




Reinforcement Learning Informs Harvest Control Rule
Design in Complex Fishery Settings

Project Lead: Christopher Cahill
Contact info: cahill11@msu.edu
QFC Collaborators: Christopher Cahill
Collaborators: Felipe Montealegre-Mora (UC-Berkeley), Carl Boettiger (UC-Berkeley), Carl Walters (University of British Columbia)
Funding Agency: National Science Foundation
Active Dates: 2022 – 2025



Caption: In reinforcement learning, an agent acts in an environment and receives rewards and state feedback in response. Photo credit: Wikipedia

Goal: Many feedback policies are rooted in analytical or dynamic programming solutions developed in the 1970s–1980s, which assume simple ecological dynamics and are subsequently applied to far more complex ecological and management contexts through approaches such as Management Strategy Evaluation. In this paper, we address the problem of harvest control rule design for partially observed, age-structured populations with highly variable and spasmodic recruitment using tools from reinforcement learning (RL).

Objective: The objective of this work was to evaluate the performance of feedback policies in the presence of spasmodic and highly variable recruitment dynamics.

Management Implications: This research improved understanding of the harvesting theory underlying feedback policies applied in management systems throughout the world. We showed that for simpler management objectives like yield or consistency in catch, RL was unable to outperform classic solutions. However, when management objectives were more complex, policies found using RL outperformed standard or traditional feedback policies used in many fisheries management systems.

Methods:

- We simulated age-structured population dynamics based on Walleye with highly variable recruitment
- We used a combination of simulation, Bayesian optimization, and proximal policy optimization via deep neural networks to solve the feedback policy problem for populations with highly variable and spasmodic recruitment dynamics
- Methods were computationally demanding, coded in Python, and training times for deep neural networks took approximately 2h on a commercial Graphics Processing Unit

Policy	Yield utility	HARA utility	Trophy fishing utility
oPP	252.82 ± 25.31	401.73 ± 20.93	88.51 ± 7.17
cPP	228.64 ± 25.03	364.55 ± 25.94	91.34 ± 8.76
U _{MSY}	237.28 ± 22.68	400.58 ± 20.67	96.44 ± 8.24
1RL	250.25 ± 22.82	402.43 ± 21.14	92.73 ± 7.13
2RL	249.47 ± 23.86	391.27 ± 21.51	126.90 ± 12.80

Caption: Summary statistics of the reward distributions for optimised policies.

Key Findings:

- We found that mean weight was useful in the trophy fishing setting scenario, and surprisingly, not when managers were interested in maximizing yield or risk-averse utility

Deliverables: Montealegre-Mora, F., Boettiger, C., Walters, C., and C.L. Cahill. 2025. Machine Learning Informs Harvest Control Rule Design in Complex Fishery Settings. Fish and Fisheries 26:1004–1020.

QFC Supporting Partners

