Machine learning based forecasting and decision

Fearghal O’Donncha, Muneeza Asmat, Malvern Madondo, Peimeng Guan, Arun Bawa, Raya Horesh, Michael Jacobs, Raghavan Srinivasan

Water Availability

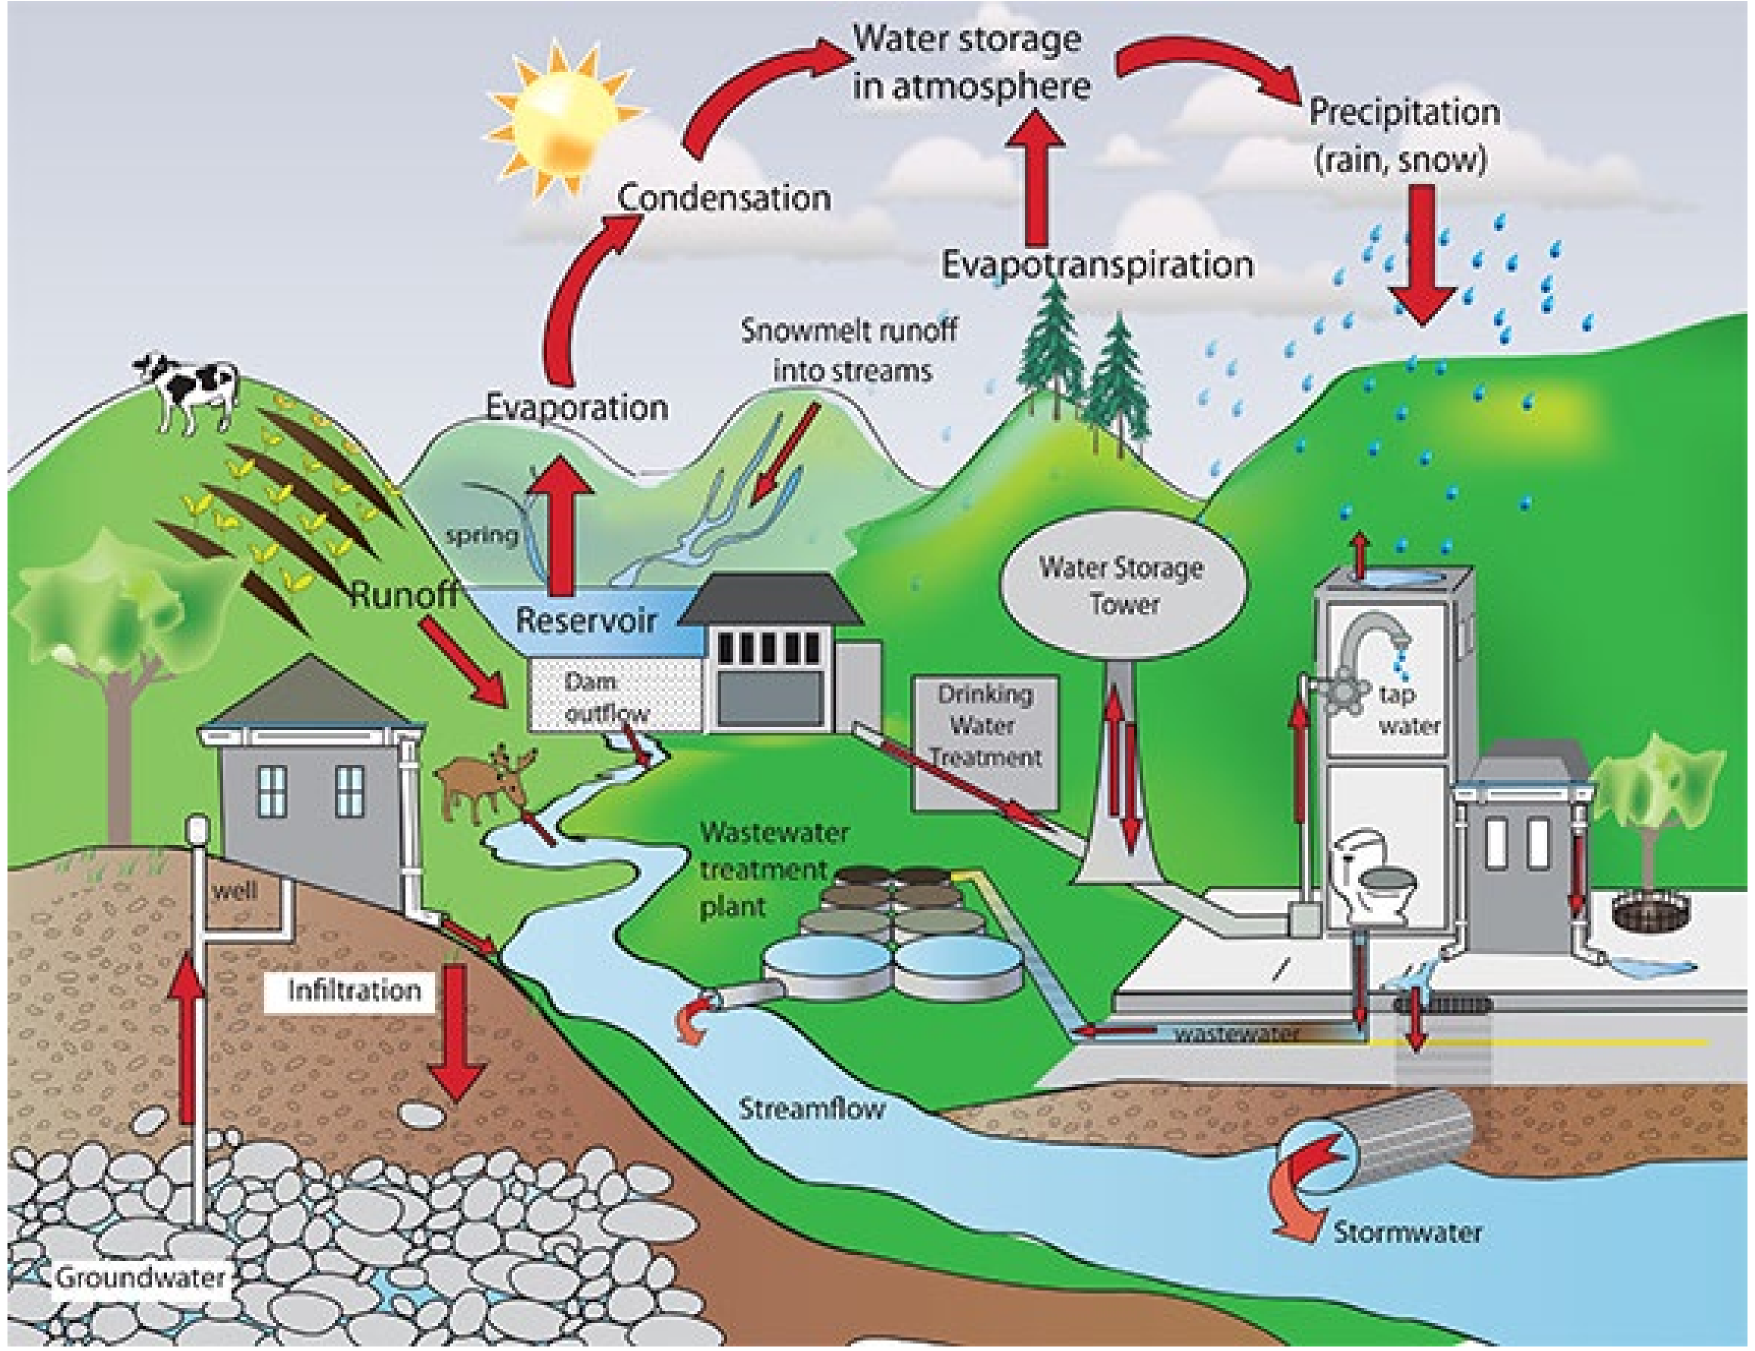

Source: Cary Institute

47% watersheds face water shortages by 2071

37% of total water usage from agriculture

87% decline in freshwater species since 1970

25% increase in global population by 2050

Existing Methods for Soil Moisture (SM) Forecast

Input

HRUs

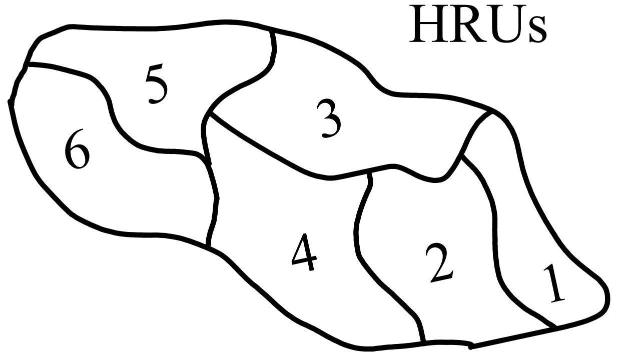

- Hydrological properties
- Climatic properties
- Soil properties
- Land Management

Hydrological Response Units (HRUs) delineation of spatial extents

Grouping

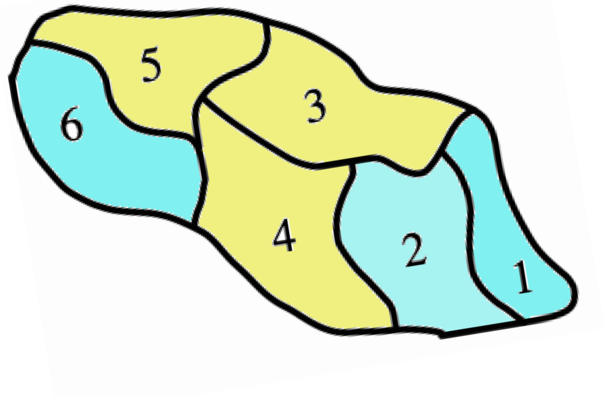

Grouping of HRUs that share same soil, land use, and slope characteristics

SWAT

1

Physics based hydrological model termed the Soil & Water Assessment Tool

Output

Temporal hydrological response features

Computationally Expensive

Oversight from domain experts

[1] J. G. Arnold, J. Kiniry, R. Srinivasan, et.al. *"Soil and Water Assessment Tool input/output file documentation version 2012. Technical report"*, Texas Water Resources Institute, 2012

Our Domain-Inspired Approach for SM Forecast

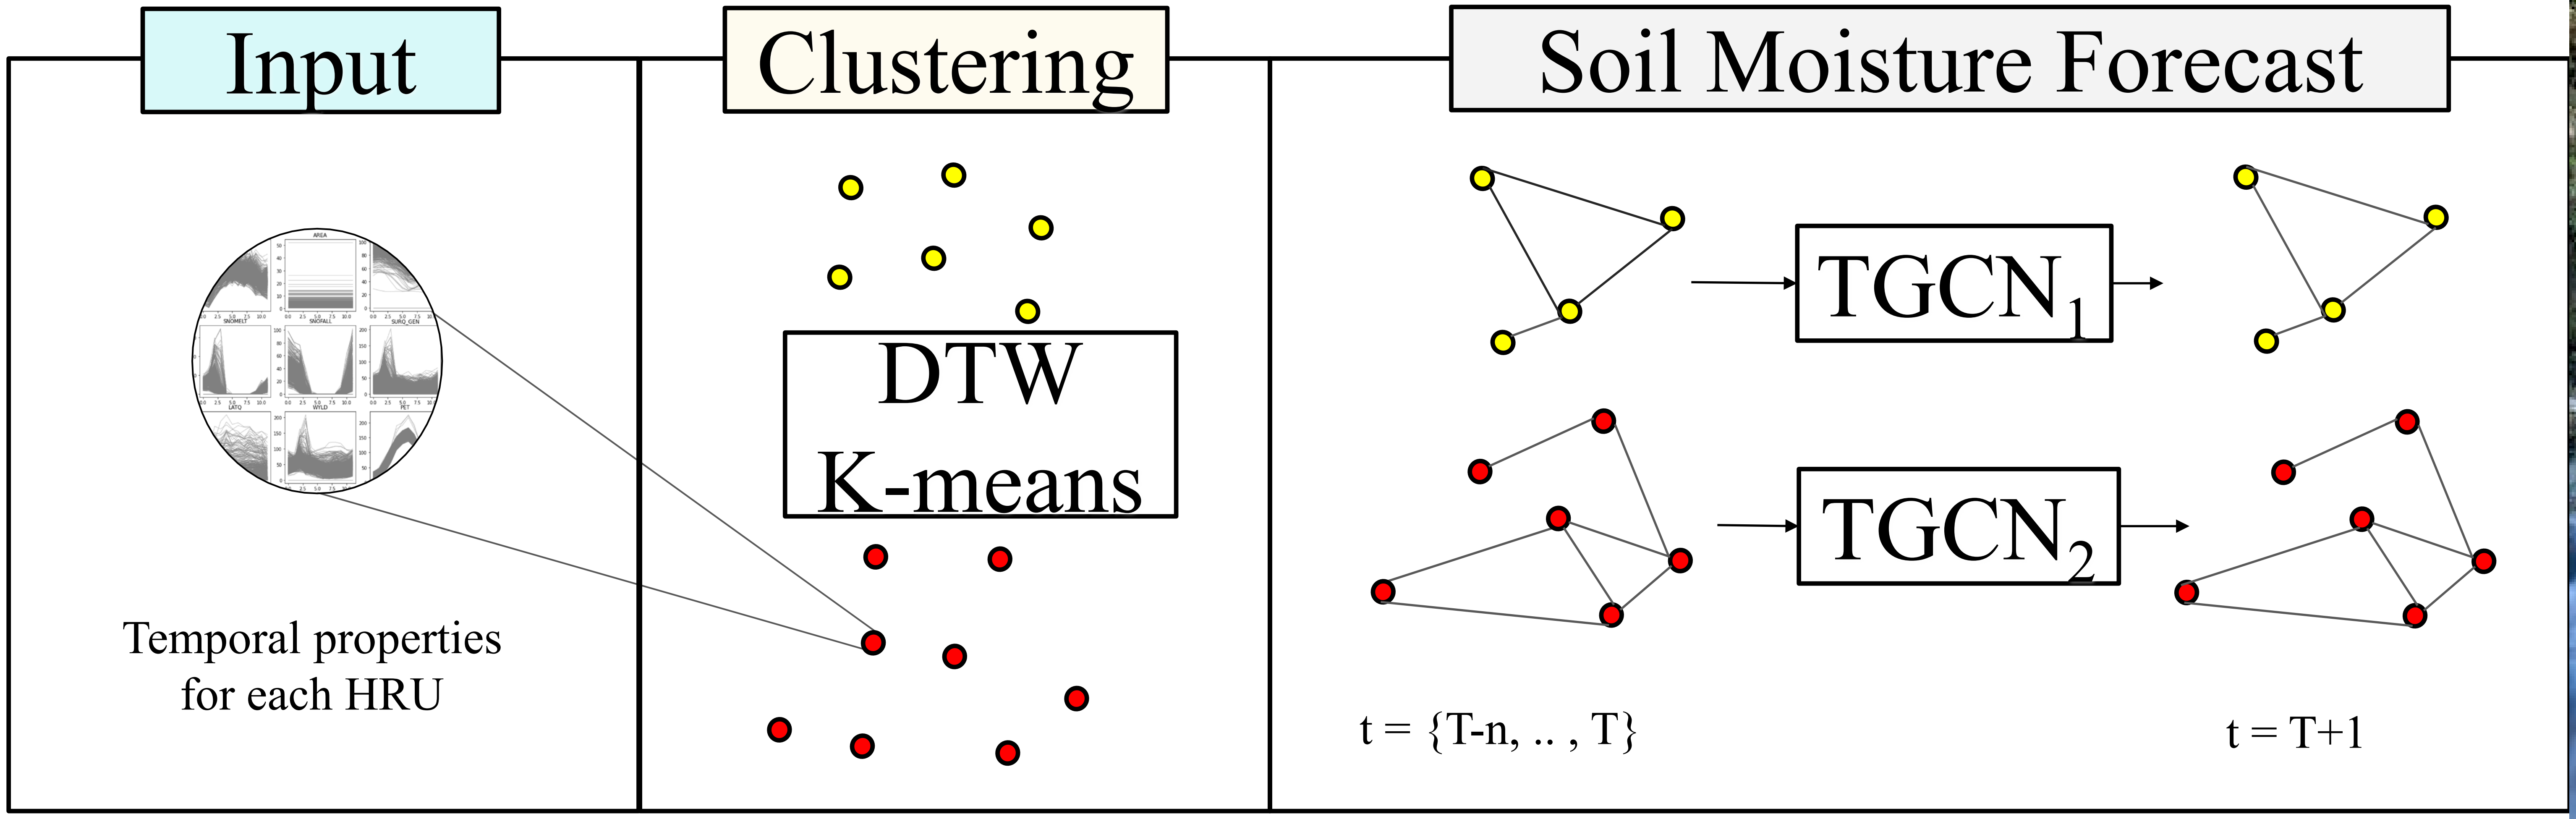

~99k HRUs
81 Temporal Features
Monthly data for 38 years

Dynamic Time Warping (DTW)

Temporal Graph Convolution Neural Network (TGCN)

Dynamic Time Warping

A phase invariant metric for measuring similarity between two discrete time series. [1]

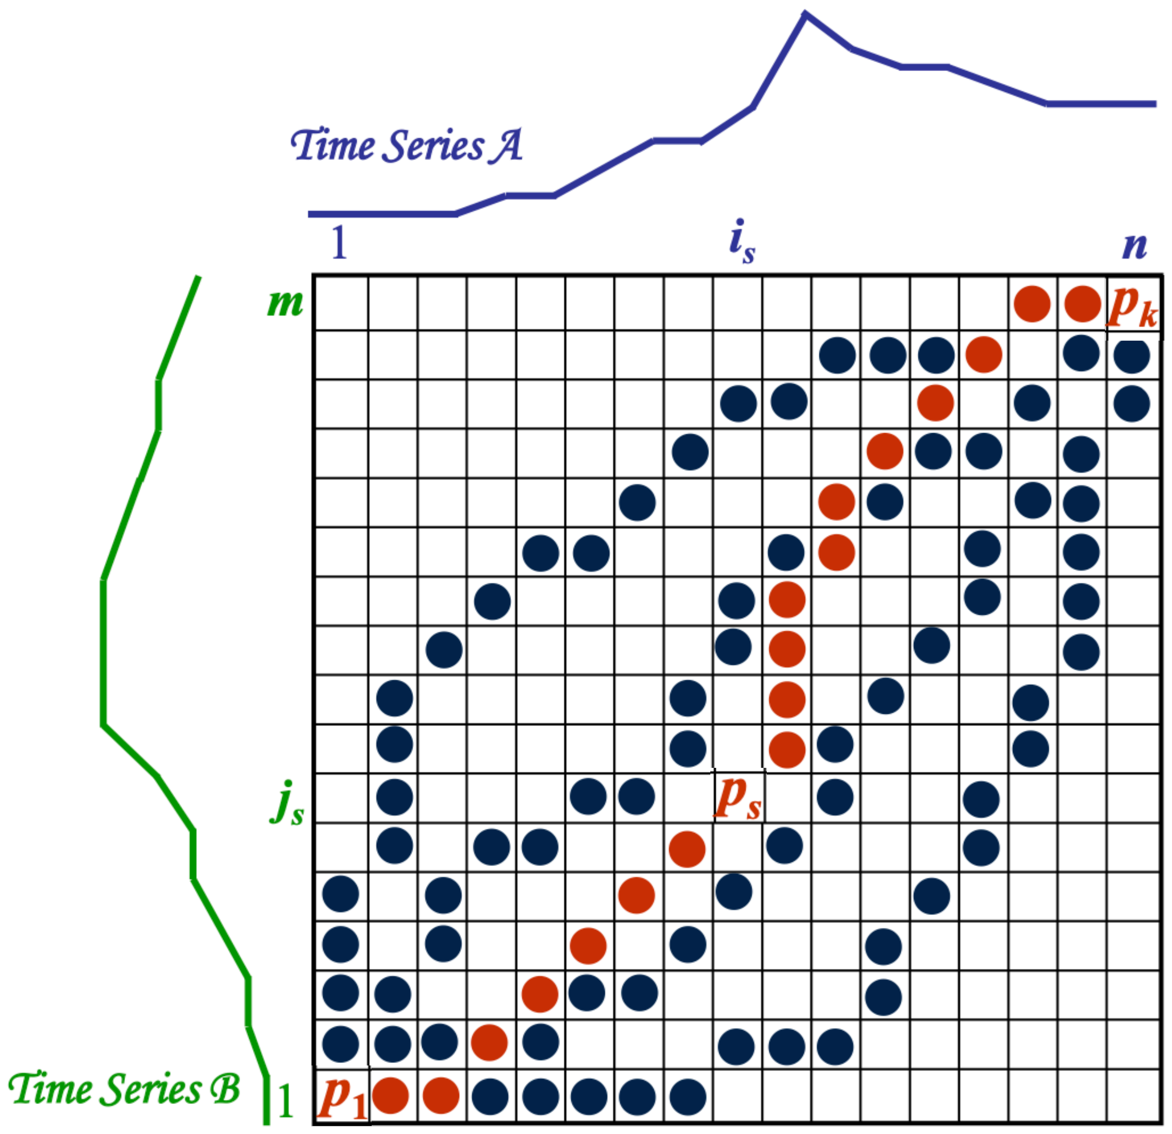

$P = (p_1, \dots, p_s, \dots, p_k)$

To find the optimal match between time series A and B:

1. Populate the distance matrix.
2. Find ***P*** (warping function) such that it minimizes the total distance between A and B subject to regularity conditions:
   1. Monotonicity.
   2. Continuity.
   3. Boundary conditions.
   4. Slope Constraint.

[1] https://www.math.emory.edu//~lxiong/cs730_s13/share/slides/searching_sigkdd2012_DTW.pdf

Clustering Results

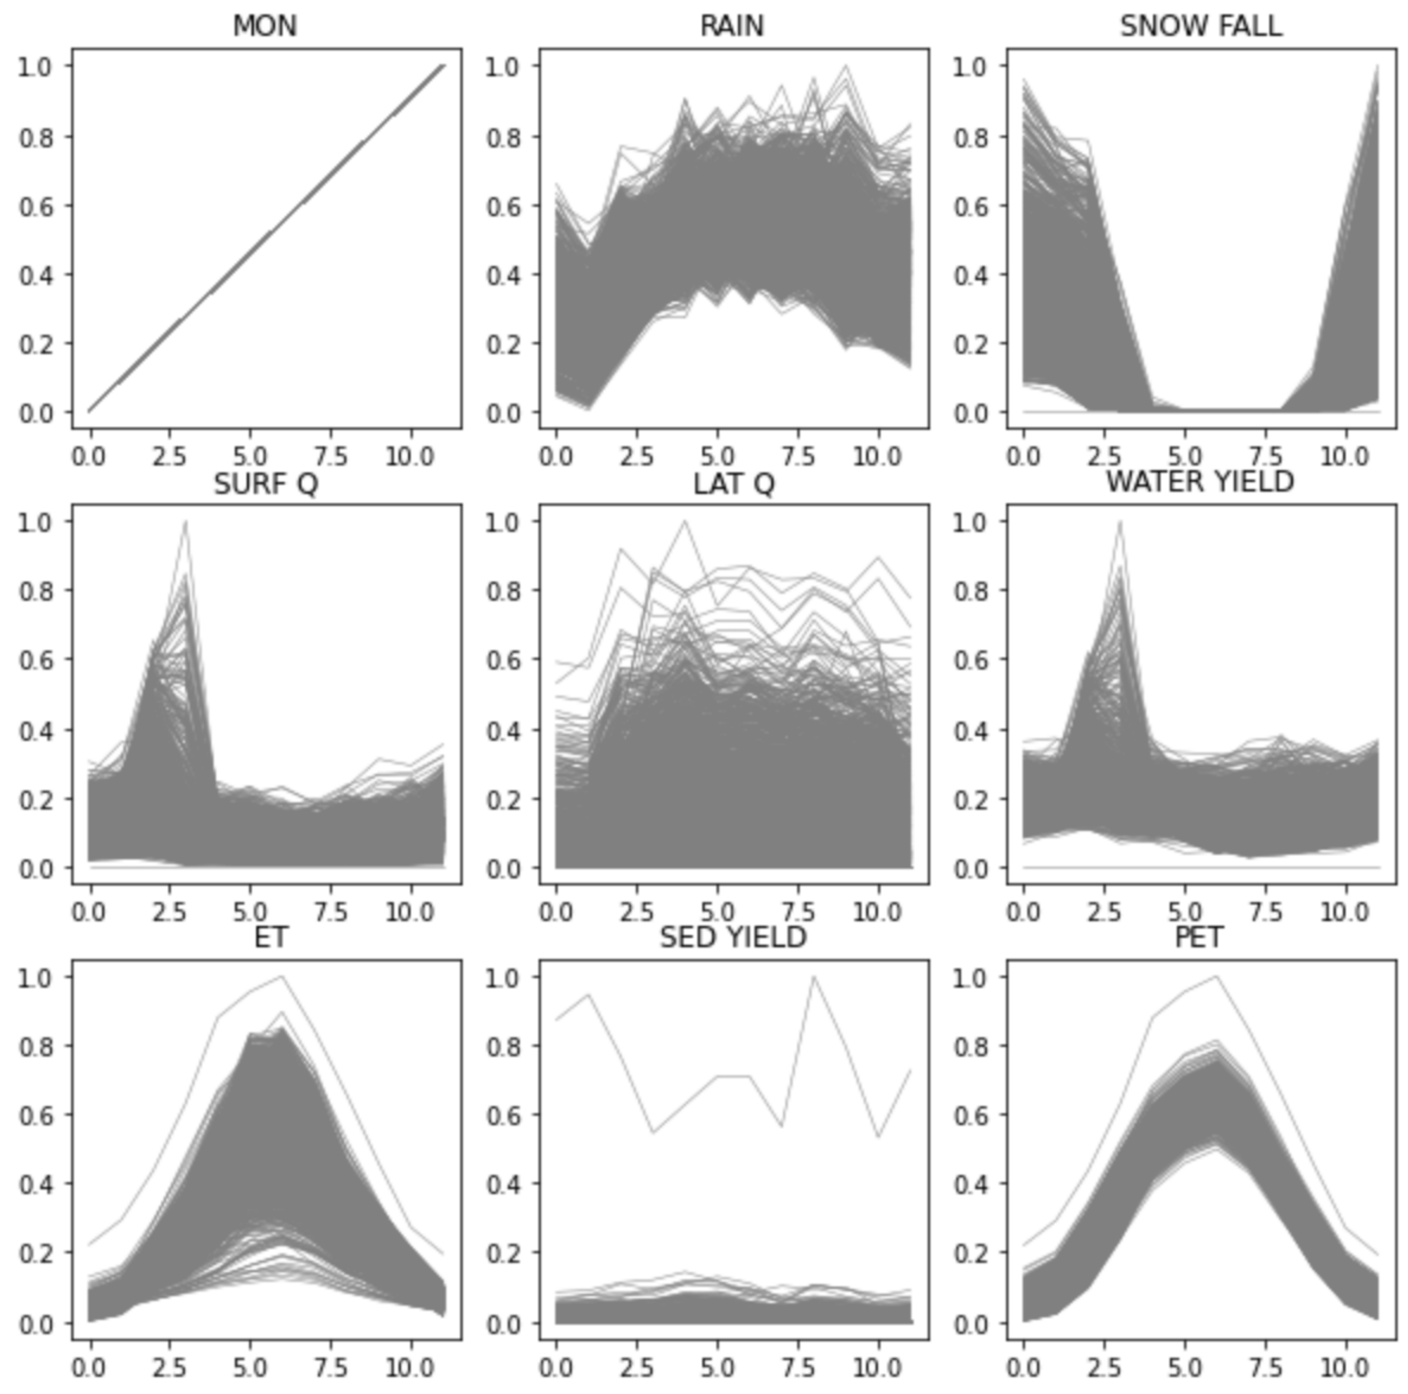

Scaled

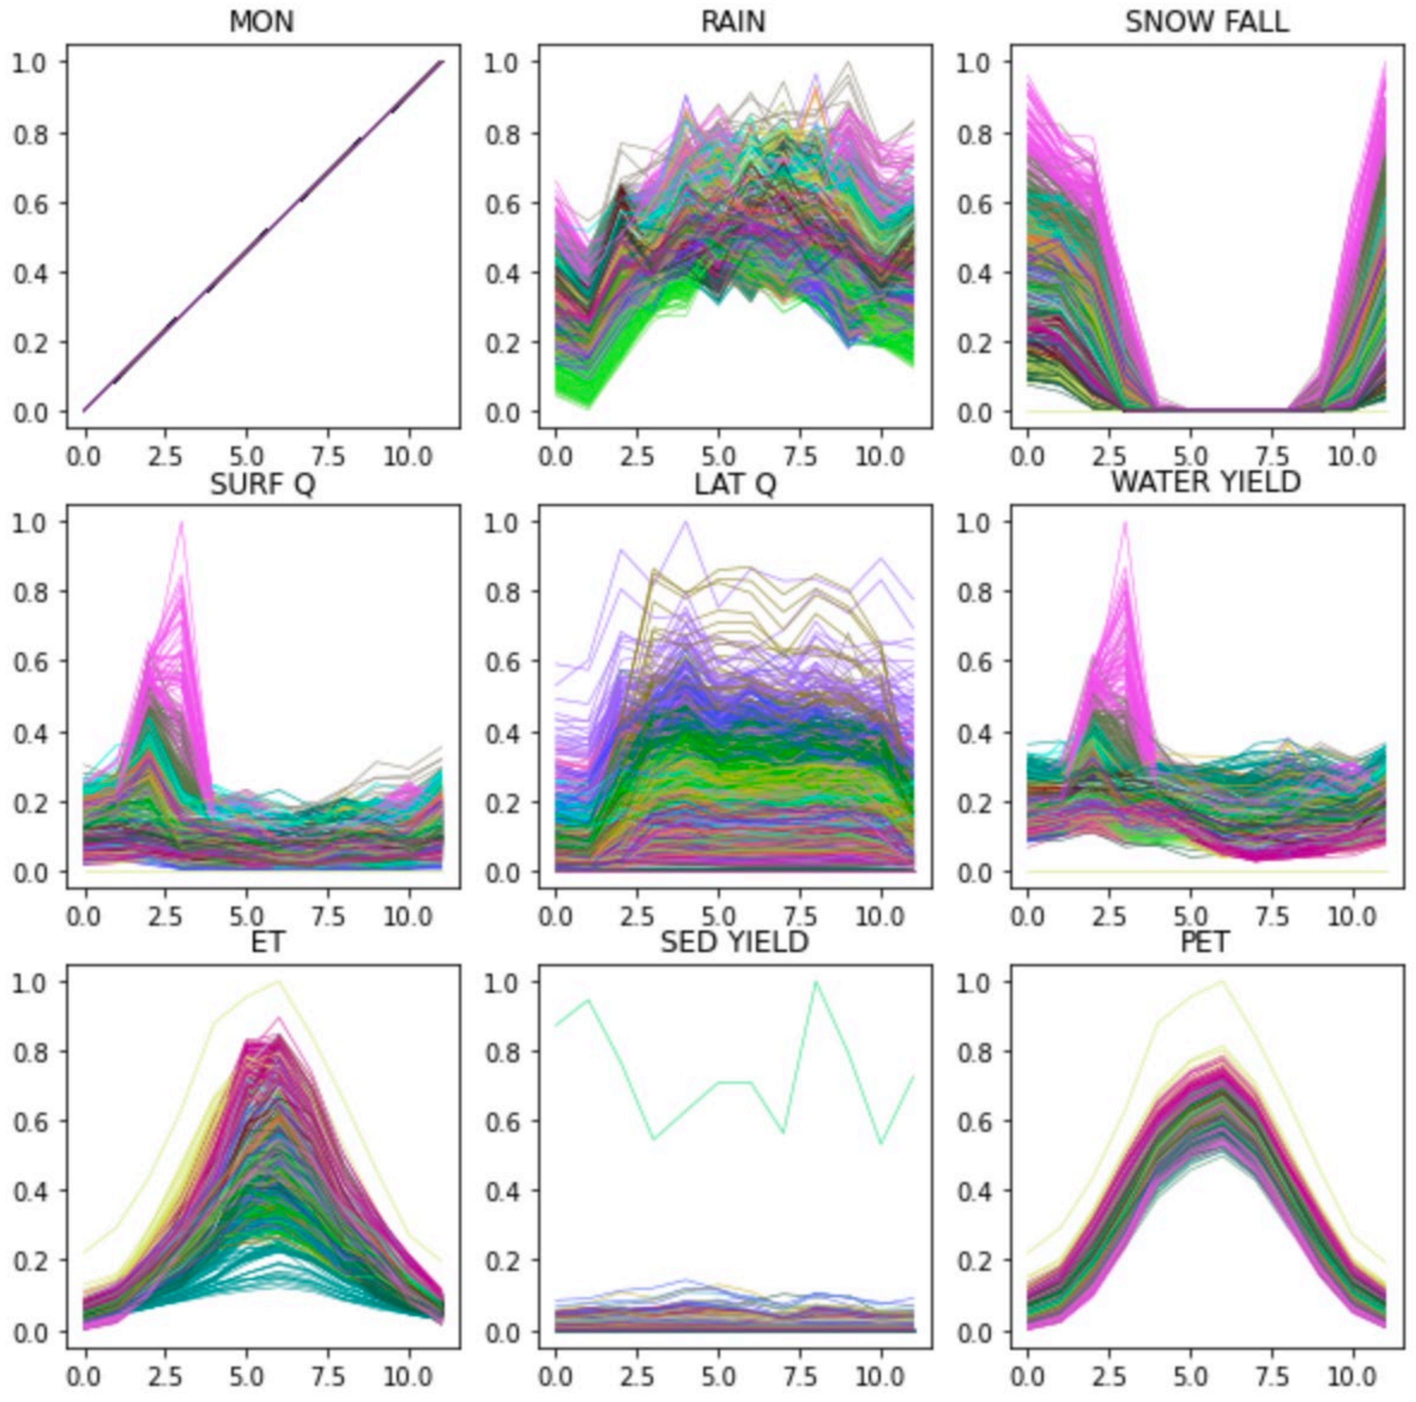

DTW + K-means

Feature Values

Months

Temporal Graph Convolutional Network (TGCN)

A graph is a data structure consisting of two components: nodes (vertices) and edges.

Graph Neural Networks (GNNs) are a class of deep learning methods designed to perform inference on data described by graphs.

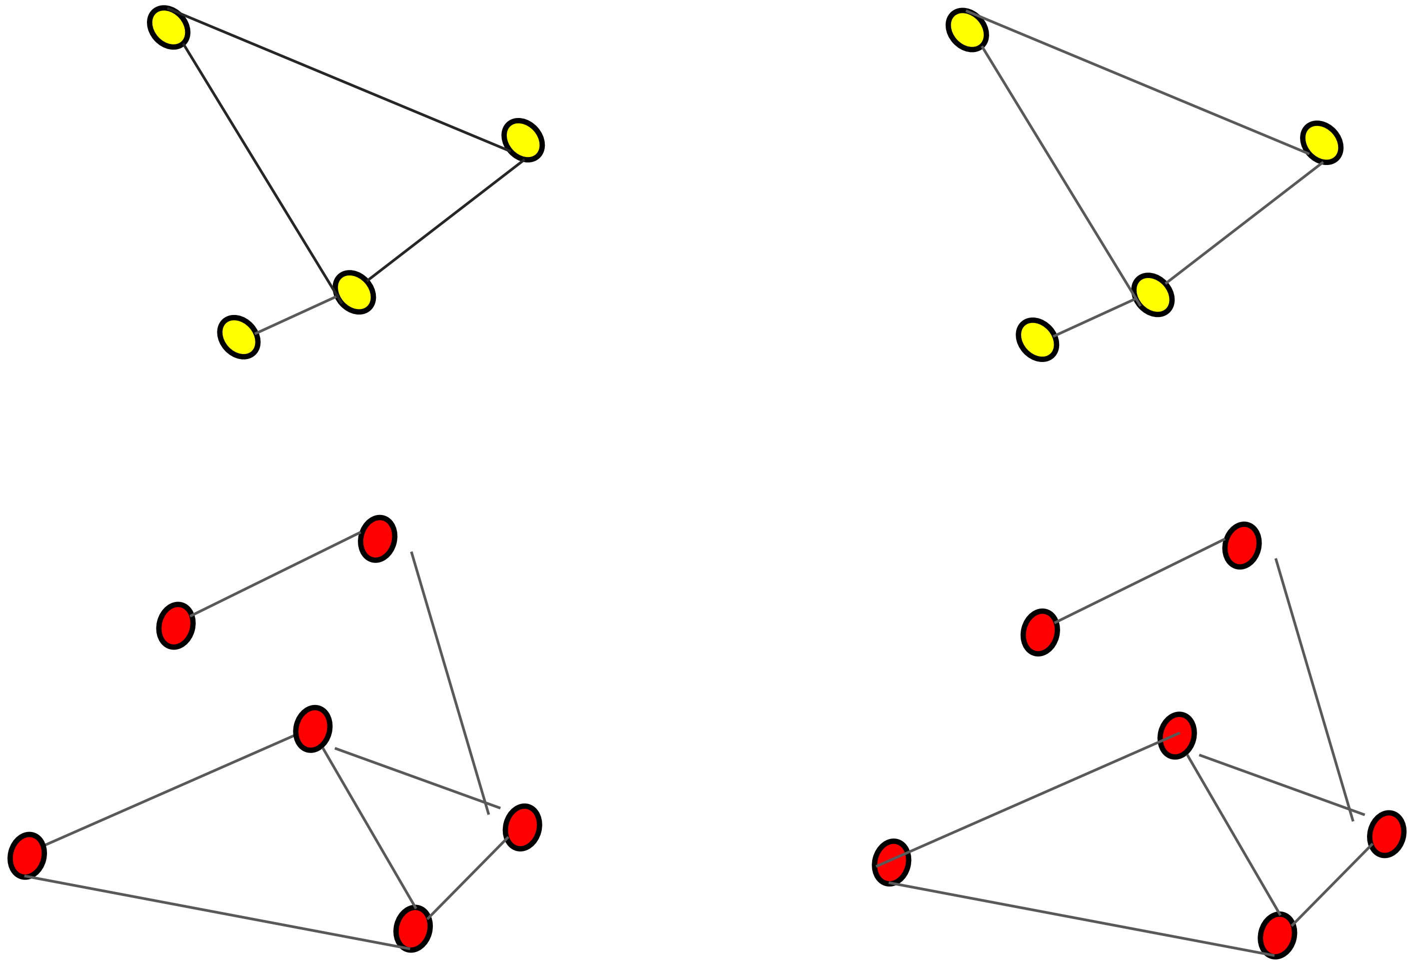

TGCN Implementation

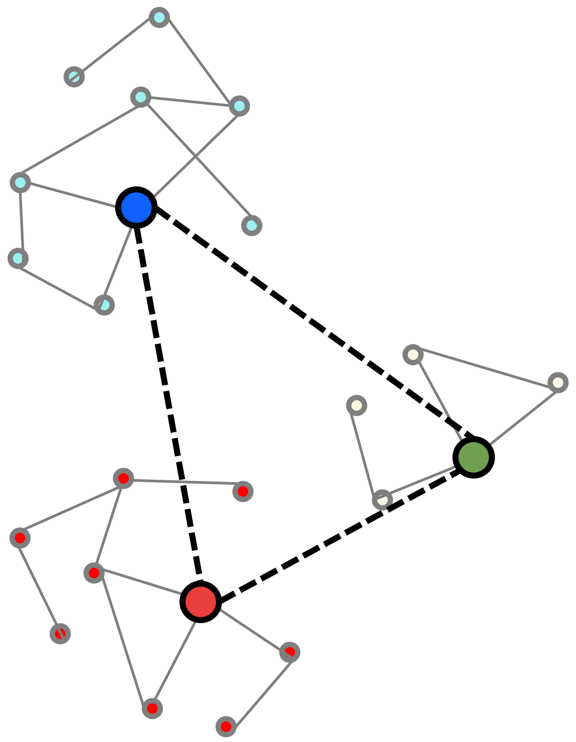

MPNN

Update hidden features of the *central nodes* using inter-graph Message Passing Neural Network (MPNN).

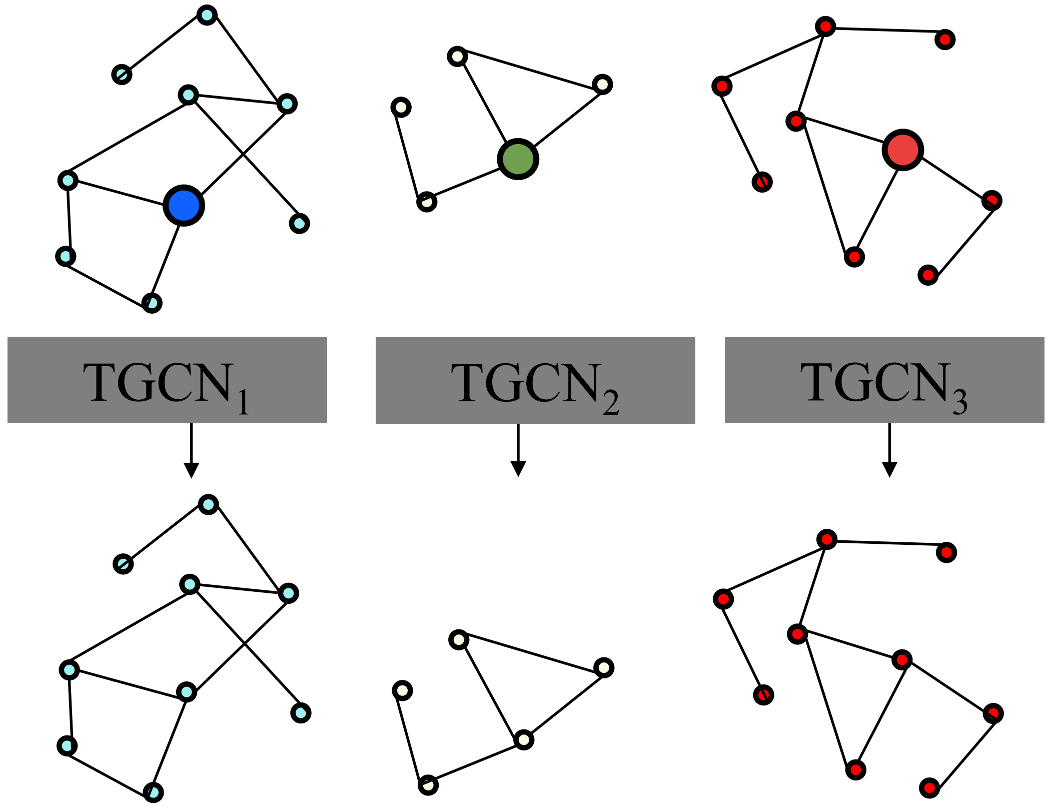

Pre-Trained TGCN

Update node features using a pre-trained Temporal Graph Convolution Neural Network (TGCN) and get soil moisture predictions for each node.

$$\text{Total loss } (\mathbf{L}) = \mathbf{L_1} + \mathbf{L_2} + \mathbf{L_3}$$

Prediction Loss

Aggregate prediction loss across all graphs to get total loss.

Update

Use total loss to update parameters of MPNN.

Pilot Site

Layout of the Mid-Atlantic basin along with its stream network and HUC12 watersheds.

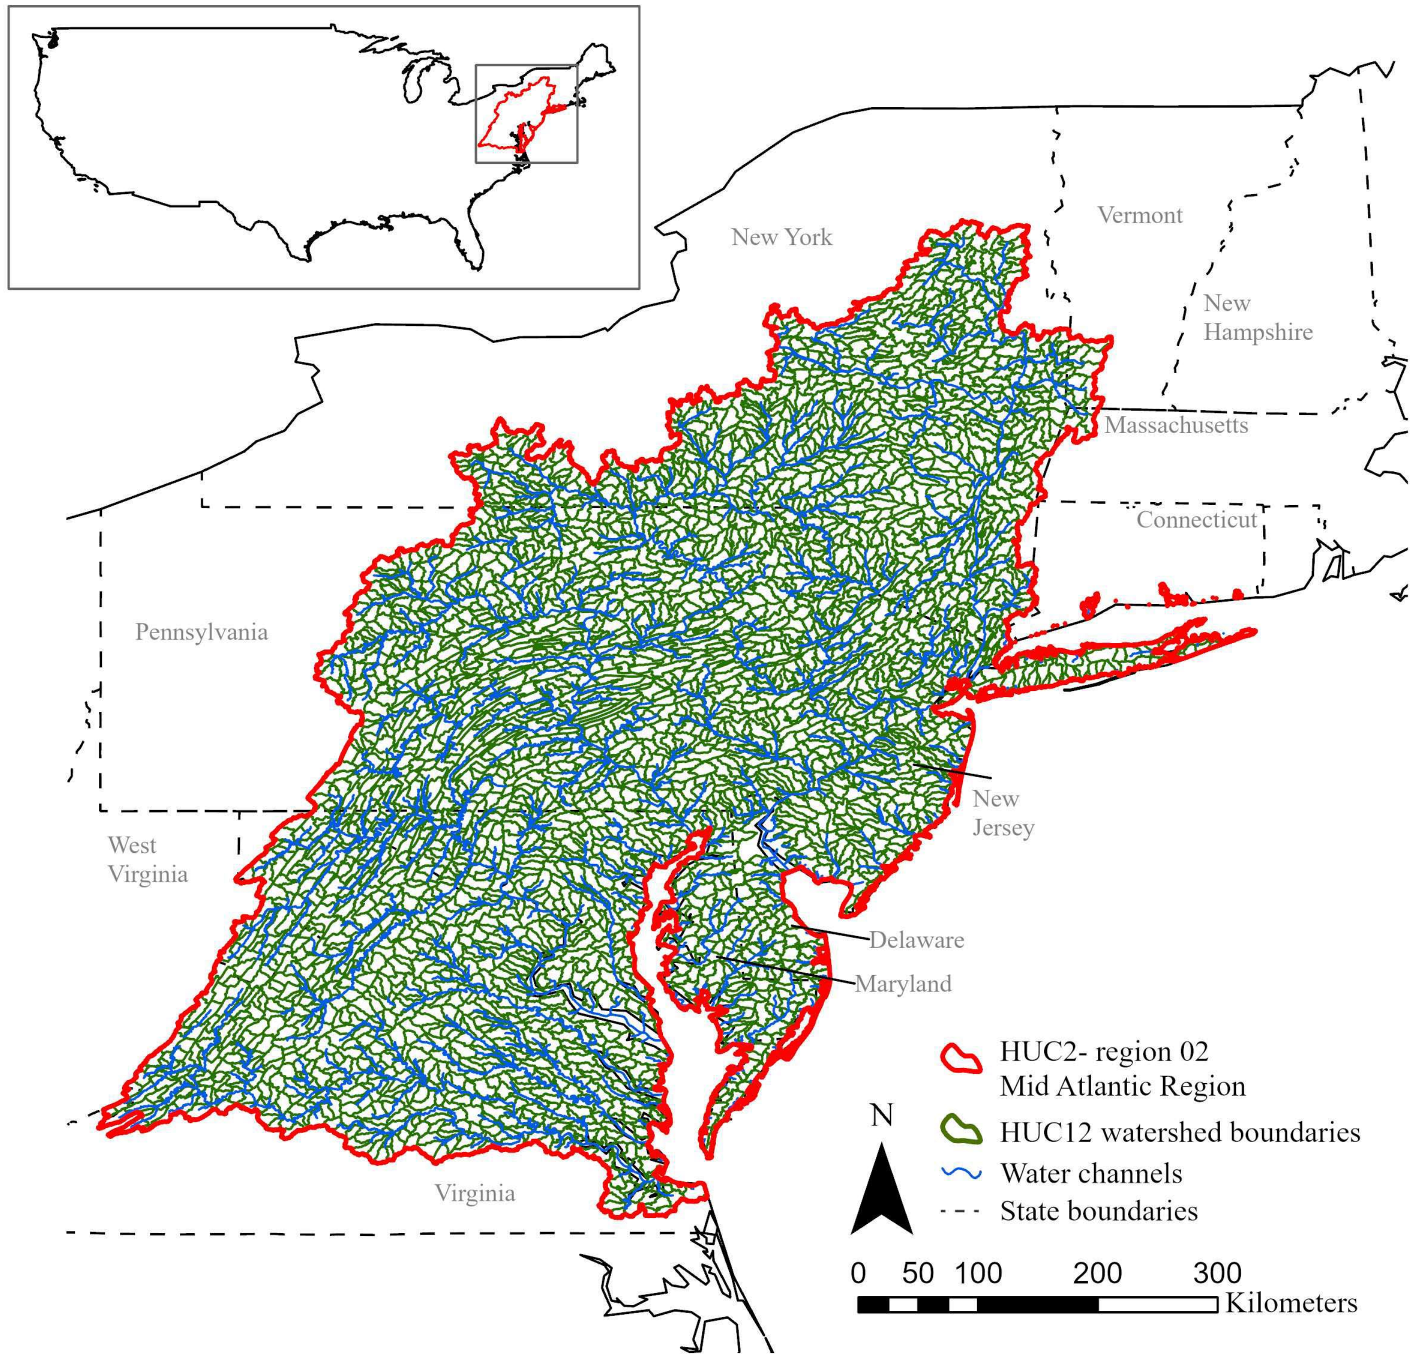

Plot of true soil moisture values of 20 randomly sub-sampled HRUs in selected clusters. Soil moisture in different clusters exhibits distinct seasonal trend.

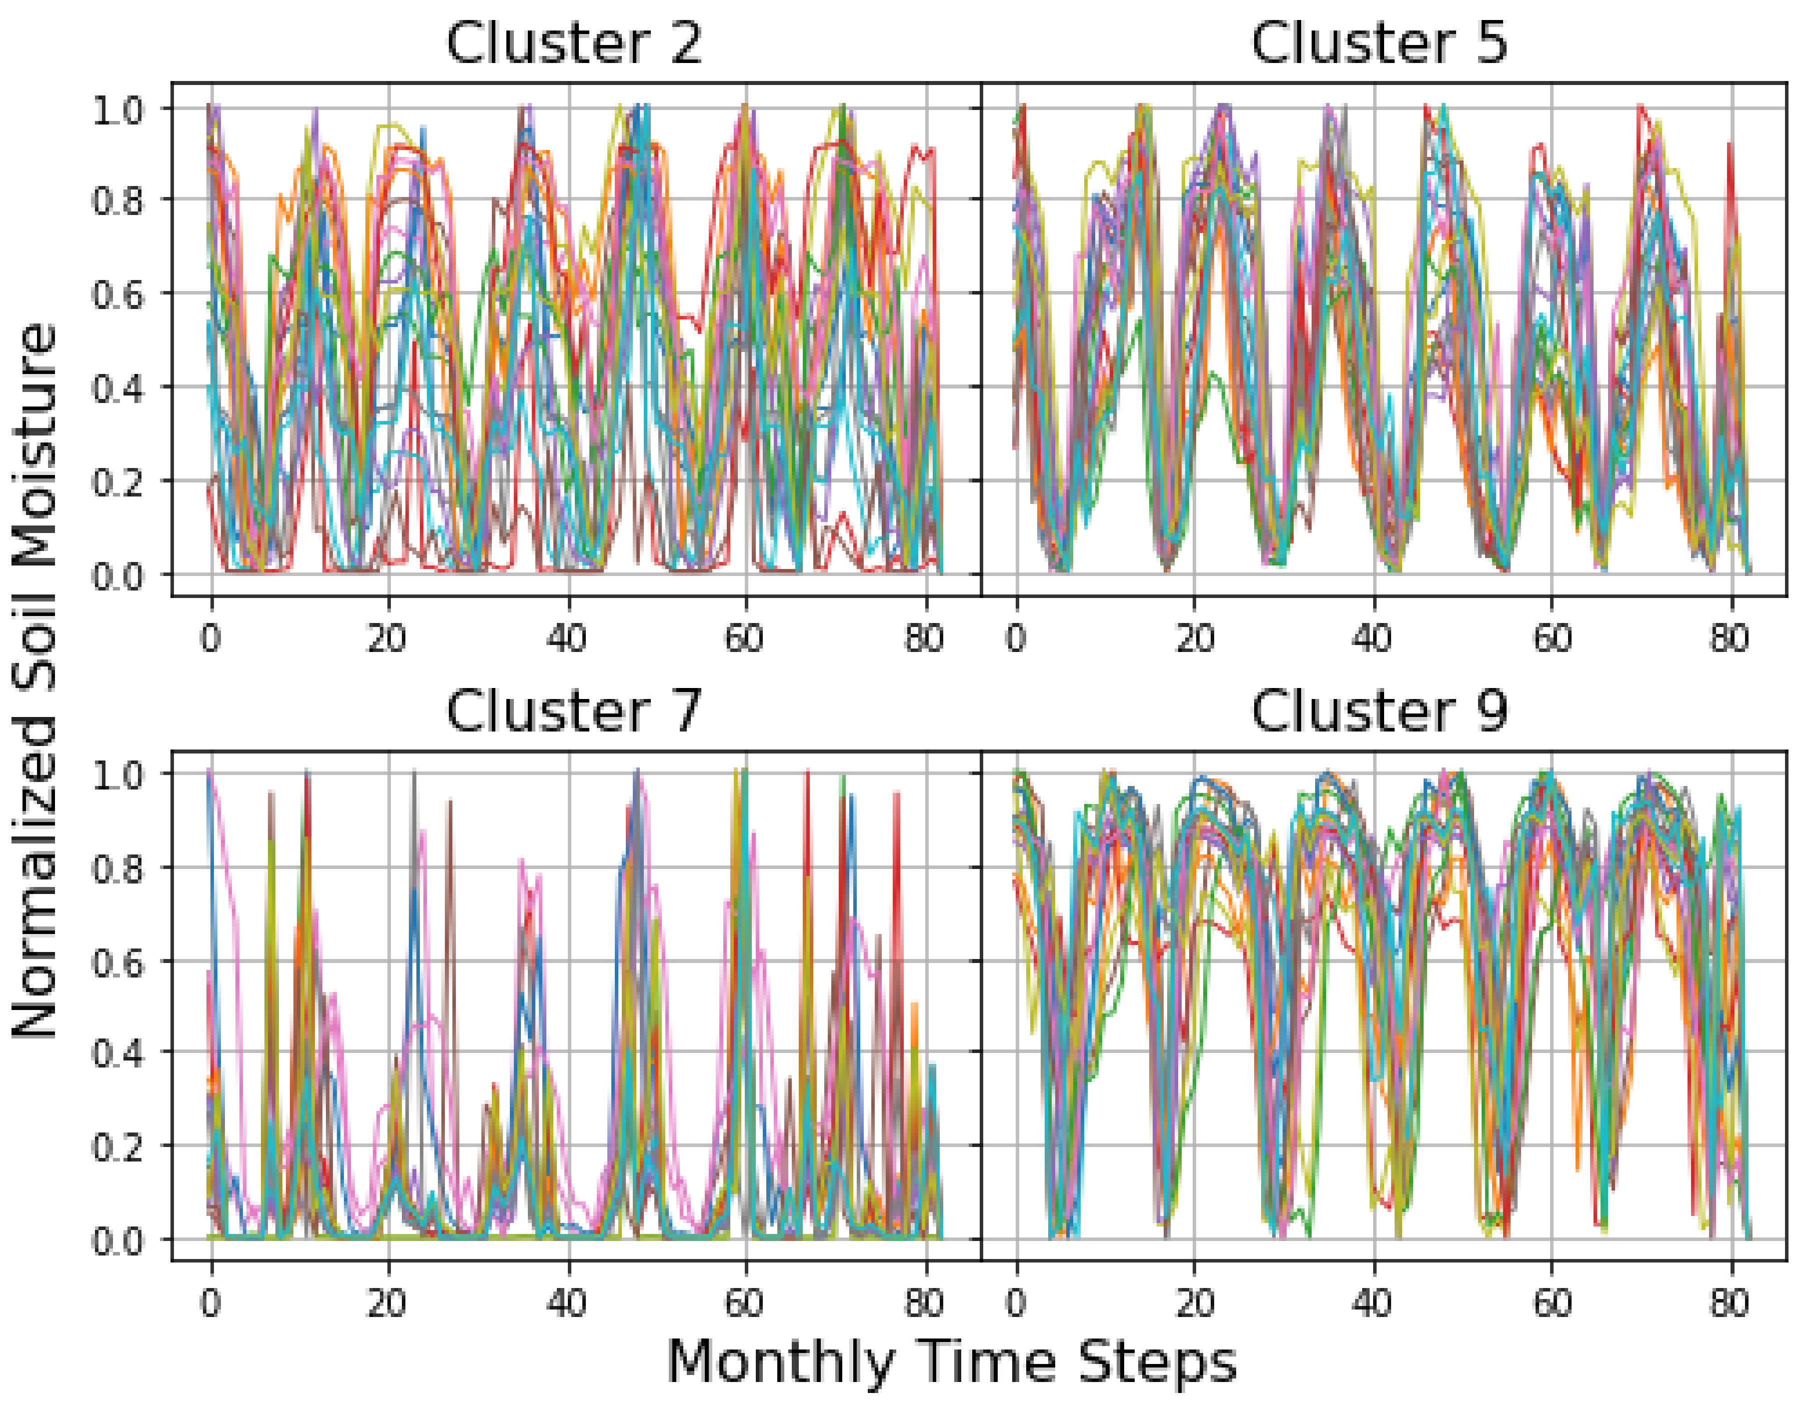

TGCN model performance

Cluster ID	LSTM MSE	C+TGCN MSE	Relative MSE Reduction
1	0.3433	0.0549	82.93%
2	0.3815	0.0573	[illegible]
3	0.3588	[illegible]	[illegible]
4	0.3057	[illegible]	79.60%
5	0.3677	[illegible]	86.06%
6	[illegible]	[illegible]	86.19%
7	[illegible]	0.0417	94.29%
8	[illegible]	0.0393	91.10%
9	0.4227	0.0560	87.42%
10	0.3847	0.0591	83.43%

~80% reduction in MSE

- Captures dynamic properties of soil moisture
- High Goodness-of-fit measure (KGE ~ 0.6)
- Outperforms classic LSTM model by ~80%

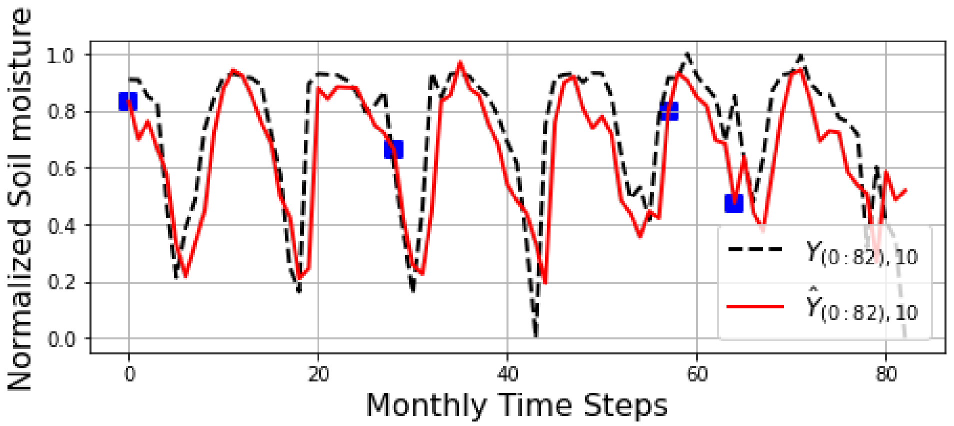

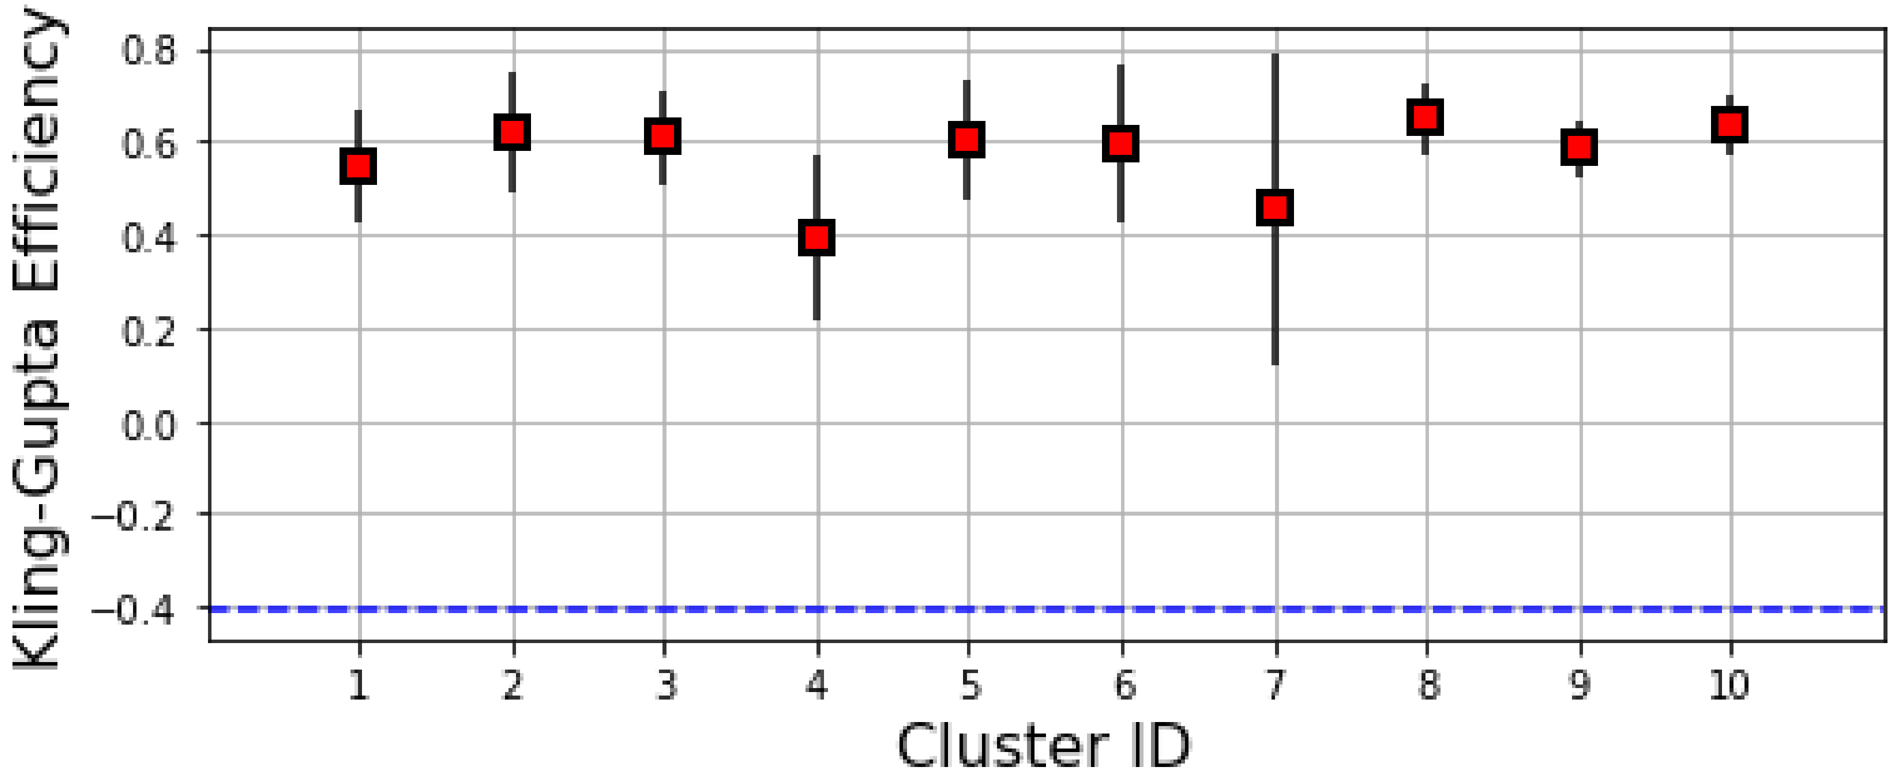

Hitchhiker's Guide to Reinforcement Learning

RL involves a decision-making agent interacting with a virtual environment

- environments often abstract real-world models
- **maximize a cumulative reward by selecting actions that achieve a certain outcome**
- agent balances exploring new actions and exploiting past effective actions

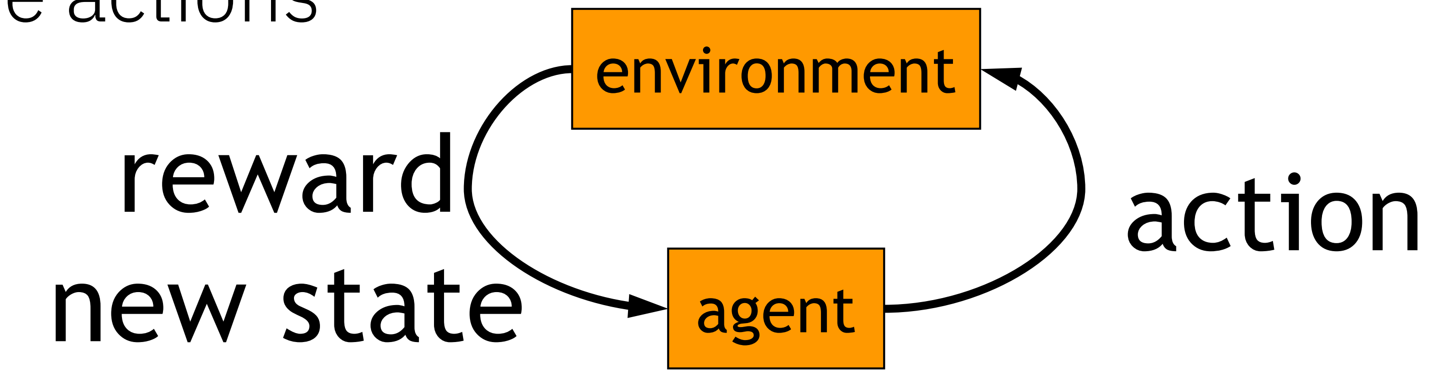

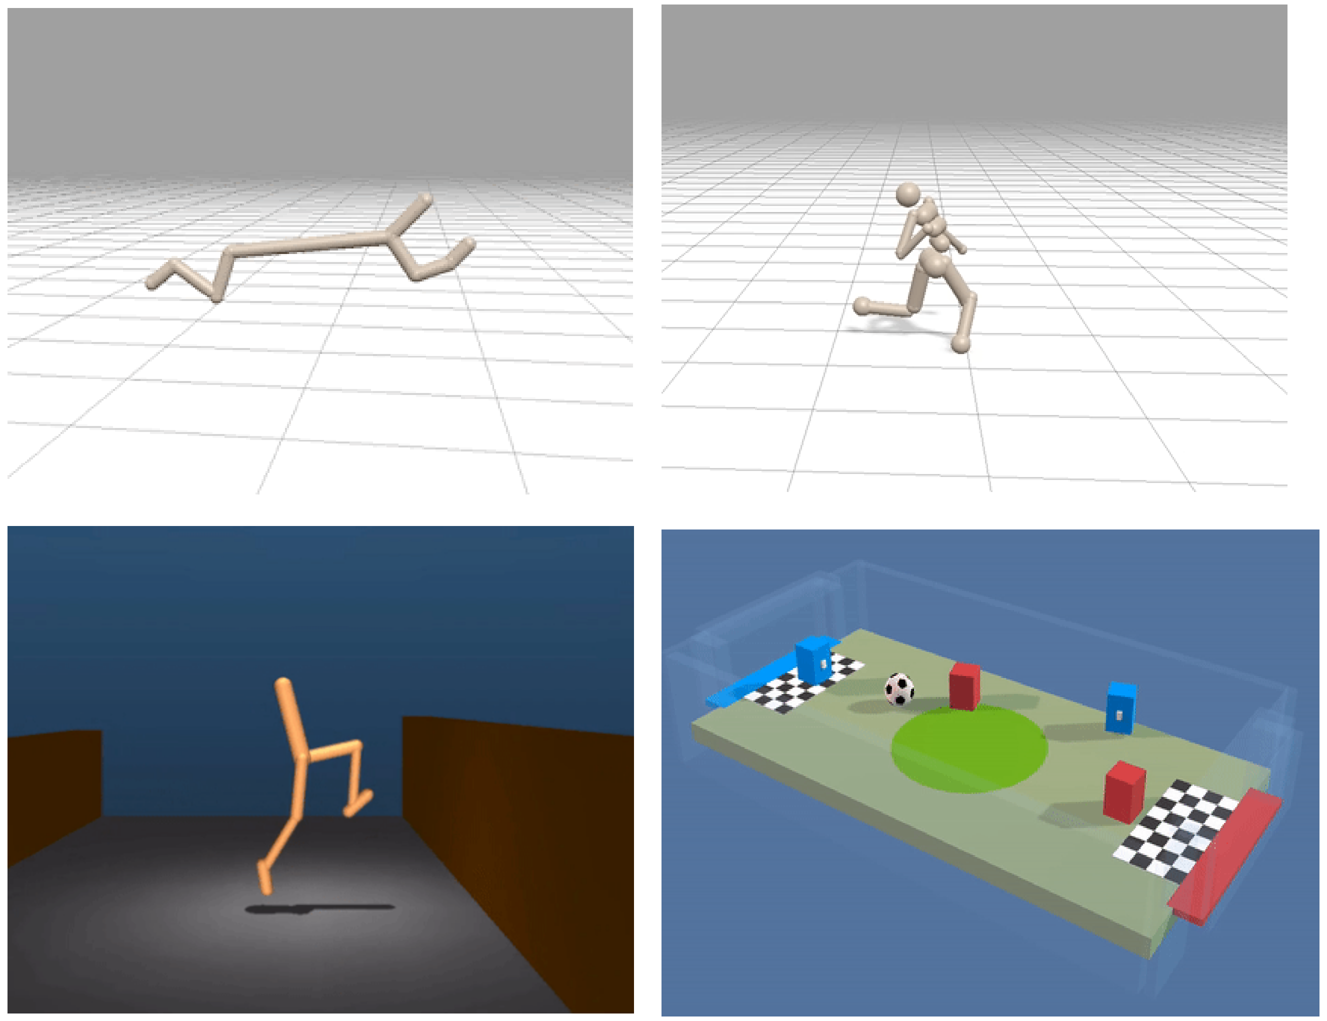

Images from Google AI, DeepMind & endtoend.ai

Objective

A decision support system to aid farmers increase crop yields while minimizing use of inputs e.g., fertilizer and irrigation

Approach

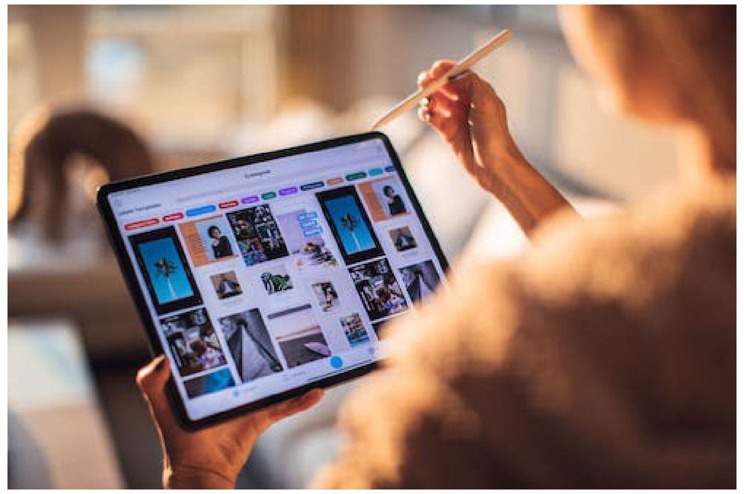

Recommend optimal farm inputs based on Reinforcement Learning.

1. *Simulate crop growth using the Soil & Water Assessment Tool (SWAT)*
2. *Implement various decision-making strategies:*
   - *standard farming practice,*
   - *reactive strategy,*
   - *deep deterministic policy gradient (DDPG)*

Reinforcement Learning Agent-Environment Interface

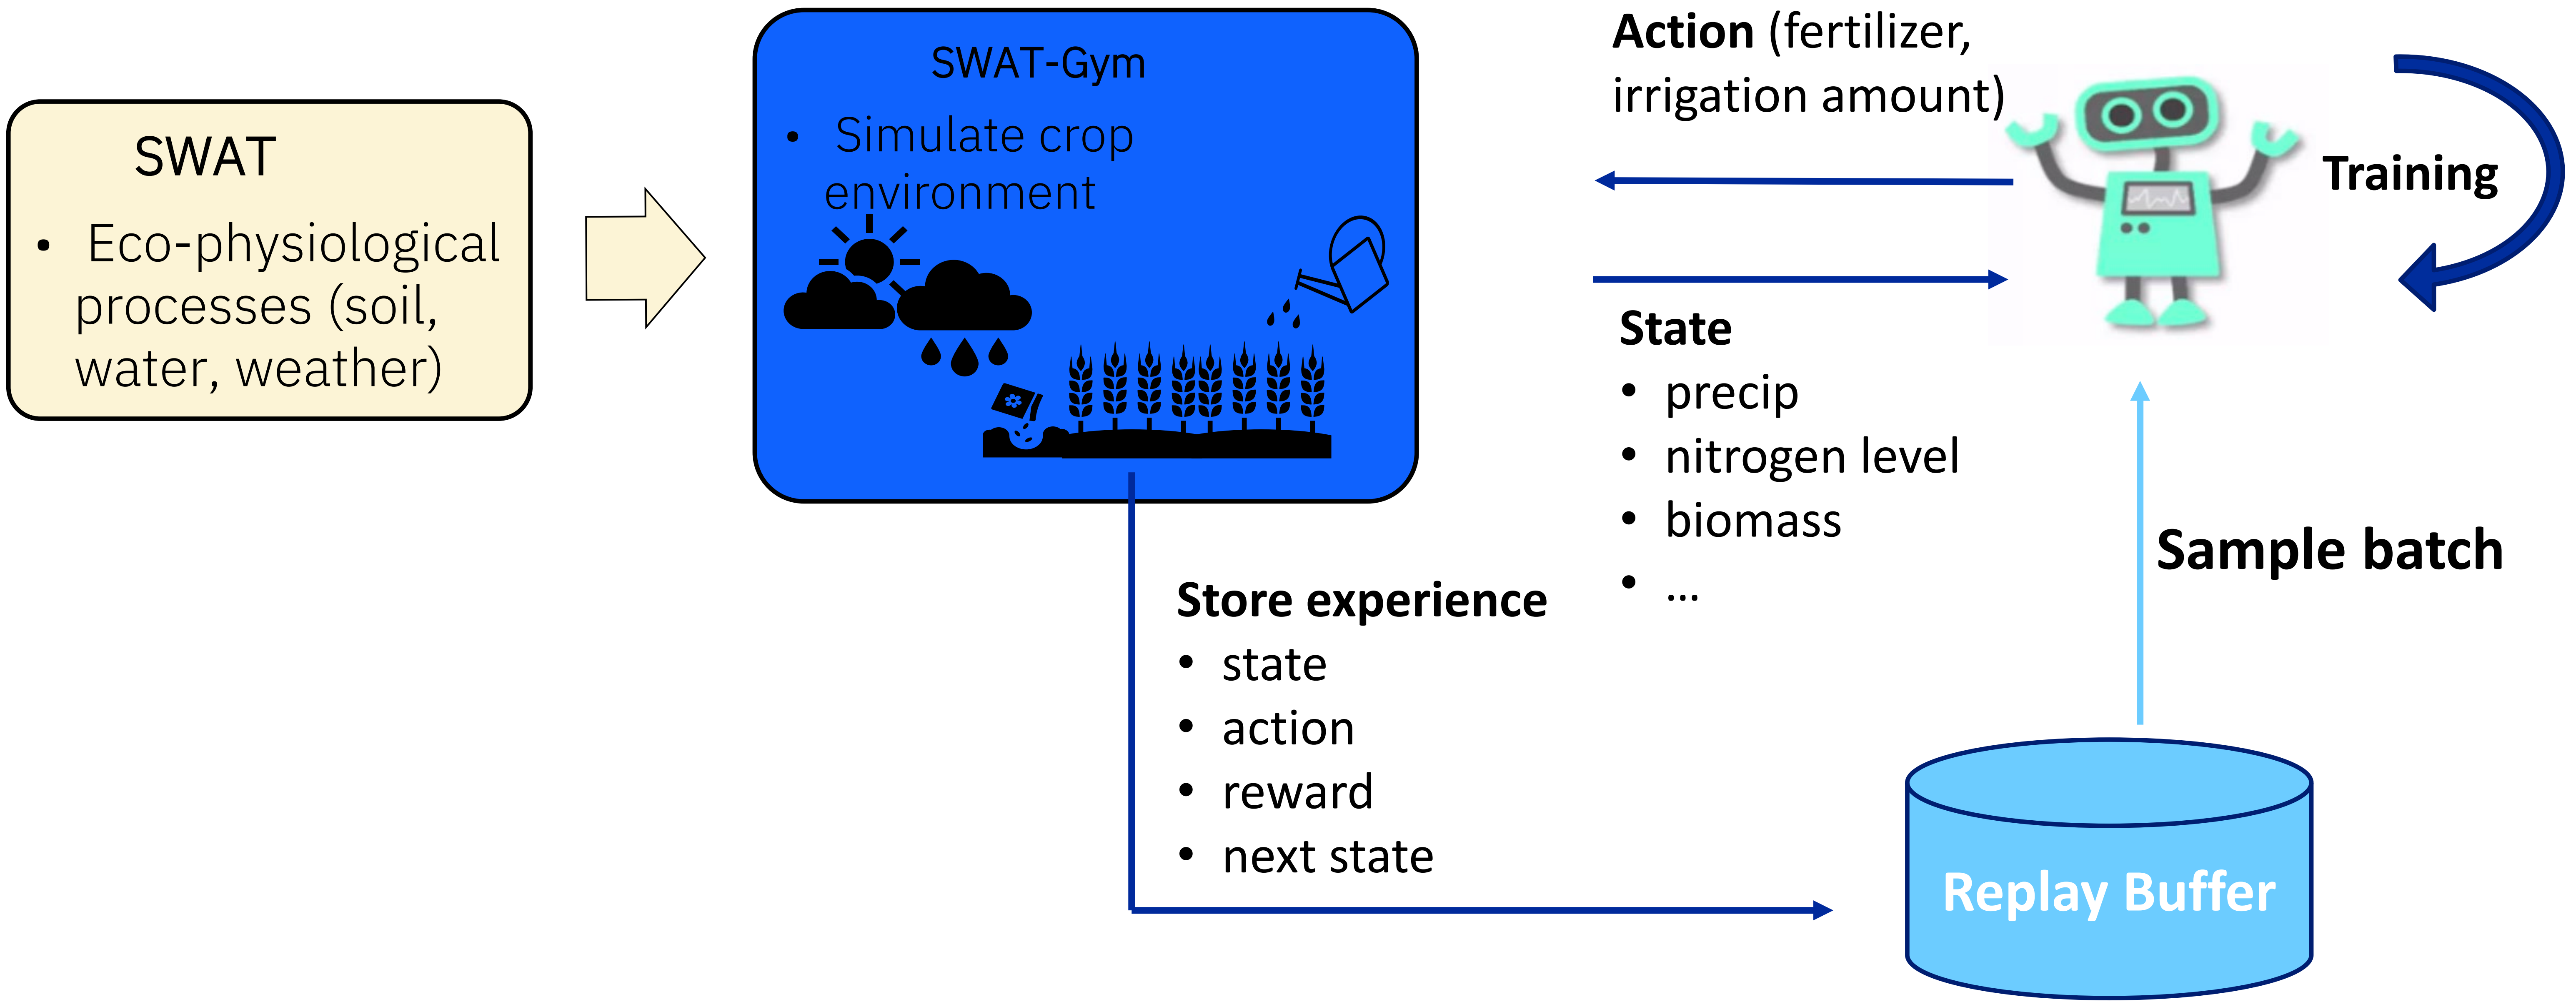

SWATGym Variables

Inputs

- Location
- Simulation duration
- Solar radiation
- Avg air temperature
- Precipitation
- Reference evapotranspiration
- Fertilizer
- Irrigation

\+

SWAT Dynamics

- Radiation Use Efficiency
- Heat units
- Canopy height
- Leaf area index
- Root development
- Evapotranspiration
- Cumulative biomass
- Soil water balance
- Nutrient balance
- Growth stress factors

=

Output

Yield

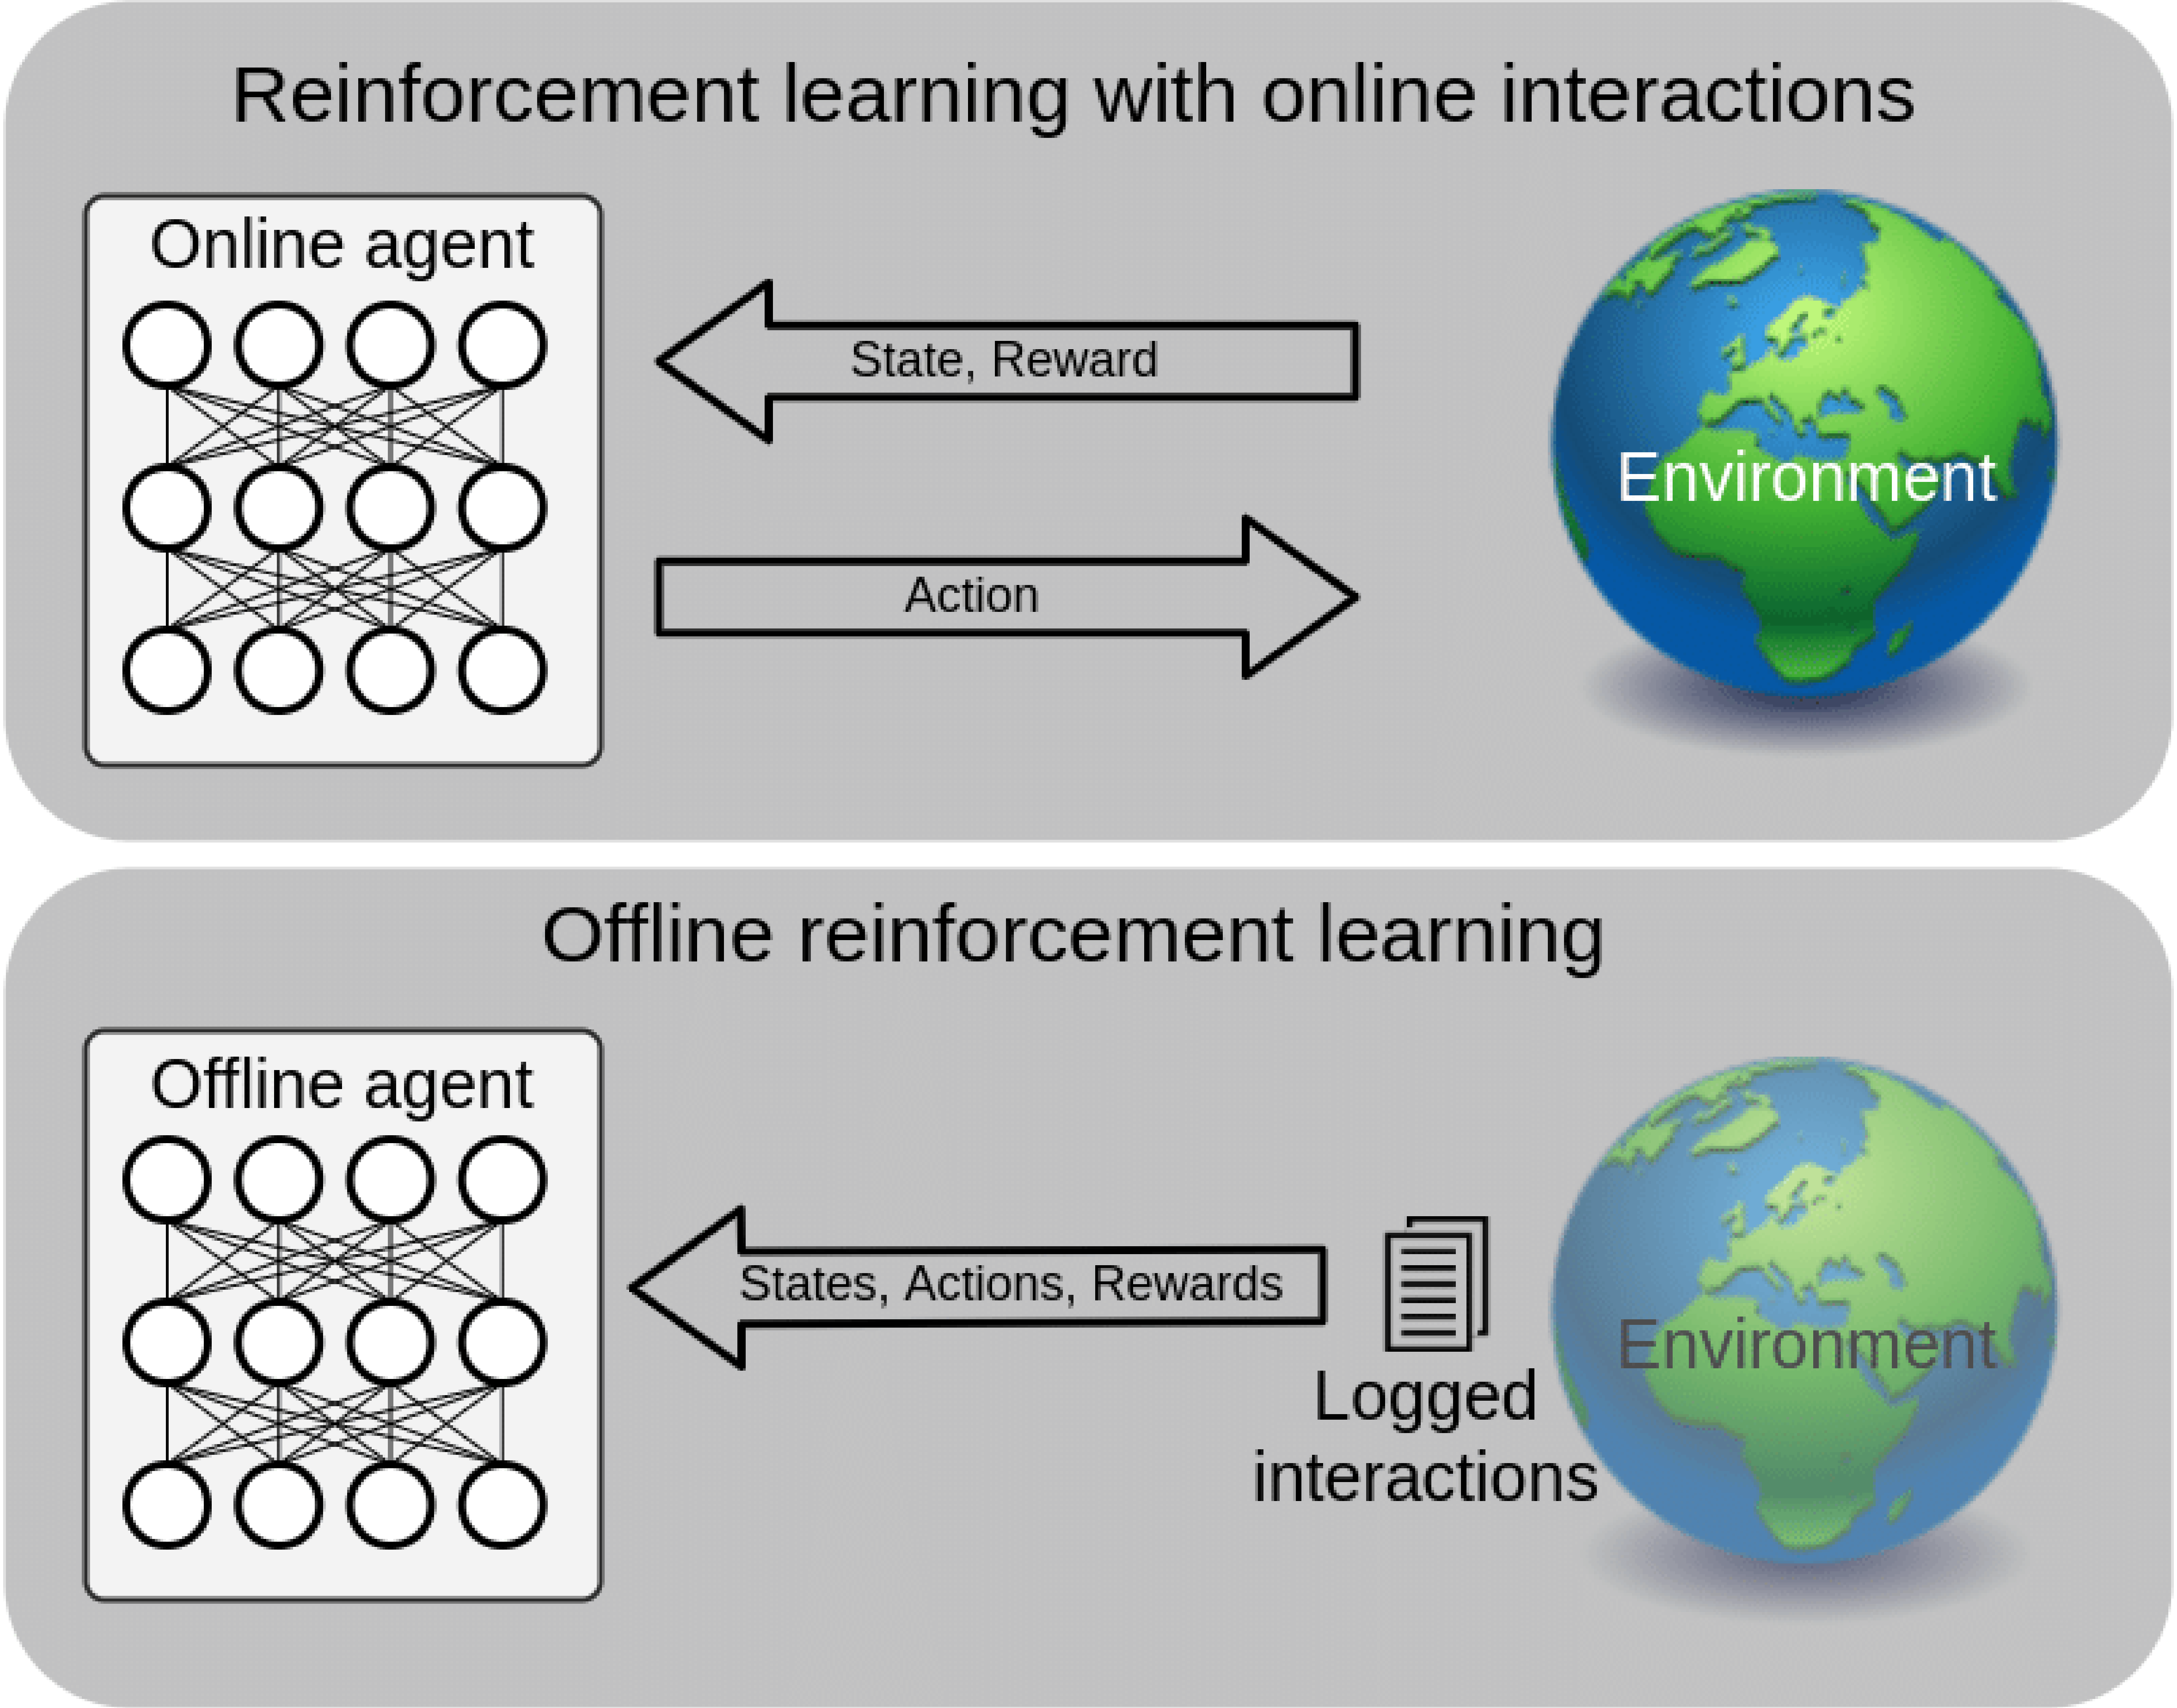

Reinforcement learning with online interactions
Online agent
State, Reward
Action
Environment
Offline reinforcement learning
Offline agent
States, Actions, Rewards
Logged interactions
Environment

Example with Random policy

- **Markov Decision Process (MDP):**
- Set of states S, set of actions A, transition model $P(s, a, s')$, reward function $R(s, a\ )$
- At each time step $s \in S$, agent:
  - receives state s of the env
  - selects an action $a \in A$ to take
  - a reward $R_a(s, s')$ encapsulating evaluating action
  - environment transitions to a new state

```python
from swat_env import SWATEnv

env = SWATEnv()
# env  = gym.make('SWATGym')

state, reward, done, info = env.reset()
rewards = []
max_timesteps = 120
for t in range(max_timesteps):
    action = env.action_space.sample()
    observation, reward, done, info = env.step(action)
    rewards.append(reward)

    if done:
        print(f"Success! Episode finished after {t+1} timesteps")
        break
```

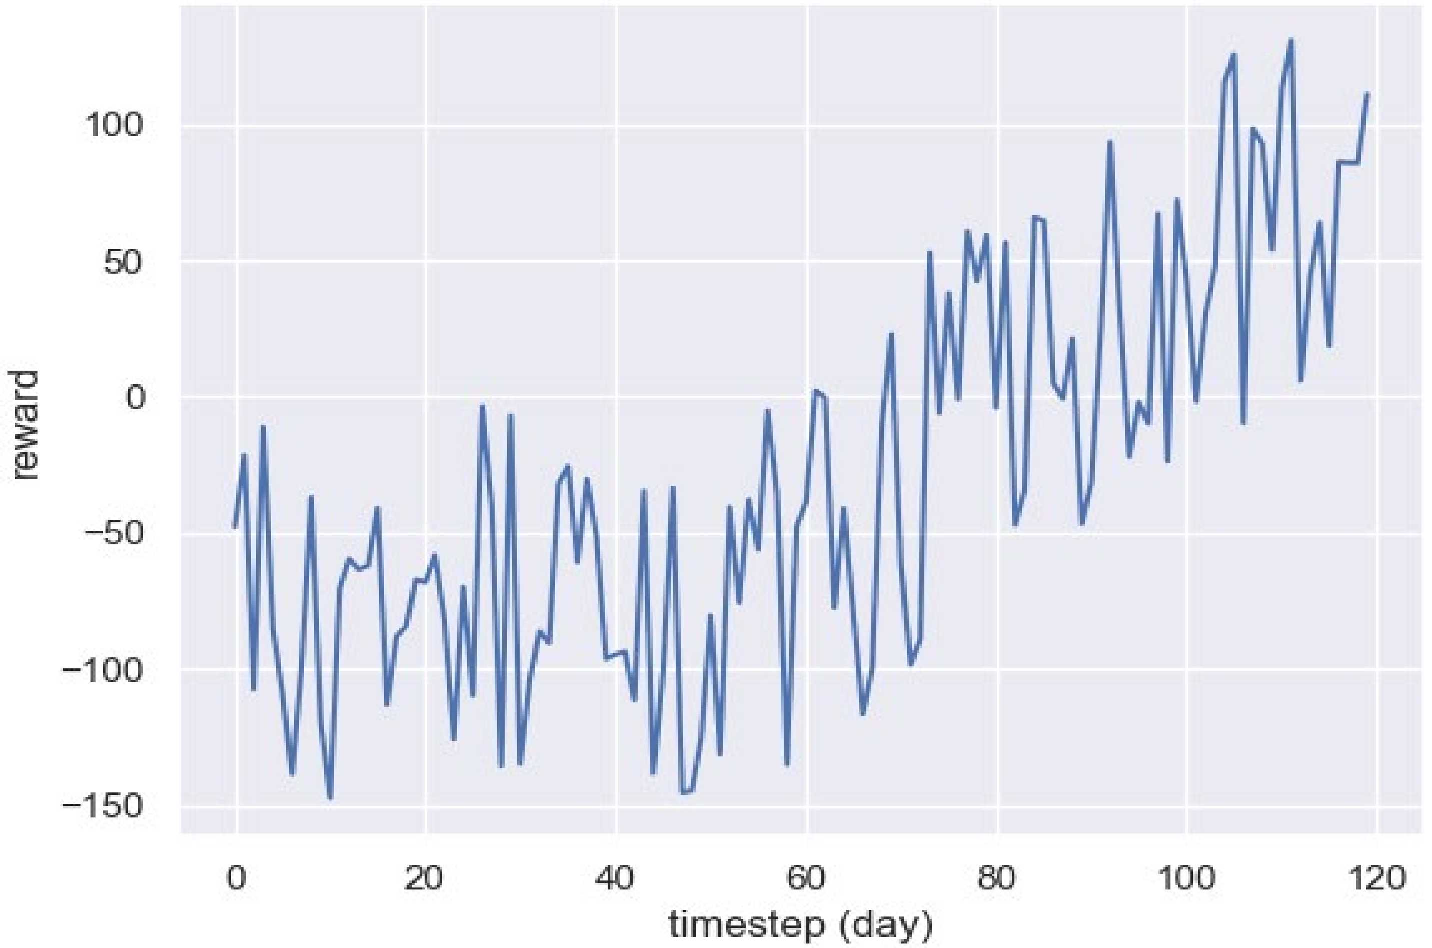

Decision-making Agents

Standard Practice

Applies 3 apps. of 60 kg N/ha and 25 mm H20/ha

Reactive Agent

Applies 120 kg/ha when N drop below 5kg/ha and 50mm SW below 25 mm/ha

Deep Deterministic Policy Gradient [5]

Learns policy to optimally select N and H20 amounts

Conclusions

- Machine Learning can provide unprecedented degree of insight into soil health
- Critical that AI algorithms are adapted to the spatiotemporal properties of geophsysical systems
- Accurate estimation of soil properties critical to improved decision making
- Unsupervised learning and human-centered AI can accelerate evidence-based decision making in agriculture

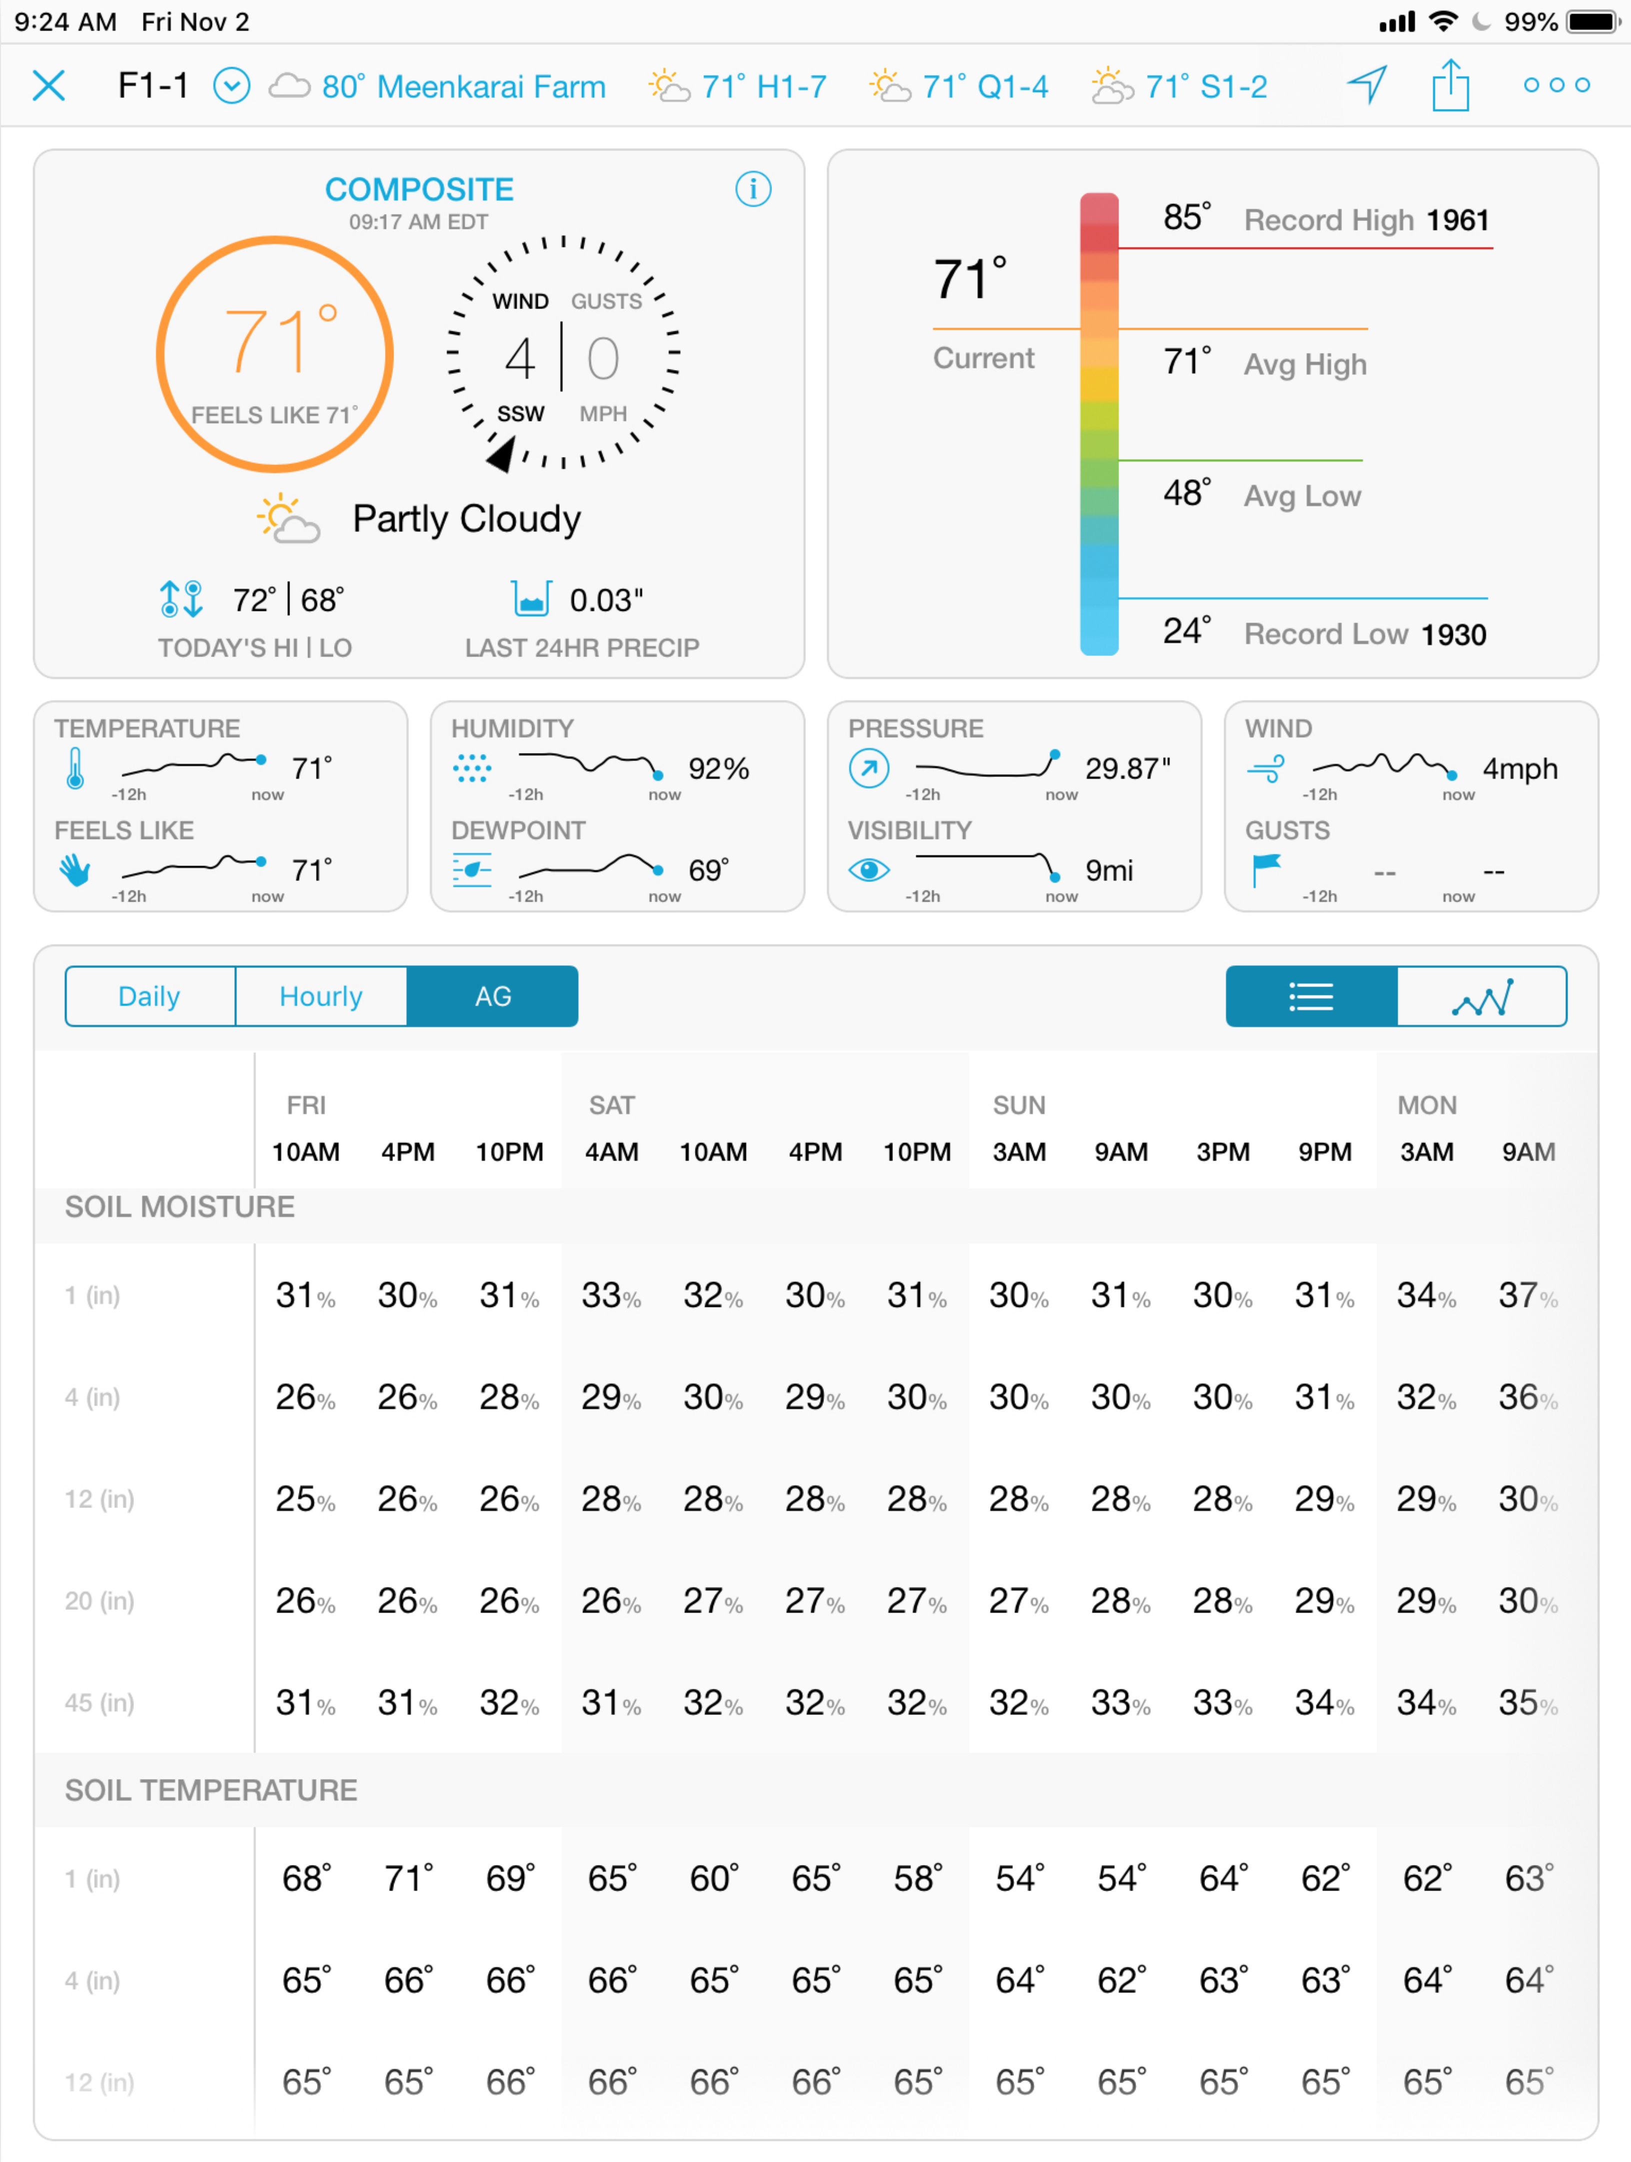

Related Publications:

1) Malvern Madondo, et. al. “A Reinforcement Learning Framework Built Within a SWAT Model Physical Environment to Inform Crop Management” American Geophysical Union Fall Meeting, December 11–17, https://www.agu.org/fall-meeting (2022).
2) Muneeza Azmat, et. al. “Forecasting Soil Moisture Using Domain-Inspired Temporal Graph Convolution Neural Networks to Guide Sustainable Crop Management.” IJCAI, August 19–25 (2023). https://arxiv.org/abs/2212.06565
3) Malvern Madondo, et. al. “A SWAT_based Reinforcent Learning Framework for Crop Management” AAAI’23 AI 4 Social Good Workshop February 11–17, 2023. PrePrint

TEXAS A&M AGRILIFE

IBM