

Genetic and Evolutionary Computation Conference



San José, Costa Rica
July 13-17, 2026

Leveraging Evolutionary Surrogate-Assisted Prescription in Chlorination Control Systems

Rivaaj Monsia
University of Texas at Austin
rivaaj@utexas.edu

Daniel Young
University of Texas at Austin
danyoung@utexas.edu

Olivier Francon
Cognizant AI Labs
olivier.francon@cognizant.com

Risto Miikkulainen
University of Texas at Austin,
Cognizant AI Labs
risto@cs.utexas.edu



Chlorine in Water Networks

- Water Distribution Systems (WDS) integral to **urban infrastructure**
- Safe water requires stable chlorine
  - Levels must stay within tight bounds
- Current Cl. control systems based on sensors: **costly, need extensive planning**
- Goal: **enable efficient, adaptive control to support urban growth via evolutionary ML**

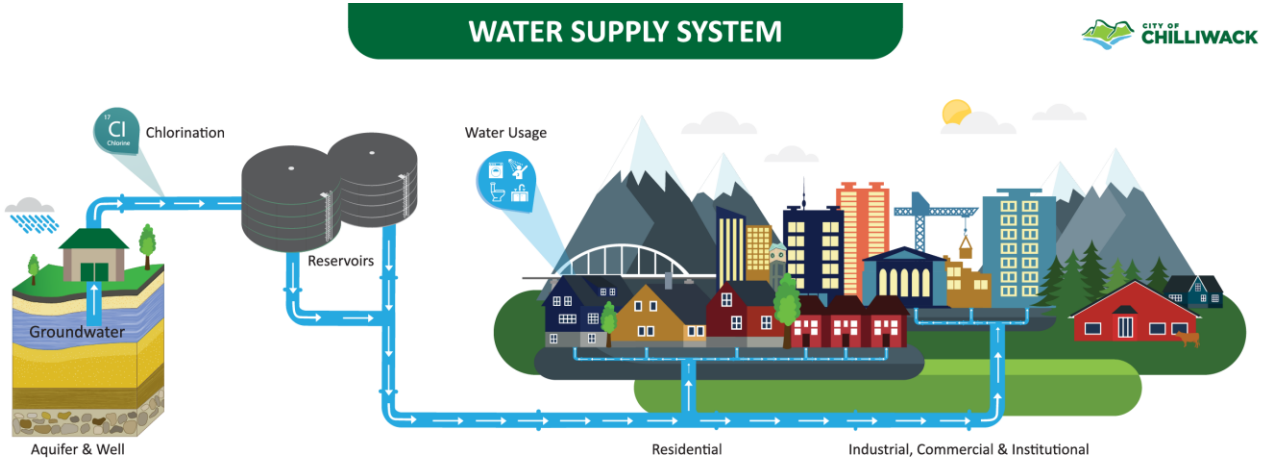

Motivation

- Infeasible to develop new systems on real-world WDS
  - Led to development of **computational simulators: EPANET**
- Need **fast predictive models**
  - EPANET is accurate but slow
  - Not suitable for RL or online control
- **Surrogate ≈ lightweight simulator**

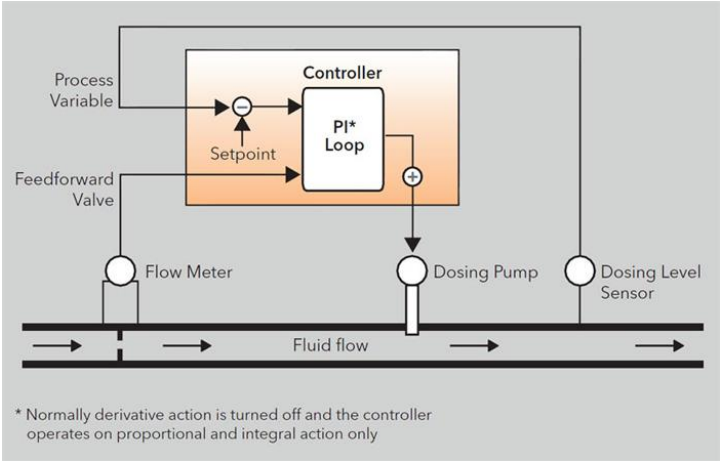

real-world control loop

- Approach: utilize **surrogate to enable agent learning via neuroevolution**

Background

- **Water networks: graph**
  - nodes, pipes, tanks
- **Dynamics:** transport + decay
  - Physical constraints
- **Actions:** injection rates
- **Observations: chlorine + flows at sensors**

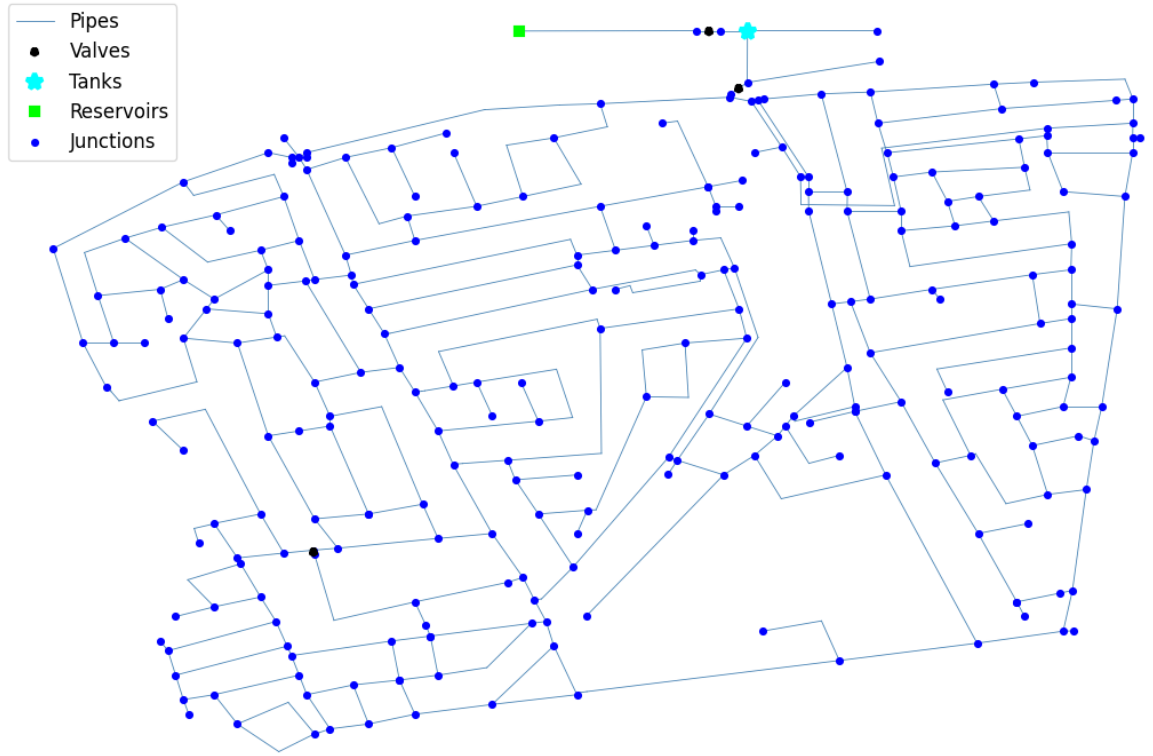

Topology of WDS network used in this research

Key Challenges

- Nonlinear transport + decay dynamics
- Must generalize over contamination events
- **Multi-objective** optimization

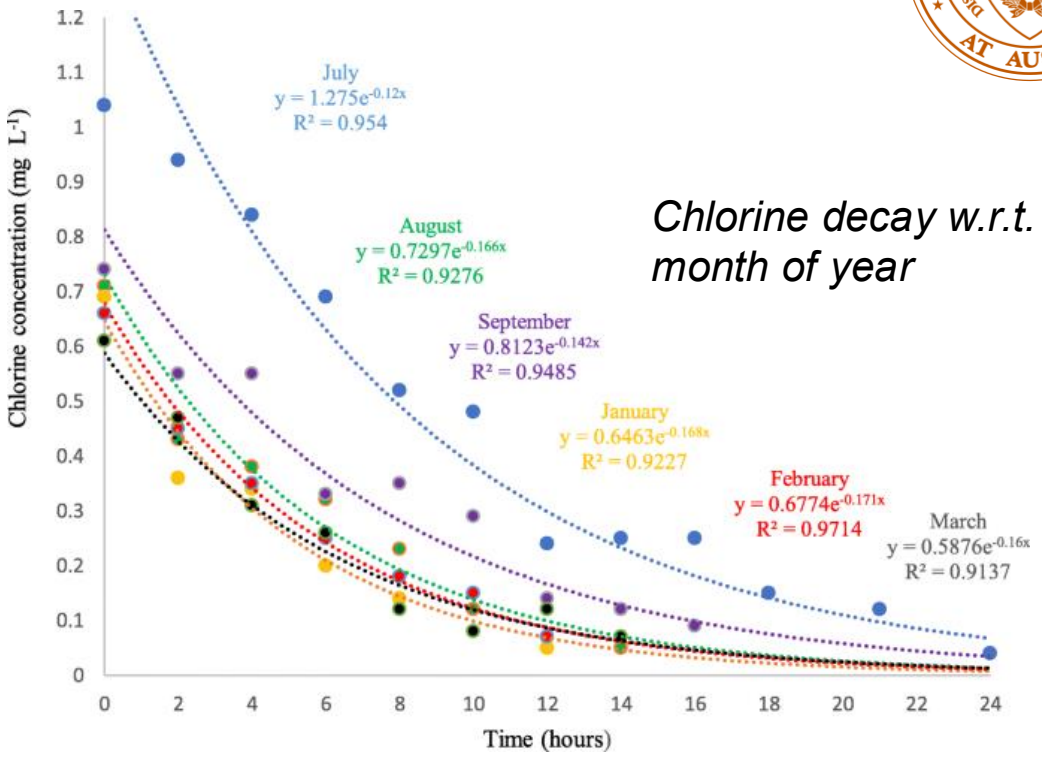

Chlorine decay w.r.t. month of year

- Prompted competitions like **1st AI for Drinking Water Chlorination Challenge @ IJCAI 2025**
  - 2nd place with rudimentary version of proposed approach

Approach

ESP

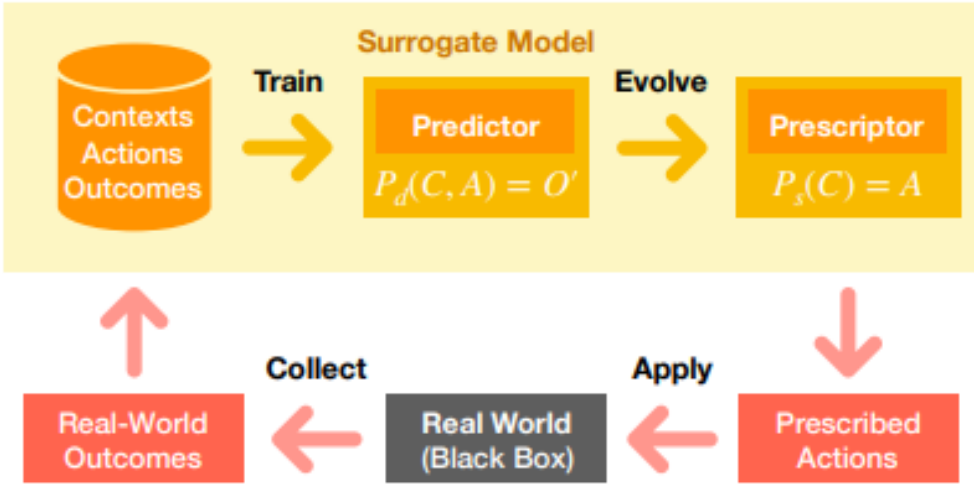

NEAT

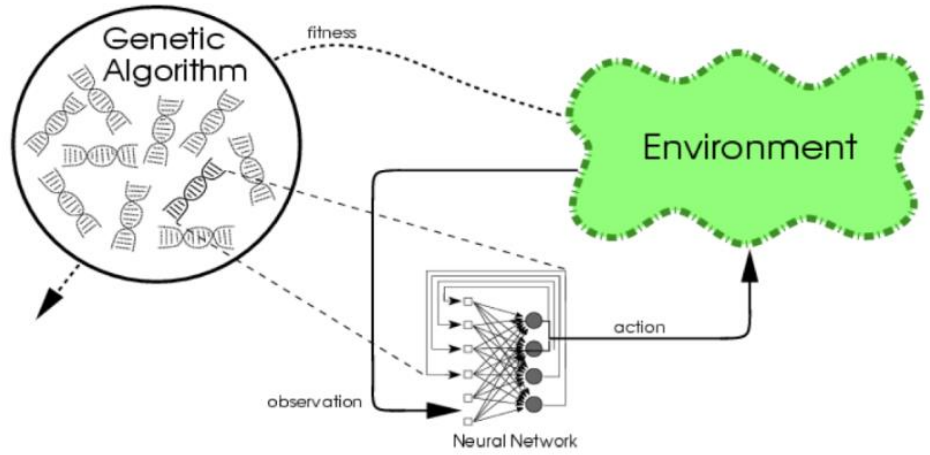

- Evolutionary Surrogate-Assisted Prescription (**ESP**)
  - Student-Teacher **Distillation**
  - Applications: land-use, COVID-19, diabetes
- Agent evolution: **NEAT**
- **NSGA-II** multi-objective optimization (*Pareto-dominance*)
  - **Curriculum Learning**

Surrogate Model

- Context C_t = *{x[t−H+1], …, x[t]}*
- *x[t] = [flow, conc., action]*
- *y[t] = [flow, conc.]*

- Predicts state increment (**Δ**)
- Residual: *y[t+1] = y[t] + **Δ***

- Time-aware + causal? **LSTM**

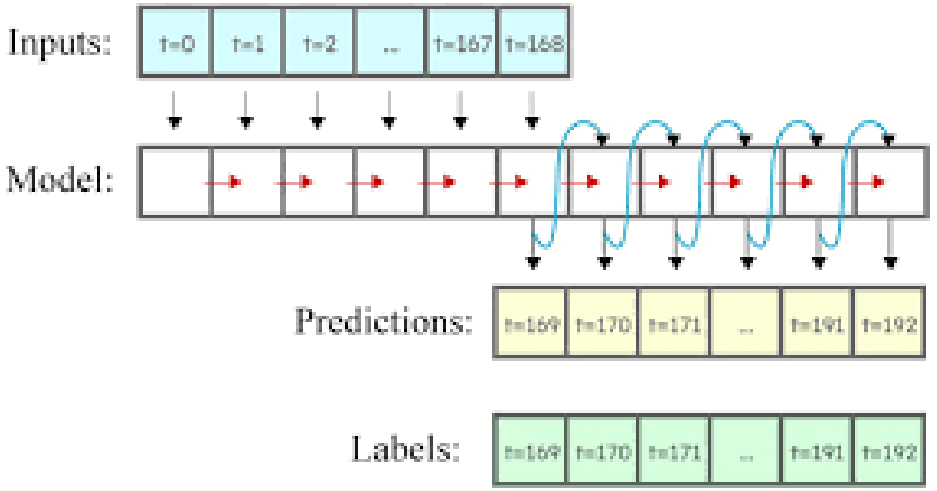

Autoregressive nature of LSTM

Surrogate Distillation

- Need *uni*-LSTM to **understand future dependencies**
  - Use BiLSTM with Context $C_t' = \{x[t-H+1], \ldots, x[t-1], x[t+1], \ldots, x[t+L]\}$
- BiLSTM **non-causal** and infeasible for online, autoregressive predictions
- **Distillation!**
  - Teacher: **BiLSTM**
  - Student (**surrogate**): smaller **uni-LSTM**
  - Train w/ *hard* + *soft* terms

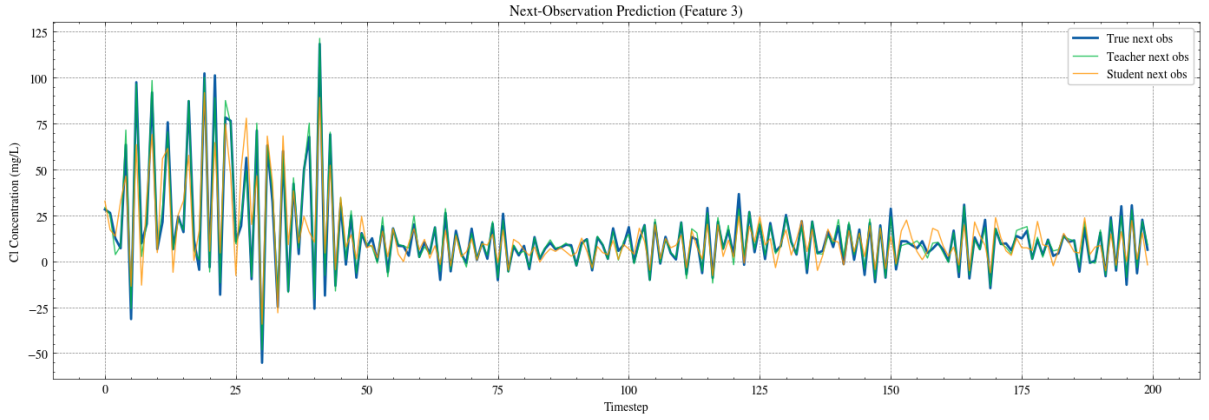

$$\mathcal{L}(\theta) = \lambda_{\text{hard}} \underbrace{\|\hat{\Delta}_t^{(S)} - \Delta_t\|_2^2}_{\text{ground-truth loss}} + \lambda_{\text{soft}} \underbrace{\|\hat{\Delta}_t^{(S)} - \hat{\Delta}_t^{(T)}\|_2^2}_{\text{teacher-loss}} + \lambda_{\text{feat}}\mathcal{L}_{feat} + \lambda_{\text{rollout}}\mathcal{L}_{\text{rollout}}$$

ESP Setup

(1) Train surrogate via distillation with simulator data pairs

- Actions ~ U(0, 1000)

(2) NSGA-II + NEAT for 50 generations w/ **possible new objective**

(3) Evaluate 5 best genomes on Pareto front on 5 scenarios

(4) Use data to **fine-tune surrogate** (*T: 20 epochs → S: 10 epochs*)

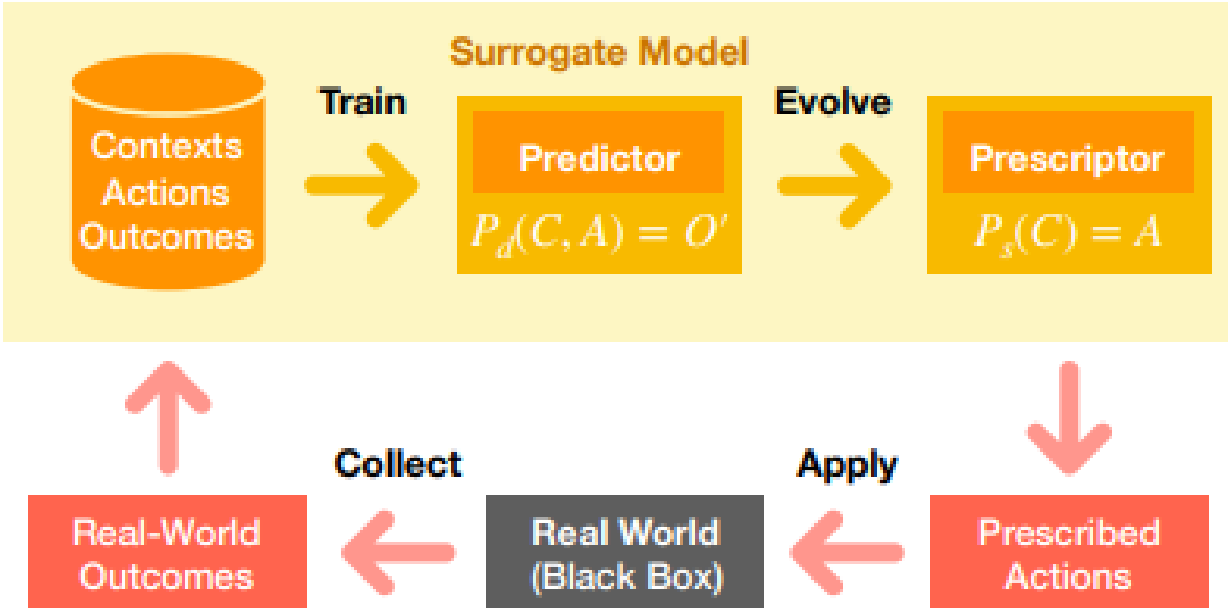

Fundamental ESP approach

Agent Baselines

Table 2: Performance of various controllers across five evaluation objectives on held-out scenario. Lower is better for all objectives.

Agent	Bound	Fairness	Smoothness	Cost	Infection Risk
Constant Injection					
0.3 mg/L	0.114	**0.194**	—	**1296.0**	5.160
10 mg/L	**0.105**	0.196	—	4.32×10^4	5.160
250 mg/L	0.840	8.915	—	1.08×10^6	3.612
Random Injection					
$a_t \sim U(0, 250)$	0.408	4.289	87.024	5.334×10^5	4.431
$a_t \sim U(0, 10000)$	16.673	185.722	3478.980	2.165×10^7	**0.106**
Evolved Agents					
NEAT	0.113	0.195	**0.452**	3456.000	5.160
NSGA-II	0.401	24.014	17.693	7.500×10^5	5.160
NSGA-II (*Curricular*)	0.463	5.180	6.852	4.843×10^5	3.627
PPO Agent					
PPO	0.114	0.195	3.658	3636.568	5.160

- **PPO** *(2.71 M steps)* **+ NEAT** *(50 generations, 350 p.s.)* collapse to null injections
  - Trained with composite reward
- Curricular obtains **2-5x better tradeoffs** in most objectives than non-curricular

Pareto Evolution

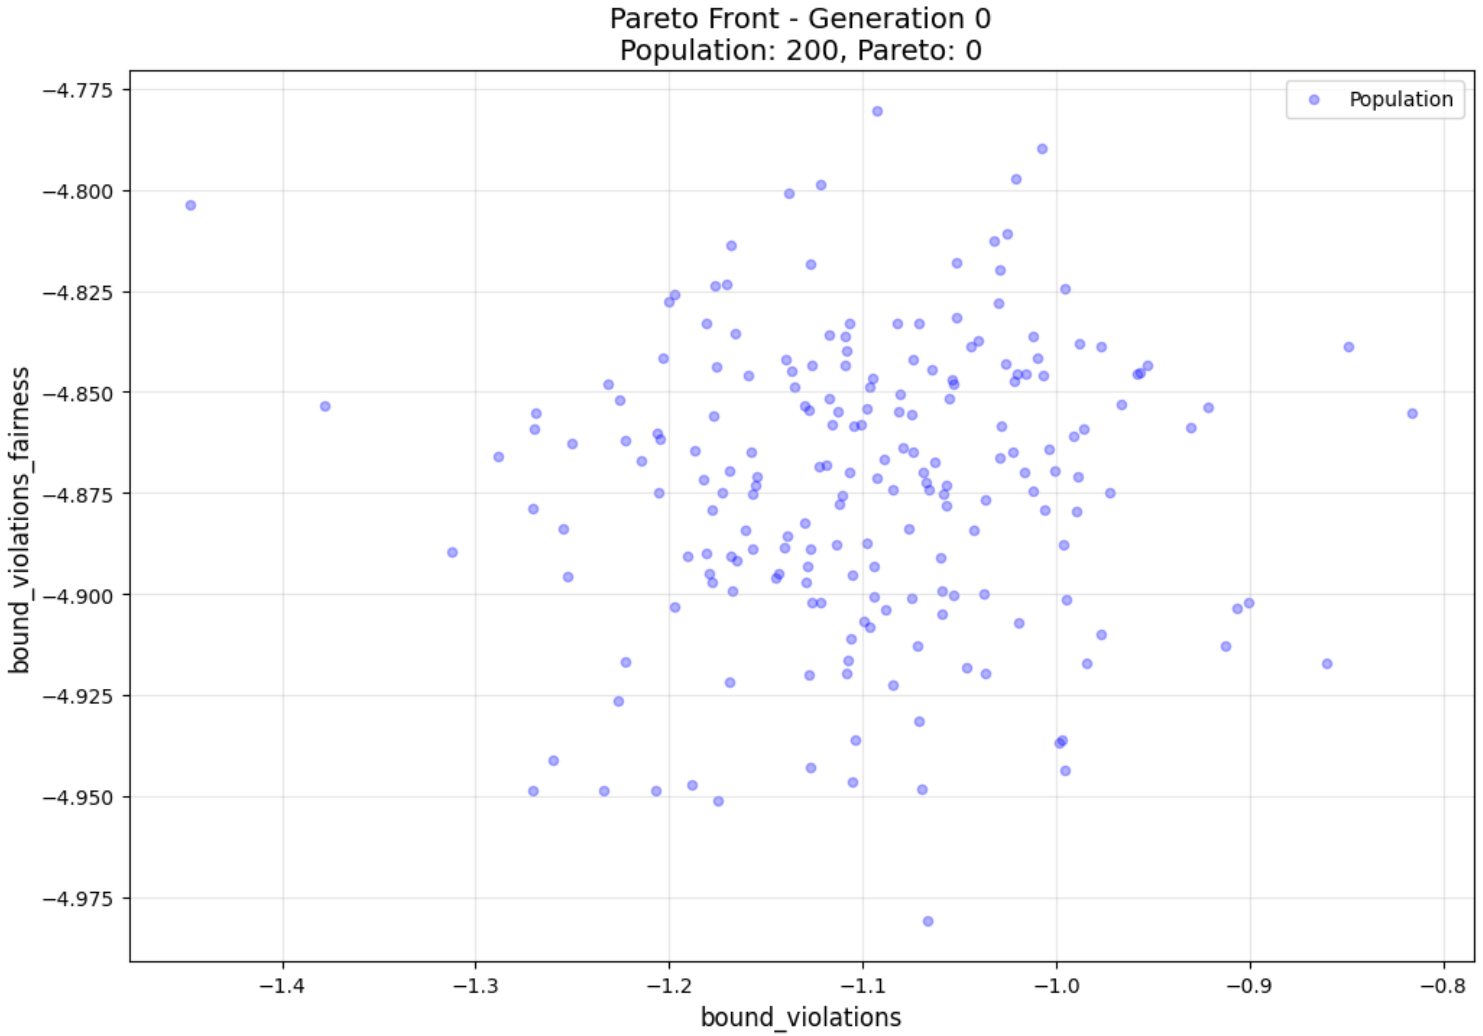

ESP iteration 0

Initially **no pareto front**

Pareto Evolution (cont.)

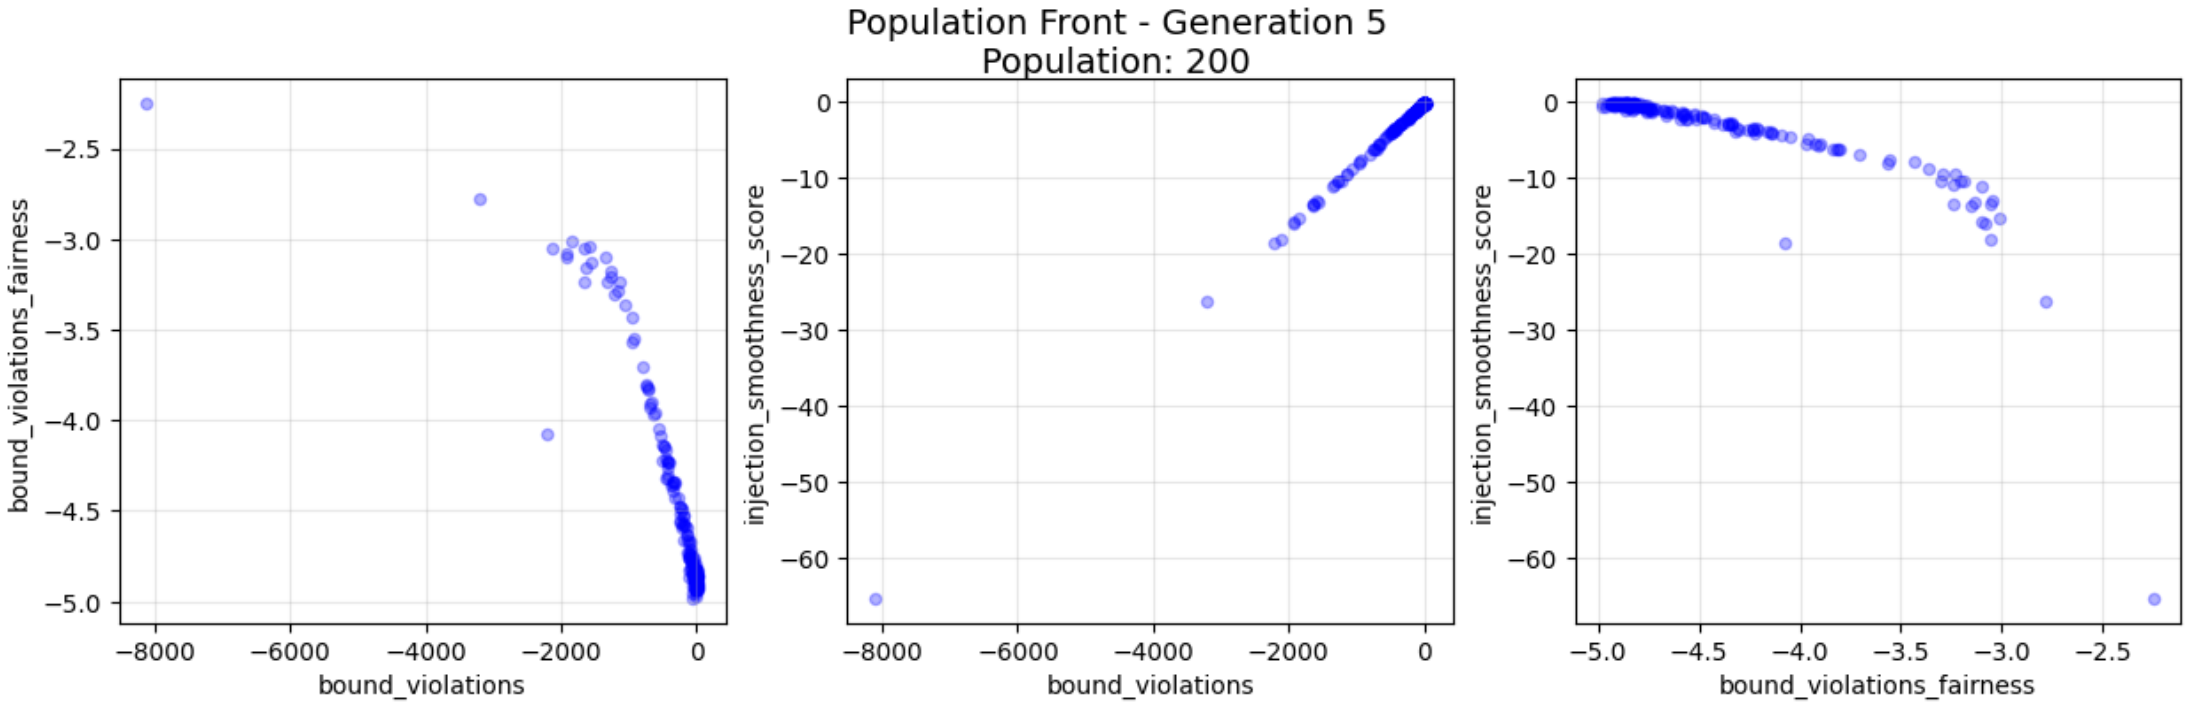

ESP iteration 3, after adding new objective

Pareto fronts are visible, but **do not encompass the entire tradeoff region**

Pareto Evolution (cont.)

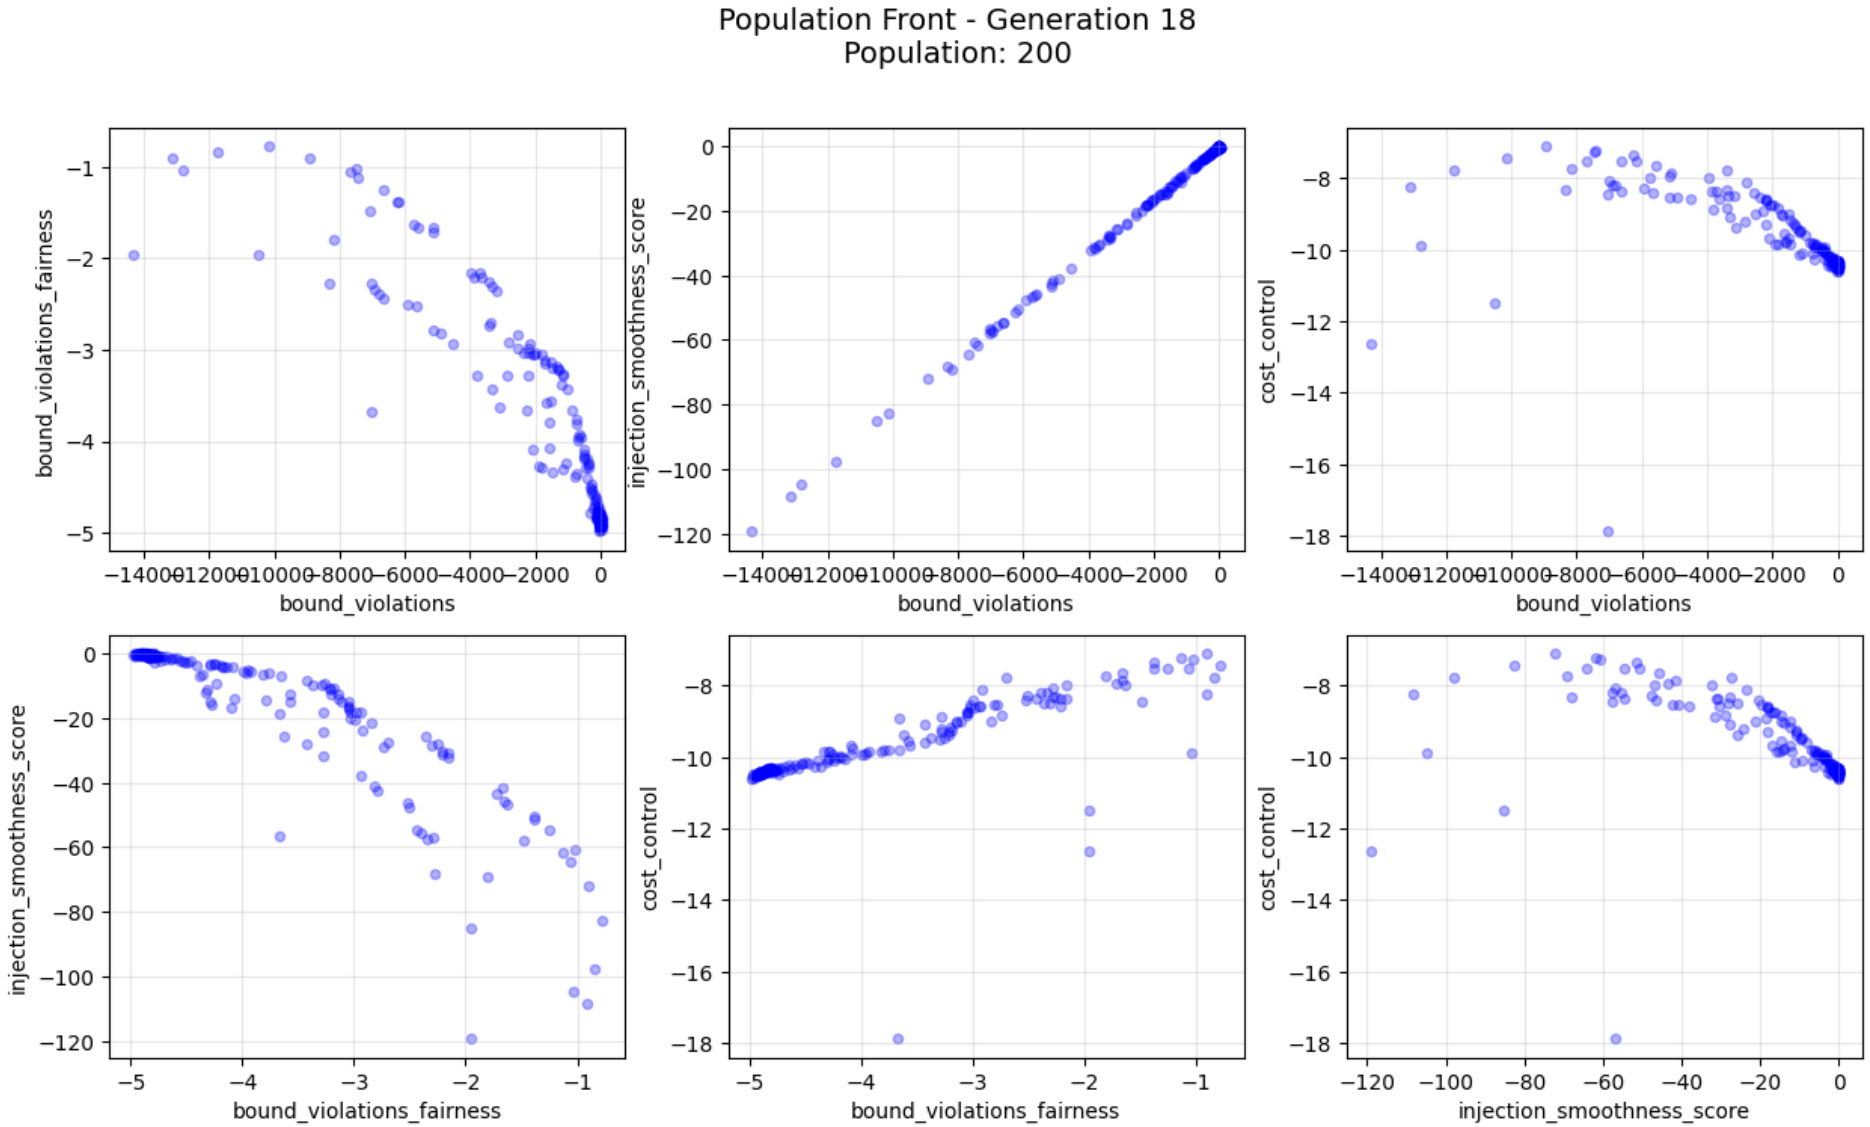

ESP iteration 10

Pareto front shows **entire tradeoff region populated**

Curriculum Ablations

- Does **changing order in which objective introduced change** final agent performance? **YES**

Ordering	Bound	Fairness	Smoothness	Cost	Inf. Risk
(BC)SF	**0.154**	**1.195**	**5.277**	$\mathbf{1.224 \times 10^5}$	5.14
(BC)FS	0.246	3.267	6.598	2.783×10^5	5.079
(FC)BS	0.615	3.864	11.045	7.215×10^5	**3.794**
(FC)SB	0.422	5.432	7.978	3.792×10^5	3.855

B=*bounds*, C=*cost*, F=*fairness*, S=*smoothness*

Table 3: Objective evaluations on different curriculum orderings of Pareto-optimal genomes. Parentheses denote the initial two objectives and subsequent objectives are in order of introduction. Lower is better for all.

- **Primary two objectives matter most**
  - *Adversarial objectives hinder evolution*
- Needs further population-wise investigation over all 12 orderings

Semantic Results

Surrogate fine-tuning enables **emergence of new NEAT species** over unexplored objective landscape

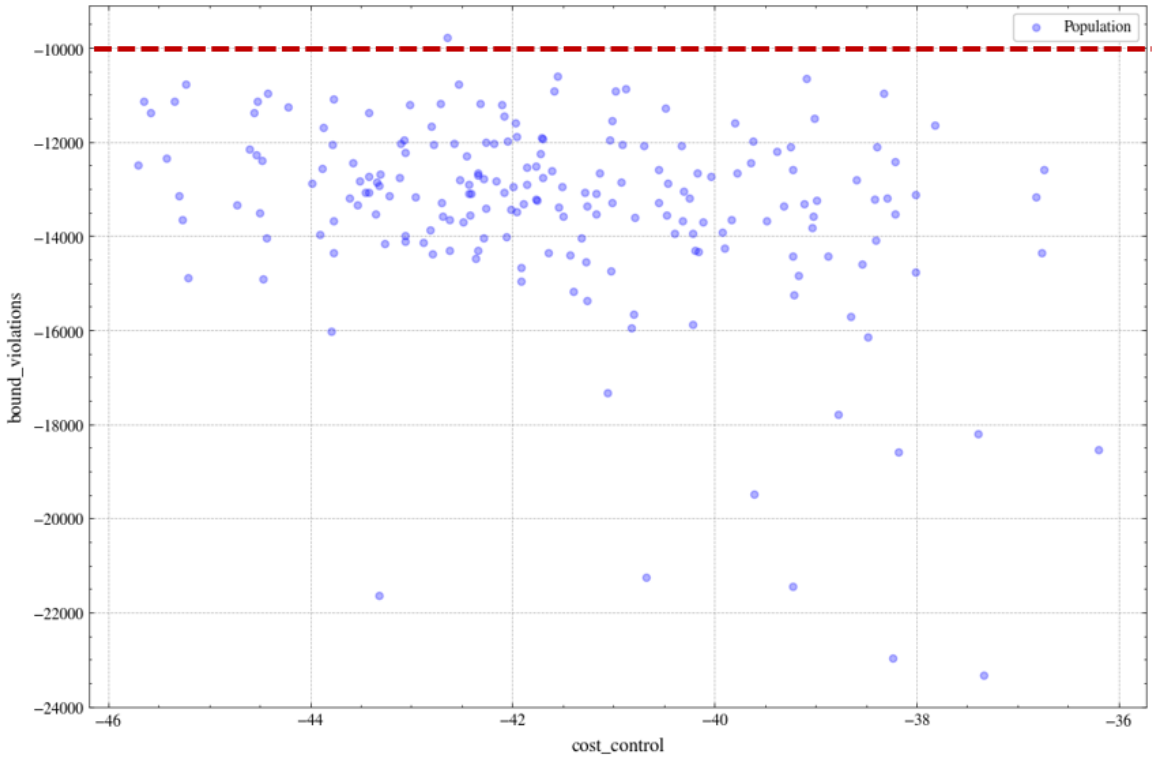

(a)

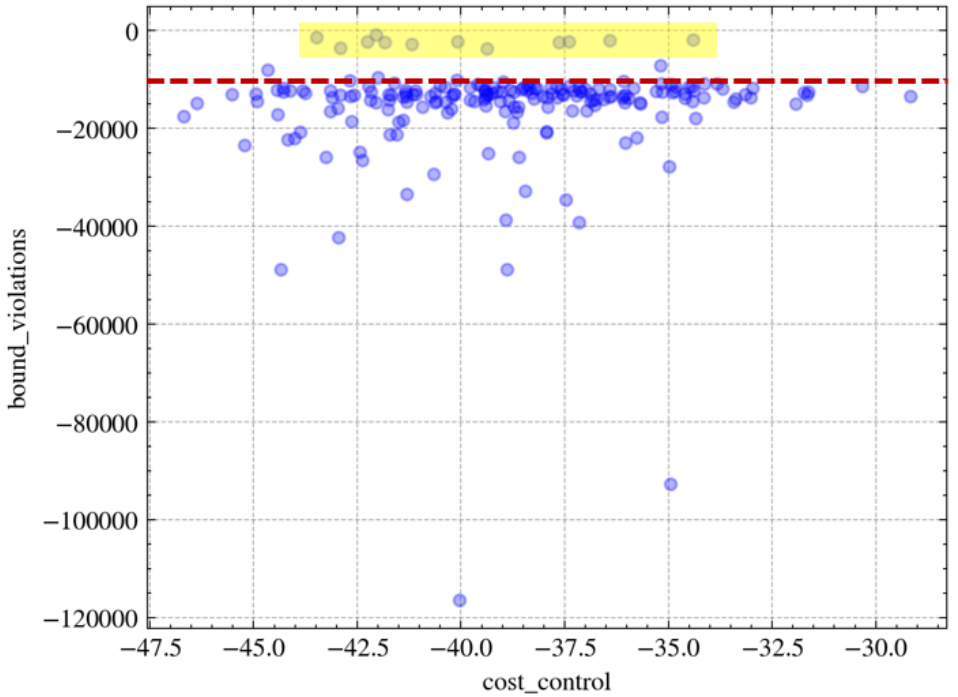

(b)

Semantic Results (cont.)

Surrogate becomes pseudo-***"rolling-mean estimator"*** due to **synergistic co-learning** of predictor and prescriptor population

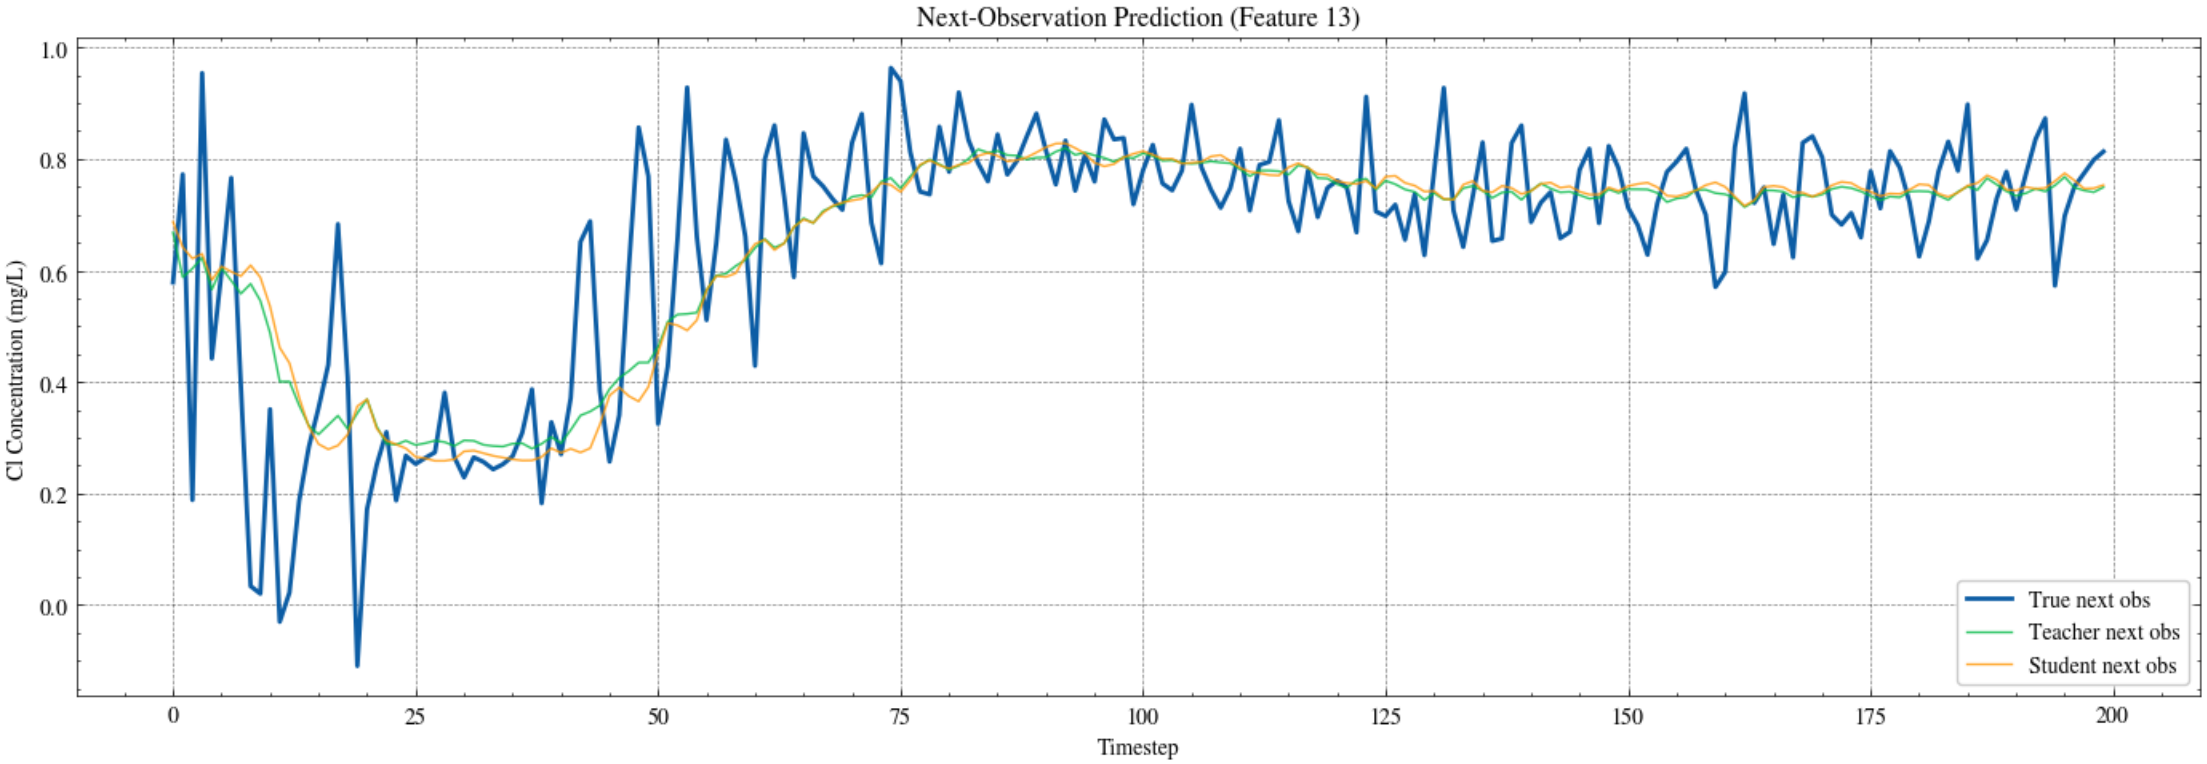

Future Directions

- Implement **physical constraints**
- Use specialized architectures that **capture spatial heuristics**
- **Longer scenarios**

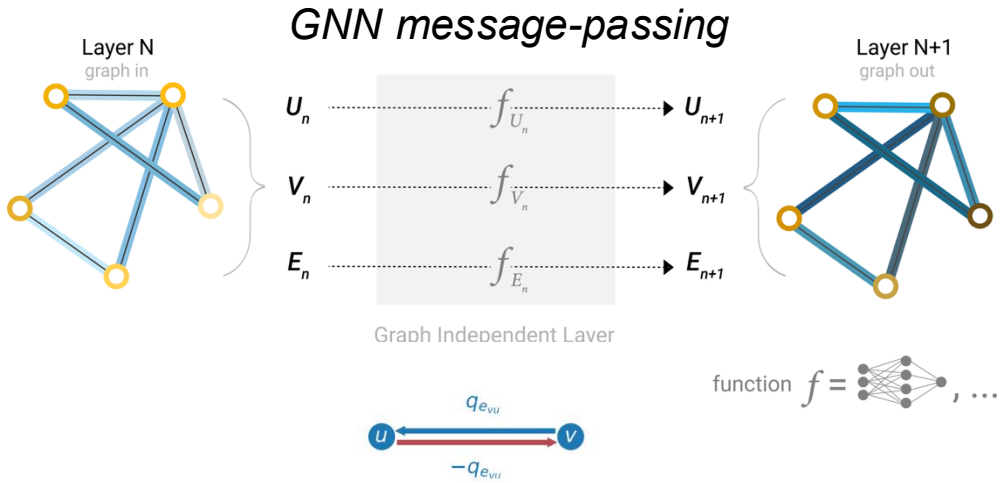

Figure 2: The flow $q_{e_{uv}}$ from node u to v has to be equal to the negative of the flow $q_{e_{vu}}$ from v to u.

Principles of fluid dynamics

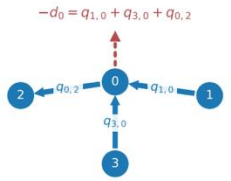

Figure 3: The sum of inflows $q_{1,0}, q_{3,0}$ and outflow $q_{0,2}$ must be equal to the negative of the demand d_0 at node 0.

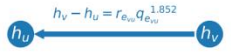

Figure 4: The relationship between pressure heads at neighbouring nodes and the flow in connecting pipe.

- Participate in *2nd AI for Drinking Water Chlorination Challenge* (and win!)

Conclusions

- Curricular, multi-objective optimization **learn agents over entire objective space**
- **ESP** can discover **unique, real-world strategies for water management**
- **Hydraulic simulators can be used to learn agents** that matter
  - may require a surrogate for learning (meta-surrogate)

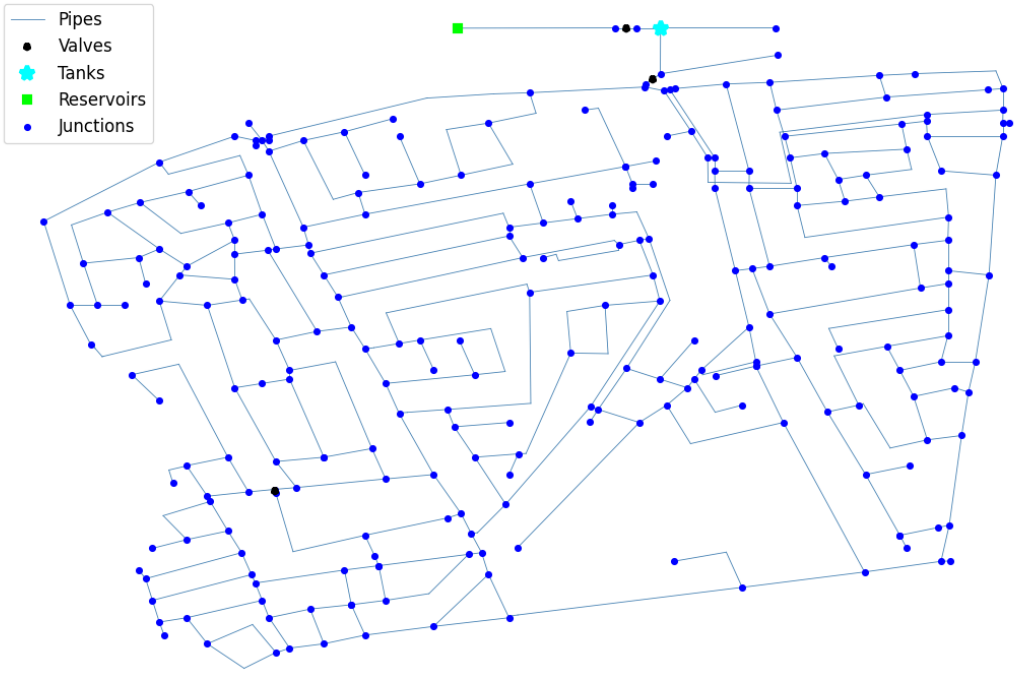



Thank You!

Questions?