Sequential decision-making in complex partially observable environments CRCS Workshop on AI-Assisted Decision-Making for

Conservation

Andrew Perrault Asst. Prof., Computer Science and Engineering The Ohio State University (formerly postdoc, Harvard CRCS)



The dream of AI methods for PA management









/ILDME

. . .



Figure from Silvestro et al., Improving biodiversity protection through artificial intelligence, Nature Sustainability, 2022.





Accelerating fusion science through learned plasma control

February 16, 2022



Mastering the Game of No-Press Diplomacy via Human-Regularized Reinforcement Learning and Planning

Anton Bakhtin, David J Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H Miller, Noam Brown



Degrave et al., Magnetic control of tokamak plasmas through deep reinforcement learning, Nature, 2022. Paquette et al., No Press Diplomacy: Modeling Multi-Agent Gameplay, NeurIPS 2019.







Reinforcement learning without a trusted simulator



Other domains: medication adherence (NeurIPS 2020, AAMAS 2021a/b) Patient procedure prep Restaurant inspections (ICML Workshop 2022)

> Wang, S., C., P., D.-V., Tambe. Learning MDPs from Features: Predict-Then-Optimize for Sequential Decision Problems by Reinforcement Learning. NeurIPS 2021.

Reinforcement learning without a trusted simulator





P., W., E., D., M., Tambe. AAAI 2020. Wang, S., C., P., D.-V., Tambe. NeurIPS 2021. Xu, **P.**, C., F., Tambe. UAI 2021. Byun, **P.** ICLR 2022. Work on noisy rewards (ongoing)



How to estimate deterrence from ranger patrols?



Xu, P., P., D., W., R., Tambe. arXiv, 2020.



How to estimate deterrence from ranger patrols?



Xu, P., P., D., W., R., Tambe. arXiv, 2020.







Guo, Xu, P., Davies, Tambe, et al. (in prep)







Historical patrol effort Predicted historical patrol effort Predicted study patrol effort Historical poaching risk Distance to boundary Distance to patrol posts





Guo, Xu, P., Miratrix, Davies, Tambe, et al. (in prep)



Questions

- 1. Other data with a shock/disruption?
- 2. PAs with aligned patrolling and camera trap data?
- 3. Better ways to infer poaching levels from patrol outcomes