactions on Robotics (T-RO)*, 10.1109/TRO.2022.3168733, May 2022.
11. Z. Gao, **A. Koppel**, and A. Ribeiro, “Balancing Rates and Variance via Adaptive Batch-Size for Stochastic Optimization Problems,” *IEEE Trans. Signal Process.*, 10.1109/TSP.2022.3186526. (to appear), Jun. 2022.
12. A. S. Bedi, K. Rajawat, V. Aggarwal, **A. Koppel**. “Escaping Saddle Points in Successive Convex Approximation,” in *IEEE Trans. Signal Processing* Volume: 70, pp. 307-321 , Issue: 3, May 2022. DOI: 10.1109/TSP.2021.3138242.
13. **A. Koppel**, A. S. Bedi, V. Elvira, and B.M. Sadler. “Nearly Consistent Finite Particle Estimates in Streaming Importance Sampling,” in *IEEE Trans. Signal Processing*, 2022. DOI: 10.1109/TSP.2021.3120512, Dec 2021.
14. **A. Koppel**, H. Pradhan, and K. Rajawat. “Consistent Online Gaussian Process Regression Without the Sample Complexity Bottleneck,” in *Statistics and Computing*, Springer, Sept. 2021.
15. J. Zhang, A. S. Bedi, M. Wang, and **A. Koppel**. “Cautious Reinforcement Learning via Distributional Risk in the Dual Domain” in *IEEE Journal on Selected Areas in Information Theory: Special Issue on Sequential, Active, and Reinforcement Learning*, May 2021.
16. R. Pradhan, A. S. Bedi, **A. Koppel**, and K. Rajawat. “Adaptive Kernel Learning in Heterogeneous Networks ” in *IEEE Trans. Signal and Info. Process. over Networks*, 2021.
17. D. S. Kalhan, A. S. Bedi, **A. Koppel**, K. Rajawat, H. Hassani, A. Gupta, and A. Banerjee. “Dynamic Online Learning via Frank-Wolfe Algorithm” in *IEEE Trans. Signal Process.* 2021.
18. A. S. Bedi, **A. Koppel**, P. Sanyal, and K. Rajawat. “Nonparametric Compositional Stochas-

- tic Optimization: Algorithms for Robust Online Learning with Kernels,” in *IEEE Trans. Signal Process.* 2021.
19. A. Mokhtari and **A. Koppel**. “High-Dimensional Nonconvex Stochastic Optimization by Doubly Stochastic Successive Convex Approximation” in *IEEE Trans. Signal Process.*, Volume 68 (2020): 6287-6302. Oct. 2020.
 20. Y. Tian, **A. Koppel**, and A. S. Bedi, and J. How, “Asynchronous and Parallel Distributed Pose Graph Optimization,” in *IEEE Robotics and Automation Letters* , Volume 5, no. 4 (2020): 5819-5826, Oct. 2020., **Honorable Mention for Best Paper Award, IEEE Robotics and Automation Society**
 21. A. Mokhtari, **A. Koppel**, M. Takáč, and A. Ribeiro, “A Class of Doubly Random Parallel Stochastic Methods for Large Scale Learning,” in *Journal of Machine Learning Research*, 21(120):1-51, 2020
 22. K. Zhang, **A. Koppel**, H. Zhu, and T. M. Başar. “Global Convergence of Policy Gradient Methods to (Almost) Locally Optimal Policies” in *SIAM Journal on Control and Optimization* DOI 10.1137/19M1288012 , Dec. 2020.
 23. **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*, DOI 10.1109/TAC.2020.3029315. Volume 66, Issue: 4 Pages: 1856-1863, April 2021.
 24. **A. Koppel**, A. S. Bedi, K. Rajawat, and B.M. Sadler. “Optimally Compressed Nonparametric Online Learning: Tradeoffs between memory and consistency,” in *IEEE Signal Processing Magazine*, Volume 37 , Issue: 3, May 2020.
 25. 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*, Volume 5, no. 3 (2019): 479-494. 2019
 26. **A. Koppel**, K. Zhang, H. Zhu, and T. M. Başar. “Projected Stochastic Primal-Dual Method for Constrained Online Learning with Kernels” in *IEEE Trans. Signal Process.*, Volume 67, no. 10 (2019): 2528-2542, May 2019.
 27. A. S. Bedi, **A. Koppel**, and K. Rajawat, “Asynchronous Online Learning in Multi-Agent Systems with Proximity Constraints,” *IEEE Transactions on Signal and Information Processing over Networks*, Volume 5, no. 3, pp. 479-494, Sept. 2019
 28. **A. Koppel**, S. Paternain, C. Richard, and A. Ribeiro. ”Decentralized online learning with kernels.” *IEEE Trans. Signal Process*, Volume 66, no. 12 (2018): 3240-3255.
 29. **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
 30. **A. Koppel**, B. Sadler, and A. Ribeiro, “Proximity without Consensus in Online Multi-Agent Optimization,” in *IEEE Trans. Signal Proc*, Volume 65 No. 12 , Page 3062-3077, June 15, 2017.
 31. **A. Koppel**, G. Warnell, E. Stump, and A. Ribeiro, “D4L: Decentralized Dynamic Discriminative Dictionary Learning,” *IEEE Trans. Signal Info. Process. over Networks*, Volume 3, no. 4 (2017): 728-743..
 32. 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, No. 11. Nov, 2017.
 33. 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*, Volume 64, no. 17 (2016): 4576-4591.

34. **A. Koppel**, F. Jakubeic, and A. Ribeiro, "A saddle point algorithm for networked online convex optimization," *IEEE Trans. Signal Process.*, Volume 63, no. 19 (2015): 5149-5164.

Conference Papers

1. A. Lahoti, S. Senapati, K. Rajawat, **A. Koppel**, "Sharpened Lazy Incremental Quasi-Newton Method" in 2024 Artificial Intelligence and Statistics (AISTATS), preprint available as arXiv:2305.17283 Search... .
2. S. Chakraborty, A. Bedi, **A. Koppel**, H. Wang, D. Manocha, M. Wang, F. Huang, "PARL: A Unified Framework for Policy Alignment in Reinforcement Learning" in 2024 International Conference on Learning Representations (ICLR), preprint available as arXiv:2308.02585v2.
3. M. Zaman, M. Laurière, **A. Koppel**, T. Başar, "Robust Cooperative Multi-Agent Reinforcement Learning: A Mean-Field Type Game Perspective" in 2024 Learning for Dynamics and Control (L4DC)
4. D. Goktas, A. Greenwald, S. Zhao, **A. Koppel**, S. Ganesh "Generative Adversarial Inverse Multiagent Learning" in 2024 International Conference on Learning Representations (ICLR), preprint available on Openreview.
5. D. Ying, Y. Ding., Y. Zhang, **A. Koppel**, J. Lavaei, "Scalable Primal-Dual Actor-Critic Method for Safe Multi-Agent Reinforcement Learning with General Utilities" in 2023 Neural Information Processing Systems (NeurIPS).
6. A. Mishler, Mohsen Ghassemi, **A. Koppel**, S. Ganesh, "Model Robustness and Active Learning with Missing-Not-At-Random Outcomes," in 2023 UAI Workshop on Epistemic Uncertainty in Artificial Intelligence
7. S. Chakraborty, A.S. Bedi, P. Tokekar, **A. Koppel**, M. Wang, F. Huang "Principal-Driven Reward Design and Agent Policy Alignment via Bilevel-RL" in ILHF Workshop at ICML 2023
8. H. He, **A. Koppel**, A. S. Bedi, D. Stilwell, M. Farhood, B. Biggs, "Bi-Level Nonstationary Kernels for Online Gaussian Process Regression" in 2023 *IEEE 19th International Conference on Automation Science and Engineering (CASE)*, Aug 26 - 30, 2023.
9. W. A. Suttle, A.S. Bedi, B. Patel, B. M. Sadler, **A. Koppel**, D. Manocha, "Beyond Exponentially Fast Mixing in Average-Reward Reinforcement Learning via Multi-Level Monte Carlo Actor-Critic" in 2023 International Conference on Machine Learning (ICML)
10. S. Chakraborty, A.S. Bedi, P. Tokekar, **A. Koppel**, M. Wang, F. Huang, D. Manocha, "STEERING: Stein Information Directed Sampling for Efficient Exploration in Model-Based Reinforcement Learning" in 2023 International Conference on Machine Learning (ICML)
11. M. A. Zaman, M. Laurière, **A. Koppel**, T. Başar, "Receding Horizon Policy Gradient for Zero-Sum Mean-Field Type Games," in 2023 57th Annual Conference on Information Sciences and Systems (CISS)
12. W. Suttle, **A. Koppel**, J. Liu, "Information-Directed Policy Search in Sparse-Reward Settings via the Occupancy Information Ratio," in 2023 57th Annual Conference on Information Sciences and Systems (CISS)
13. A. S. Bedi, C. Fan, **A. Koppel**, A. K. Sahu, F. Huang, D. Manocha, "Federated Learning Beyond Consensus," in 2023 57th Annual Conference on Information Sciences and Systems (CISS)
14. D. Ying, Y. Ding., **A. Koppel**, J. Lavaei, "Scalable Multi-Agent Reinforcement Learning with General Utilities" in 2023 *IEEE American Control Conference (ACC)*.
15. S. Chakraborty, A.S. Bedi, P. Tokekar, **A. Koppel**, B.M. Sadler, F. Huang, D. Manocha, "Posterior Coreset Construction with Kernelized Stein Discrepancy for Model-Based Reinforcement Learning" in 2023 Conference on Artificial Intelligence (AAAI). arXiv preprint arXiv:2206.01162.

16. W. A. Suttle, **A. Koppel**, J. Liu, "Occupancy Information Ratio: Infinite-Horizon, Information-Directed, Parameterized Policy Search," Workshop paper at 2023 AAAI Symposium on Machine Learning for Dynamical Systems.
17. S. Chakraborty, A.S. Bedi, P. Tokekar, **A. Koppel**, B.M. Sadler, F. Huang, D. Manocha, "Stein Information Directed Sampling for Efficient Exploration in Model-Based Reinforcement Learning" Workshop paper at 2023 AAAI Workshop on Reinforcement Learning Ready for Production.
18. M. A. Zaman, **A. Koppel**, S. Bhatt, and T. Başar. "Oracle-free Reinforcement Learning in Mean-Field Games along a Single Sample Path" in 2023 Artificial Intelligence and Statistics (AISTATS), arXiv preprint arXiv:2208.11639.
19. S. Chakraborty, A.S. Bedi, K. Mudiyansele, P. Poddar, **A. Koppel**, P. Tokekar, D. Manocha, "Dealing with Sparse Rewards in Continuous Control Robotics via Heavy-Tailed Policy Optimization" in *International Conference on Robotics and Automation (ICRA)*, May 29 - June 2, 2023.
20. H. He, **A. Koppel**, A. S. Bedi, D. Stilwell, M. Farhood, B. Biggs, "Decentralized Multi-agent Exploration with Limited Inter-agent Communications" in *International Conference on Robotics and Automation (ICRA)*, May 29 - June 2, 2023.
21. Y. Tian, A. S. Bedi, **A. Koppel**, M. Calvo-Fullana, D. M. Rosen, J. P. How, "Distributed Riemannian Optimization with Lazy Communication for Collaborative Geometric Estimation," in *2022 IEEE Conference on Intelligent Robotics and Systems (IROS)*. Available as arXiv:2203.00851.
22. **A. Koppel**, A. S. Bedi, B. Ganguly, V. Aggarwal, "Convergence Rates of Average-Reward Multi-agent Reinforcement Learning via Randomized Linear Programming," in *2022 IEEE Conference on Decision and Control (CDC)*.
23. Q. Jin, **A. Koppel**, K. Rajawat, A. Mokhtari "Sharpened Quasi-Newton Methods: Faster Superlinear Rate and Larger Local Convergence Neighborhood" **Spotlight** in *International Conference on Machine Learning (ICML)*, July 19-21, 2022.
24. A. S. Bedi, S. Chakraborty, A. Parayil, B. M. Sadler, P. Tokekar, **A. Koppel**, "On the Hidden Biases of Policy Mirror Ascent in Continuous Action Spaces," **Spotlight** in *International Conference on Machine Learning (ICML)*, July 19-21, 2022
25. H. Pradhan, **A. Koppel**, K. Rajawat, "On Submodular Set Cover Problems For Near-Optimal Online Kernel Basis Selection," in *2022 IEEE International Conference on Acoustics, Speech, & Signal Processing (ICASSP)*, May 2022.
26. J. Di, E. Zobeidi, A. Koppel, N. Atanasov, "Distributed Gaussian Process Mapping for Robot Teams with Time-varying Communication," in *2022 IEEE American Control Conference (ACC)*, June 2022.
27. W. Suttle, **A. Koppel**, and J. Liu, "Policy gradient for ratio optimization: a case study," in *2022 IEEE Conference on Information Science and Systems (CISS)*.
28. Q. Bai, A. S. Bedi, M. Agarwal, **A. Koppel**, V. Aggarwal, "Achieving Zero Constraint Violation for Constrained Reinforcement Learning via Primal-Dual Approach" in 2022 Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22).
29. J. Zhang, A. S. Bedi, M. Wang, **A. Koppel**, "Multi-Agent Reinforcement Learning with General Utilities via Decentralized Shadow Reward Actor-Critic" in 2022 Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22).
30. **A. Koppel**, J. Zhang, A. S. Bedi, M. Wang, "Intermittent Communications in Decentralized Shadow Reward Actor-Critic," in *IEEE Conference on Decision and Control (CDC)*, Austin, TX, Dec. 13-15 2021
31. S. Bhatt, W. Mao, **A. Koppel**, T. Başar, "Semiparametric Information State Embedding for

- Policy Search under Imperfect Information,” in *IEEE Conference on Decision and Control* (CDC), Austin, TX, Dec. 13-15 2021
32. A. Chakraborty, K. Rajawat, and **A. Koppel**, “Projected Pseudo-Mirror Descent in Reproducing Kernel Hilbert Space” in *Asilomar Conference on Signals, Systems, and Computers*, Oct. 31 - Nov 3rd, 2021.
 33. E. Noorani, Y. Savas, **A. Koppel**, J. S. Baras, U. Topcu, B. M. Sadler, “Collaborative Beamforming for Agents with Localization Errors,” in *Asilomar Conference on Signals, Systems, and Computers*, Oct. 31 - Nov 3rd, 2021.
 34. **A. Koppel**, A. S. Bedi, B. Ganguly, and V. Aggarwal, “Randomized Linear Programming for Tabular Average-Cost Multi-agent Reinforcement Learning” in *Asilomar Conference on Signals, Systems, and Computers*, Oct. 31 - Nov 3rd, 2021.
 35. M. Kepler, **A. Koppel**, A. S. Bedi, D. Stillwell “Wasserstein-Splitting Gaussian Process Regression for Heterogeneous Online Bayesian Inference,” in *IEEE International Conference on Intelligent Robotics and Systems(IROS)*, Sept. 27-Oct. 1, 2021.
 36. **A. Koppel**, A. S. Bedi, and V. Krishnamurthy, “A Dynamical Systems Perspective on Online Bayesian Nonparametric Estimators with Adaptive Hyperparameters,” in *Int. Conf. Acoustics Speech Signal Process (ICASSP)*, 6-11 June 2021.
 37. A. Parayil, A. S. Bedi, and **A. Koppel**, “Joint Position and Beamforming Control via Alternating Nonlinear Least-Squares with a Hierarchical Gamma Prior,” in *IEEE American Control Conf. (ACC)*, May 26-28, 2021
 38. J. Zhang, A. S. Bedi, M. Wang, and **A. Koppel**, “Beyond Cumulative Returns via Reinforcement Learning over State-Action Occupancy Measures,” in *IEEE American Control Conf.*, May 26-28, 2021
 39. Z. Gao, **A. Koppel**, and A. Ribeiro, “Incremental Greedy BFGS: An Incremental Quasi-Newton Method with Explicit Superlinear Rate,” in *Neural Information Processing Systems (NeurIPS)*, Workshop on Optimization for Machine Learning, **Spotlight**, 11-12 Dec. 2020
 40. J. Zhang, **A. Koppel**, A. S. Bedi, C. Szepesvári, M. Wang, “Variational Policy Gradient Method for Reinforcement Learning with General Utilities,” in *Neural Information Processing Systems (NeurIPS)*, **Spotlight**, 6-12 Dec. 2020
 41. H. Pradhan, A. Bedi, **A. Koppel**, K. Rajawat, “Conservative Multi-agent Online Kernel Learning in Heterogeneous Networks,” in *IEEE Proc. Asilomar Conf. Signals, Systems, Computers*, Pacific Grove, CA, Nov. 8-11, 2020.
 42. E. Zobeidi, **A. Koppel**, and N. Atanasov, “Dense Incremental Metric-Semantic Mapping via Sparse Gaussian Process Regression,” in *IEEE Proc. Int. Conf. Intelligent Robotics and Systems (IROS)*, Las Vegas, NV, Oct. 25-9, 2020.
 43. A.S. Bedi, **A. Koppel**, K. Rajawat, B.M. Sadler, “Trading Dynamic Regret for Model Complexity in Nonstationary Nonparametric Optimization,” in *IEEE American Control Conference*, Denver, Colorado, Jul. 1-3, 2020.
 44. A. S. Bedi, D. Peddireddy, V. Aggarwal, **A. Koppel**, “Efficient Large-Scale Gaussian Process Bandits by Believing only Informative Actions,” in 2nd Annual Conference on Learning for Dynamics and Control (L4DC), Berkeley, CA, Jun. 6-9, 2020.
 45. D. S. Kalhan, A. S. Bedi, **A. Koppel**, K. Rajawat, A. Gupta, and A. Banerjee, “Projection Free Dynamic Online Learning,” in *IEEE Proc. Int. Conf. Acoustics Speech Signal Process (ICASSP)*, Barcelona, Spain, May. 4-8, 2020.
 46. Z. Gao, **A. Koppel**, and A. Ribeiro, “Balancing Rates and Variance via Adaptive Batch-Sizes in First-Order Stochastic Optimization,” in *IEEE Proc. Int. Conf. Acoustics Speech Signal Process (ICASSP)*, Barcelona, Spain, May. 4-8, 2020.

47. **A. Koppel**, A.S. Bedi, V. Elvira, B.M. Sadler, "A Projection Operator to Balance Consistency and Complexity in Importance Sampling," in *Neural Information Processing Systems, Symposium on Advances in Approximate Bayesian Inference*, Vancouver, Canada, Dec. 8, 2019.
48. H. Kumar, **A. Koppel**, Alejandro Ribeiro, "On the Sample Complexity of Actor-Critic for Reinforcement Learning," in *Neural Information Processing Systems, Workshop on Optimization Foundations of Reinforcement Learning*, Vancouver, Canada, Dec. 14, 2019.
49. S. Bhatt, **A. Koppel**, V. Krishnamurthy, "Policy Gradient using Weak Derivatives for Reinforcement Learning," in *IEEE Conference on Decision and Control (CDC)*, Nice, France, Dec. 11-13, 2019.
50. K. Zhang, **A. Koppel**, H. Zhu, T. Başar, "Convergence and Iteration Complexity of Policy Gradient Methods for Infinite-horizon Reinforcement Learning," in *IEEE Conference on Decision and Control (CDC)*, Nice, France, Dec. 11-13, 2019.
51. 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.
52. K. Zhang, **A. Koppel**, H. Zhu, T. Başar, "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.
53. **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.
54. **A. Koppel**, "Consistent Online Gaussian Process Regression Without the Sample Complexity Bottleneck," in *American Control Conference*, Philadelphia, PA, July 10-12, 2019.
55. H. Pradhan, A. S. Bedi, **A. Koppel**, and K. Rajawat, "Exact Decentralized Online Non-parametric Optimization," in *IEEE Global Conf. on Signal and Info. Processing (GlobalSIP)*, Anaheim, CA, Nov. 26-28, 2018.
56. **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.
57. 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.
58. K. Zhang, H. Zhu, T. Başar, 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.
59. 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.
60. 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.
61. **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.
62. **A. Koppel**, S. Paternain, C. Richard, and A. Ribeiro, "Decentralized Efficient Nonparametric Stochastic Optimization", in *IEEE Global Conference on Signal and Information Process-*

ing, Montreal, Canada, Nov. 14-16, 2017.

63. 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. **Best Paper Finalist, IEEE Signal Processing Society**
64. M. Fazylib, **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.
65. 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.
66. **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.
67. **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.
68. **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.
69. **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
70. 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.
71. 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.
72. **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.
73. 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.
74. **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.
75. 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.
76. **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.
77. **A. Koppel**, G. Warnell, E. Stump, and A. Ribeiro, "D4L: Decentralized Dynamic Discriminative Dictionary Learning," in *Proc. Int. Conf. Intelligent Robotics and Systems*, Hamburg, Germany, Sep 28-Oct2 2015.

78. **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.
79. **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.

Technical Reports and Theses

1. **A. Koppel**, "Stochastic Optimization For Multi-Agent Statistical Learning And Control." Phd Dissertation, Dept. of Electrical and Systems Engineering, University of Pennsylvania, June 2017.
2. **A. Koppel**, "Parameter Estimation in High-Dimensions using Doubly Stochastic Approximation." MS Thesis, Dept. of Statistics, The Wharton School, University of Pennsylvania, May 2017.
3. V. Ganesan, **A. Koppel**, S. Han, J. Conroy, A. Wickenden, R. Murray, and W. Nothwang. "Implementation and Validation of Bioplausible Visual Servoing Control." ARL-TR- 6387; U.S. Army Research Laboratory: Adelphi, MD, March 2013.
4. **A. Koppel**, E. Stump, W. Nothwang, and B. Sadler. "An Adaptive Stochastic Differential System for Multi-Agent Coordination." Army Research Laboratory Technical Report. Aug. 2012.
5. **A. Koppel**, V. Ganesan, A. Wickenden, W. Nothwang. "Slow Computing Simulation of Bioplausible Control." ARL-TR-5959; U.S. Army Research Laboratory: Adelphi, MD, March 2012
6. **A. Koppel** and R. Feres. "Stochastic Methods for the Lotka-Volterra Model with Migration." Bachelors Honors Thesis. Washington University in St. Louis, Mar. 2011.

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: Undergraduates**

- Andi Johnson, “Pseudo-input search in Online Gaussian Processes” Summer 2020
Affiliation: BA student, Math, Eastern New Mexico University
- Gisselle Contreras, “Adapting Hyper-parameters in Kernel Regression” Summer 2020
Affiliation: BA student, ECE, UT San Antonio
- Dylan Scott, “Adaptive Gradient Methods for Non-stationary Learning” Summer 2020
Affiliation: BA student, CS, Hampton University

ARL Research Associates: Graduate Fellows

- Bingjia Wang, “Compressed GPs for Policy Search” Summer 2020 - present
Affiliation: MS student, ECE, Cornell University
- James Berneburg, “RL for Event-Triggered Control” Summer 2020 - present
Affiliation: Phd student, ECE, George Mason University
- Kaiqing Zhang, “Non-convex optimization approaches to policy search” Summer 2018
Affiliation: Phd student, ECE University of Illinois Urbana-Champaign
- Hrusiksha Pradhan, “Compressions of Gaussian Processes” Fall 2018 - present
Affiliation: Phd student, EE, India Institute of Technology, Kanpur
- Sujay Bhatt, “Variance-reduced policy gradients via weak-derivatives” Fall 2018
Affiliation: Phd student, ECE, Cornell University

ARL Research Associates: Postdoctoral Fellows

- Amrit Singh Bedi, “Risk in Bayesian & Reinforcement Learning” Winter 2019 - present
- Anjaly Parayil, “Control Design for Decentralized Beamforming” Winter 2020 - present

Other Mentorship Experience: Graduate Fellows

- Yagiz Savas, “Beamforming with Localization Errors” Summer 2019 - present
Affiliation: Phd student, University of Texas at Austin
- Erfan Noorani, “Beamforming with Localization Errors” Summer 2019 - present
Affiliation: Phd student, University of Maryland, College Park
- Ehsan Zobeidi, “Gaussian Processes for Occupancy Mapping” Spring 2019 - present
Affiliation: Phd student, University of California, San Diego
- Zhan Gao, “Covarying Step and Batch-sizes in Learning” Spring 2019 - present
Affiliation: Phd student, ESE, UPenn
- Harshat Kumar, “Sample Complexity of Actor-Critic in RL” Summer 2018 - Fall 2019
Affiliation: Phd student, ESE, UPenn
- Mathew Rittonia, “ Defenses to Black Box Data Poisoning ” Spring - Summer 2019
Affiliation: MS student, AMS, Johns Hopkins University
- Ekaterina Tolstaya, “Compressed Kernel Q Learning” Fall 2016 - Spring 2018
Affiliation: Phd student, ESE, UPenn

REFERENCES

Alejandro Ribeiro

Rosenbluth Associate Professor
Dept. of Electrical & Systems Eng.
University of Pennsylvania
Room 411B, Warren Center for Data and Network
Sciences
3401 Walnut Street, 4th floor
Philadelphia, PA 19104
Phone: (612) 889-9217
Email: aribeiro@seas.upenn.edu

Tamer Başar

Director, Center for Advanced Study
Swanlund Endowed Chair & CAS Professor
Dept. of Electrical & Computer Eng.
University of Illinois
356 Coordinated Science Lab MC 228
1308 W. Main St.
Urbana Illinois 61801
Phone: (217) 333-3607
Email: basarl@illinois.edu

Prof. Victor Elvira

Reader (Associate Prof.) in Stat. & Data Science
School of Mathematics
University of Edinburgh
James Clerk Maxwell, King's Buildings,
Peter Guthrie Tait Road
Edinburgh, EH9 3FD, UK
Email: victor.elvira@ed.ac.uk

Brian M. Sadler

IEEE Fellow
Intelligent Systems ST
U.S. Army Research Laboratory
Adelphi Laboratory Center
2800 Powder Mill Rd.
Adelphi, MD 20793
Phone: (301) 394-1239
Email: brian.m.sadler6.civ@mail.mil

Prof. Vikram Krishnamurthy

IEEE Fellow
Dept. of Electrical & Computer Eng.
Cornell University
134 Rhodes Dr.
Ithaca, NY 14850
Phone: (607) 255-6912
Email: vikramk@cornell.edu

Prof. Ketan Rajawat

Dept. of Electrical & Computer Eng.
Indian Institute of Technology Kanpur
Kanpur UP 208016 INDIA
+91-512-259-7337,
Email: ketan@iitk.ac.in