

ERIK MIEHLING

mail: emiehling@gmail.com
web: <https://emiehling.github.io>

work experience

IBM Research *Dublin, Ireland*
Research Scientist (AI) **Aug 2022 – present**

- Contributed to emerging AI governance policy through development of human feedback elicitation mechanisms combining natural language explanations and knowledge distillation
- Pioneered and prototyped new approaches for studying algorithmic amplification of bias in RL and ML pipelines
- Played a key role in seeking government funding through involvement in research ideation and proposal writing

University of Illinois at Urbana–Champaign *Urbana, IL*
Postdoctoral Research Associate **Feb 2018 – Aug 2022**

- Advised 10 Ph.D. students and published 15 peer-reviewed articles in multi-agent reinforcement learning and game theory spanning five federally funded projects (total research funding: \$39 500 000 USD)
- Made foundational contributions to multi-agent reinforcement learning in both cooperative domains (RNN-based information embeddings) and adversarial domains (online attacker intent inference for cyber-security)
- Co-wrote a successful NSF grant valued at \$500 000 USD which funded a three year research program on modeling, learning, and control of epidemic processes

education

University of Michigan *Ann Arbor, MI*
Ph.D. – Electrical Engineering & Computer Science **Sept 2011 – Dec 2017**

University of British Columbia *Vancouver, Canada*
M.A.Sc. – Electrical & Computer Engineering **Sept 2009 – Aug 2011**
B.A.Sc. – Electrical Engineering **Sept 2006 – May 2009**

research experience

Horizon Europe: AutoFair – Human-Compatible Artificial Intelligence with Guarantees **Oct 2022 – present**

- Collaborated with multiple academic and industrial research institutions on problems of algorithmic fairness.
- Developed simulation-based models of dynamic fairness, human-aligned foundation model selection strategies, and contributed to AI liability policy.

C3.ai: Algorithms and Software Tools for Testing and Control of COVID-19 **Dec 2020 – Aug 2022**

- Initiated an interdisciplinary collaboration with epidemiologists to construct data-driven models for explaining the spread of COVID-19.
- Developed control-theoretic models for efficient allocation of testing kits in epidemics.

NSF: A Comprehensive Approach to Modeling, Learning, Analysis and Control of Epidemic Processes Over Time-Varying and Multi-Layer Networks **Oct 2020 – Aug 2022**

- Employed mean-field game theory to construct an agent-based behavioral model for describing asymptomatic viral spread.

ARO: A Multimodal Approach to Network Information Dynamics **Sept 2020 – Aug 2022**

- Constructed game-theoretic models describing how misinformation/disinformation emerges, spreads, and is sustained.

ARL: Internet of Battlefield Things Collaborative Research Alliance Feb 2019 – Aug 2022

- Employed recurrent neural networks (RNNs) to discover low-dimensional summaries of sufficient information for decentralized partially observable Markov decision processes.
- Developed game-theoretic conflict models for large-scale, uncertain, and time-critical battle-field environments.

ONR: An Analytical Framework for Actionable Defense against Advanced Persistent Threats (ADAPT) Feb 2018 – Jun 2021

- Designed provably-convergent reinforcement learning algorithms for cooperative and non-cooperative multi-agent decision problems.
- Developed game-theoretic models for online intent inference in adversarial domains.

ARO: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense Jan 2014 – Dec 2017

- Design and analysis of sequential decision-making models under both probabilistic and non-probabilistic uncertainty.
- Application to the development of dynamic defense schemes for large-scale networks.
- Extended the theory of structural results in partially observable Markov decision processes.

NSF: Foundations of Resilient Cyber-Physical Systems (FORCES) Jan 2013 – Dec 2017

- Developed provably convergent algorithms for a class of non-convex problems.
- Application to the design of decentralized electricity market mechanisms.

research
advising

PhD students

- Shubham Aggarwal (with T. Başar): mean-field games 2021 – 2022
- Raj Kiriti Velicheti (with T. Başar): misinformation, multi-agent RL 2021 – 2022
- Xiaoqi Bi (with C. Beck): epidemics, stochastic control 2020 – 2022
- S. Yagiz Olmez (with P. G. Mehta): epidemics, game theory 2020 – 2022
- Jameson Mori (with R. L. Smith): epidemics 2020 – 2021
- Weichao Mao (with T. Başar): multi-agent RL 2019 – 2022
- Kaiqing Zhang (with T. Başar): multi-agent RL 2018 – 2020
- Xiangyuan Zhang (with T. Başar): multi-agent RL, game theory 2018 – 2022
- Aneeq Zaman (with T. Başar): mean-field games 2018 – 2022
- Muhammed Sayin (with T. Başar): security, game theory 2018 – 2019

Masters students

- Danilo Dordevic (ETH Zürich, with M. Hudoba de Badyn): math finance 2022

teaching
experience

University of Illinois at Urbana–Champaign Urbana, IL

ECE 486 - Control Systems, *Guest Lecturer* Apr 2019

University of Michigan, College of Engineering Ann Arbor, MI

EECS 215 - Introduction to Electronic Circuits, *Discussion Lecturer* Winter 2013

EECS 216 - Signals and Systems, *Laboratory Lecturer* Fall 2012

internships

US Department of Energy, Oak Ridge National Laboratory <i>Advanced Short-Term Research Opportunity</i>	<i>Oak Ridge, TN</i> May 2013 – Sept. 2013
Defense Research & Development Canada, Radar Systems Group <i>Summer Internship</i>	<i>Ottawa, ON</i> May 2010 – Sept. 2010
Broadcom Canada, Centre for Audio Excellence <i>Summer Internship</i>	<i>Richmond, BC</i> May 2008 – Sept. 2008

**talks &
presentations**

“Mixed Signals: Information Disclosure in Persuasion, Security, and Value Alignment” <i>IBM Research, Dublin, Ireland</i>	<i>April 2022</i>
“Estimation, Modeling, and Active Testing of Uncertain Epidemic Processes” <i>C3.ai Digital Transformation Institute Research Symposium, Miami, FL</i>	<i>Mar 2022</i>
“Optimal Transport and Boundary-Constrained Control” <i>Internet of Battlefield Things (IoBT) Research Presentation (remote)</i>	<i>Feb 2022</i>
“Secure Contingency Prediction and Response for Cyber-Physical Systems” <i>4th IEEE Conf. on Control Technology and Applications, Montreal, QC, Canada (remote)</i>	<i>Aug 2020</i>
“Strategic Inference With a Single Private Sample” <i>58th IEEE Conf. on Decision and Control, Nice, France</i>	<i>Dec 2019</i>
“Optimal Switching Strategies for System Obfuscation” <i>Internet of Battlefield Things (IoBT) Bootcamp, Chicago, IL</i>	<i>Jul 2019</i>
“Online Planning for Decentralized Stochastic Control with Partial History Sharing” <i>American Control Conf. (ACC) 2019, Philadelphia, PA</i>	<i>Jul 2019</i>
“Strategic Inference Under Private Sampling” <i>EC 2019 Workshop: Learning in the Presence of Strategic Behavior, Phoenix, AZ</i> <i>8th Midwest Workshop on Control and Game Theory, St. Louis, MO</i>	<i>Jun 2019</i> <i>Apr 2019</i>
“A Bayesian Multi-armed Bandit Approach for Identifying Human Vulnerabilities” <i>Decision and Game Theory for Security (GameSec 2018)</i>	<i>Oct 2018</i>
“Monotonicity Properties of Optimal Policies for POMDPs on Partially Ordered Spaces” <i>Department of Electrical Engineering, University of Washington, Seattle, WA</i> Host: Lillian Ratliff <i>Stanford Intelligent Systems Laboratory Seminar, Stanford University</i> Hosts: Katherine Driggs-Campbell, Mykel J Kochenderfer	<i>Apr 2018</i> <i>Jan 2018</i>
“A Dependency Graph Formalism for the Dynamic Defense of Cyber Networks” <i>5th IEEE Global Conf. on Signal and Information Processing (GlobalSIP 2017)</i>	<i>Nov 2017</i>
“A Stochastic Control Approach to Dynamic Defense of Large-Scale Cyber Networks” <i>ECE Systems & Control Seminar, University of New Mexico</i> Host: Meeko Oishi	<i>Nov 2017</i>
“A Dependency Graph Formalism for the Dynamic Defense of Large-Scale Cyber Networks” <i>3rd Multidiscip. Conf. on Reinforcement Learning & Decision Making (RLDM 2017)</i>	<i>Jun 2017</i>
“A Decentralized Mechanism for Computing Competitive Equilibria in Deregulated Electricity Markets” <i>American Control Conf. (ACC 2016)</i>	<i>Jul 2016</i>
“A Formal Model for the Dynamic Defense of Cyber Networks” <i>2nd Moving Target Defense Quantification Workshop (MTDQ)</i>	<i>Jun 2016</i>
“Optimal Defense Policies for Partially Observable Spreading Processes on Bayesian Attack Graphs” <i>2nd ACM Workshop on Moving Target Defense (MTD)</i>	<i>Oct 2015</i>

Organizational & Outreach

- IBM Research – speaker at Technical Vitality Day 2022
- NSF – grant review panelist, Spring 2022 session
- CDC 2021 – session co-organizer (modeling, prediction, and control of epidemic processes)
- CDC 2020 – session co-chair (epidemics)
- CDC 2019 – session co-chair (machine learning in complex networks II)
- Allerton 2019 – session assistant (performance, reliability, and security)
- ACC 2019 – volunteer
- SPS 2019 – technical program committee member
- CDC 2018 – volunteer
- GameSec 2018 – session chair (advanced persistent threats)
- Allerton 2018 – session assistant (dynamic games); tutorial session assistant (real time optimization for estimation and control)
- ACC 2018 – volunteer

Reviewer

- ACM/IEEE International Conference on Cyber-Physical Systems
- ACM Transactions on Cyber-Physical Systems
- ACM Workshop on Moving Target Defense
- American Control Conference (ACC)
- Computers & Security
- Dynamics Games and Applications (Springer)
- IEEE Conference on Decision and Control (CDC)
- IEEE International Workshop on Signal Processing Advances in Wireless Communications
- IEEE Power & Energy Society General Meeting
- IEEE Transactions on Information Forensics & Security
- IEEE Transactions on Power Systems
- IEEE Transactions on Reliability
- IEEE/PES Transmission and Distribution Conference and Exposition
- Indian Control Conference (ICC)
- Springer/Birkhäuser
- Systems & Control Letters

Professional Membership

- Institute of Electrical and Electronics Engineers (IEEE)

Journal

- A. Zaman, E. Miehling, and T. Başar. Reinforcement Learning for Non-Stationary Discrete-Time Linear-Quadratic Mean-Field Games in Multiple Populations. *Dynamic Games and Applications*, vol. 13, no. 1, pp. 118-164, 2023.
- E. Miehling and D. Teneketzis. Monotonicity Properties for Two-Action Partially Observable Markov Decision Processes on Partially Ordered Spaces. *European Journal of Operational Research*, vol. 282, no. 3, pp. 936-944, 2019.
- E. Miehling, M. Rasouli and D. Teneketzis. A POMDP Approach to the Dynamic Defense of Large-Scale Cyber Networks. *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 10, pp. 2490–2505, 2018.

- V. Krishnamurthy, R.R. Bitmead, M. Gevers, and E. Miehling. Sequential Detection with Mutual Information Stopping Cost. *IEEE Transactions on Signal Processing*, vol. 60, no. 2, pp. 700–714, 2012.

Conference & Workshop

- E. Miehling, R. Nair, E. Daly, K. N. Ramamurthy, R. Redmond. Cookie Consent Has Disparate Impact on Estimation Accuracy. In *Advances in Neural Information Processing Systems (NeurIPS)*, to appear. 2023.
- X. Bi, E. Miehling, C. Beck, T. Başar. Approximate Testing in Uncertain Epidemic Processes. In *Conference on Decision and Control (CDC 2022)*, pp. 4339-4344. IEEE, 2022.
- S. Y. Olmez, S. Aggarwal, J. W. Kim, E. Miehling, T. Başar, M. West, and P. G. Mehta. Modeling Presymptomatic Spread in Epidemics via Mean-Field Games. In *American Control Conference (ACC 2022)*, pp. 3648-3655. IEEE, 2022.
- M. H. de Badyn, E. Miehling, D. Janak, Behçet Açıkmeşe, Mehran Mesbahi, T. Başar, J. Lygeros, R. S. Smith. Discrete-Time Linear-Quadratic Regulation via Optimal Transport. In *Conference on Decision and Control (CDC 2021)*, pp. 3060-3065. IEEE, 2021.
- S. Y. Olmez, J. Mori, E. Miehling, T. Başar, R. L. Smith, M. West, and P. G. Mehta. A Data-Informed Approach for Analysis, Validation, and Identification of COVID-19 Models. In *American Control Conference (ACC 2021)*, pp. 3138-3144. IEEE, 2021.
- R. Dong, E. Miehling, and C. Langbort. Protecting Consumers Against Personalized Pricing: A Stopping Time Approach. Extended abstract in *The Symposium on Foundations of Responsible Computing (FORC 2020)*.
- W. Mao, K. Zhang, E. Miehling, and T. Başar. Information State Embedding in Partially Observable Cooperative Multi-Agent Reinforcement Learning. In *59th Conference on Decision and Control (CDC 2020)*, pp. 6124-6131. IEEE, 2020.
- A. Zaman, K. Zhang, E. Miehling, and T. Başar. Reinforcement Learning in Non-Stationary Discrete-Time Linear-Quadratic Mean-Field Games. In *59th Conference on Decision and Control (CDC 2020)*, pp. 2278-2284. IEEE, 2020.
- E. Miehling, C. Langbort, and T. Başar. Secure Contingency Prediction and Response for Cyber-Physical Systems. In *4th IEEE Conference on Control Technology and Applications (CCTA 2020)*, pp. 998-1003. IEEE, 2020.
- A. Zaman, K. Zhang, E. Miehling, and T. Başar. Approximate Equilibrium Computation for Discrete-Time Linear-Quadratic Mean-Field Games. In *American Control Conference (ACC)*, pp. 333-339. IEEE, 2020.
- X. Zhang, K. Zhang, E. Miehling, and T. Başar. Non-cooperative Inverse Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 9482-9493. 2019.
- E. Miehling, R. Dong, C. Langbort, and T. Başar. Strategic Inference with a Single Private Sample, In *58th Conference on Decision and Control (CDC)*, arXiv:1909.06057. IEEE, 2019.
- K. Zhang, E. Miehling, and T. Başar. Online Planning for Decentralized Stochastic Control with Partial History Sharing. In *American Control Conference (ACC)*, pp. 3544–3550. IEEE, 2019.
- E. Miehling, B. Xiao, R. Poovendran, and T. Başar. A Bayesian Multi-armed Bandit Approach for Identifying Human Vulnerabilities. In *Decision and Game Theory for Security (GameSec)*, pp. 521–539. Springer, 2018.
- E. Miehling and D. Teneketzis. A Decentralized Mechanism for Computing Competitive Equilibria in Deregulated Electricity Markets. In *American Control Conference (ACC)*, pp. 4107–4113. IEEE, 2016.
- E. Miehling, M. Rasouli, D. Teneketzis. Optimal Defense Policies for Partially Observable Spread-ing Processes on Bayesian Attack Graphs. In *Proceedings of the Second ACM Workshop on Moving Target Defense*, pp. 67–76. ACM, 2015.

- **E. Miehling** and D. Teneketzis. Multilateral Trades in Interconnected Power Systems: A Local Public Goods Approach. In *Proceedings of the 3rd International Conference on High Confidence Networked Systems*. ACM, 2014.
- M. Rasouli, **E. Miehling**, D. Teneketzis. A Supervisory Control Approach to Dynamic Cyber-Security. In *Decision and Game Theory for Security*, pp. 99–117. Springer, 2014.

Book Chapters

- **E. Miehling**, M. Rasouli, and D. Teneketzis. Control-Theoretic Approaches to Cyber-Security. In *Adversarial and Uncertain Reasoning for Adaptive Cyber Defense* (Sushil Jajodia, George Cybenko, Peng Liu, Cliff Wang, Michael Wellman; Editors), pp. 12–28. Springer/Birkhäuser, 2019.
 - M. Rasouli, **E. Miehling**, D. Teneketzis. A Scalable Decomposition Method for the Dynamic Defense of Cyber Networks. In *Game Theory for Security Risk Management – From Theory to Practice of Static & Dynamic Game Theory: Foundations and Applications* (Tamer Başar, Series Editor), pp. 75–98. Springer/Birkhäuser, 2017.
-