

A SWAT-based Reinforcement Learning Framework for Crop Management

Malvern Madondo¹, Muneeza Azmat², Kelsey DiPietro³, Raya Horesh⁴, Michael Jacobs⁵, Arun Bawa⁶, Raghavan Srinivasan⁶, Fearghal O'Donncha⁴

¹ Emory University; ² Michigan State university; ³ Sandia National Laboratories; ⁴ IBM Research;

⁵ IBM Corporate Social Responsibility; ⁶ Texas A&M AgriLife

malvern.madondo@emory.edu, azmatmun@msu.edu, kdipiet@sandia.gov, {rhoresh@us., michael.jacobs@, feardonn@ie.}ibm.com, {arun.bawa@ag., r-srinivasan@}tamu.edu

Abstract

Crop management involves a series of critical, interdependent decisions or actions in a complex and highly uncertain environment, which exhibit distinct spatial and temporal variations. Managing resource inputs such as fertilizer and irrigation in the face of climate change, dwindling supply, and soaring prices is nothing short of a Herculean task. The ability of machine learning to efficiently interrogate complex, nonlinear, and high-dimensional datasets can revolutionize decision-making in agriculture. In this paper, we introduce a reinforcement learning (RL) environment that leverages the dynamics in the Soil and Water Assessment Tool (SWAT) and enables management practices to be assessed and evaluated on a watershed level. This drastically saves time and resources that would have been otherwise deployed during a full-growing season. We consider crop management as an optimization problem where the objective is to produce higher crop yield while minimizing the use of external farming inputs (specifically, fertilizer and irrigation amounts). The problem is naturally subject to environmental factors such as precipitation, solar radiation, temperature, and soil water content. We demonstrate the utility of our framework by developing and benchmarking various decision-making agents following management strategies informed by standard farming practices and state-of-the-art RL algorithms.

Introduction

In the face of increasing demand for agricultural products together with growing environmental and economic stressors, farming practices that produce higher yields with minimal resources are becoming a key strategy for climate-smart agriculture (Alexandratos and Bruinsma 2012; Muller et al. 2017; Ranganathan et al. 2018). The advent of precision agriculture, an innovative approach to modern farming that combines models and sensor data to improve decision-making, has enabled many farmers to achieve higher crop yields without extending arable land or increasing farming inputs (Dutia 2014). Each growing season, farmers make critical operational decisions such as crop selection, planting and harvesting schedules, and application of fertilizers and irrigation. Complexities are exacerbated by a changing climate and the need to achieve global food security while

minimizing environmental impacts and carbon emissions, as populations continue to grow.

The widespread adoption of agricultural technology and a burgeoning big data ecosystem has provided many machine learning researchers and practitioners with opportunities to develop innovative crop and resource management solutions. Increasingly machine learning is being applied to tasks ranging from soil moisture and crop yield prediction to crop disease detection using satellite image data (Chlingaryan, Sukkariéh, and Whelan 2018; Gandhi 2022; Liakos et al. 2018). Muller et al. (2017) proposed various sustainable food production solutions, including increasing crop production on current arable land, adopting greener farming practices, and reducing food consumption and wastage. We consider the first approach and make contributions toward optimizing crop yields using Reinforcement Learning (RL), a sub-field of machine learning that has proven its utility in decision-making across multiple domains such as healthcare, engineering, and games (Sutton and Barto 2018).

Agriculture use cases have focused on optimizing crop production subject to various resource constraints such as fertilizer application and water usage (Binas, Luginbuehl, and Bengio 2019; Elavarasan and Vincent 2020; Overweg, Berghuijs, and Athanasiadis 2021; Wu et al. 2021). This paper employs RL to autonomously learn the optimal set and distribution of actions to direct crop growth, cognizant of spatial and temporal variations. We leverage the Soil and Water Assessment Tool (SWAT) (Arnold et al. 1998; Gassman et al. 2007) numerical simulator to generate representative soil-water-plant synthetic data and to develop an RL environment on which various decision-making agents can be evaluated. Our SWAT-based model incorporates dynamic representations of economic cost, environmental impact, and crop yield. It also produces a reward that characterizes the effects of different agent actions on crop production. By favoring less frequent and minimal amounts of agricultural inputs, it generates a holistic framework to minimize environmental impact and operational costs while still recovering optimal yields.

Contributions. Our key contributions are as follows:

1. We introduce SWATGym, an OpenAI Gym environment modeled after the widely used Texas A&M SWAT model (Arnold et al. 2012). SWATGym simulates crop growth

and incorporates geographic and weather data to model the complex soil-water-plant-atmosphere system as described by the SWAT model.

2. We provide a simple API to apply custom crop management strategies and/or evaluate them against state-of-the-art RL-based strategies. This presents an opportunity to assess the impact and practicality of a strategy before real-world deployment.
3. We highlight some promising future research directions in the SWATGym environment, e.g., fine-tuning it to be self-contained and capture the spatial-temporal variability of the seasons.

Background and Related Work

Nitrogen and water are the primary limiting factors for crop production (Rimski-Korsakov, Rubio, and Lavado 2009). Limiting their amounts may hinder crop growth potential while excess amounts are not only costly but can lead to environmental harm downstream (e.g., nutrient loading in rivers and depleted groundwater levels). Further, the optimal amounts are highly sensitive to local variables such as soils and climate. Applying correct amounts of nitrogen fertilizer and irrigation can optimize yields while reducing excess nutrient runoff to local groundwater, river, and lake systems. However, determining the input application rates and schedules that meet environmental and economic goals is not easy. Complexities are amplified by uncertainties such as weather, resource availability, and costs, among other factors. Further, the optimal values are highly dependent on the stage of growth of the crop. As a result, techniques that can inform optimal crop management policy within a complex, dynamic environment are critical to more resilient and sustainable farming methods. Moreover, controlling for the environmental and financial impacts of crop management is central to the tenets of sustainable agriculture in a highly volatile market (Binas, Luginbuehl, and Bengio 2019; Dutia 2014).

Excessive fertilizer application elevates nitrate levels in the soil. Aided by percolating water, nitrogen may leach down to groundwater or get washed away by surface runoff towards rivers and streams (Follett and Hatfield 2001; Sharp-ley, Smith, and Naney 1987). High nitrate levels are considered harmful to humans and livestock and have been found to affect water quality. Excessive irrigation applications only magnify this problem and wash away fertilizer and other potentially harmful substances into rivers and streams. This leads to environmental degradation, non-optimal crop growth, and economic losses. For sustainability efforts and climate-smart agriculture, models that consider the long-term impact of management practices on water quality are vital.

Reinforcement Learning for Crop Management

Reinforcement learning (RL) involves a decision-making agent interacting with an environment in order to learn a reward-maximizing strategy (Sutton and Barto 2018). RL has only recently been applied to crop management with notable works including (Binas, Luginbuehl, and Bengio 2019)

which explores applying RL to sustainable agriculture as well as (Ashcraft and Karra 2021) and (Overweg, Berghuijs, and Athanasiadis 2021) which optimizes for crop yield subject to irrigation and fertilizer management actions.

Problem formulation. We consider an agent interacting with a simulated environment modeled as a finite horizon Markov Decision Process comprising of a tuple $(\mathcal{S}, \mathcal{A}, P, r, \gamma)$ described by the following:

- a continuous state space \mathcal{S} , consisting of soil, hydrological, and plant data as well as climate inputs.
- a continuous action space \mathcal{A} , where each action a consists of a tuple of fertilizer and irrigation amounts.
- a state transition function $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, which characterizes the probability distribution over states at time $t + 1$ given the state and action taken at time t .
- a measurable reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which describes the effect of an agent’s choice of actions on crop production.
- $\gamma \in [0, 1)$ is the discount factor applied to future returns.

The goal is to maximize the cumulative rewards over a finite horizon of length T , corresponding to a growing season. At each discrete time step $t \in T$ and given a state $s \in \mathcal{S}$, the agent selects an action $a \in \mathcal{A}$ informed by its policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$. The agent then receives a reward r and a new observation of the environment s' . The RL objective is to find the optimal policy π^* that maximizes the cumulative reward.

Crop Growth models

Simulated environments are fast becoming the key to the creation of state-of-the-art algorithms for various learning tasks as evidenced by the success of RL algorithms in games such as Go (Silver et al. 2016) and robotics. Most existing environments are designed to manage either irrigation operations or fertilizer operations. The most notable environment is the Python Crop Simulation Environment (de Wit 2018) which houses various crop models such as WOFOST (Van Diepen et al. 1989) and LINTUL3 (Shibu et al. 2010) and has inspired environments such as CropGym (Overweg, Berghuijs, and Athanasiadis 2021). One drawback of PCSE is that it requires manual calibration of various model components (soil, crop, weather) as well as the specification of the agro-management activities that will take place.

The CropGym environment simulates winter wheat growth in the Netherlands, with a particular focus on fertilizer management. Other seminal models only focus on irrigation management and include the SIMPLE model (Zhao et al. 2019), implemented as an OpenAI environment in (Ashcraft and Karra 2021) and applied to potato growth simulation in Washington State. (Chen et al. 2021) also focuses on irrigation management and simulates paddy rice growth in China. In contrast to these environments, our proposed environment focuses on both fertilizer and irrigation management and considers processes at the watershed level.

Environment	Fertilizer	Irrigation
(Ashcraft and Karra 2021)	✗	✓
CropGym (Overweg, Berghuijs, and Athanasiadis 2021)	✓	✗
(Chen et al. 2021)	✓	✗
SWATGym	✓	✓

Table 1: Comparison of crop growth reinforcement learning environments.

SWAT

The Soil and Water Assessment Tool simulates physical processes such as crop growth, soil water balance, and nutrient cycling in a watershed (Arnold et al. 1998; Gassman et al. 2007). SWAT primarily considers two production levels: potential production, which represents estimated growth under optimal conditions, and actual production, which is limited by factors such as temperature stress, and nutrient and water availability. SWAT uses a simplified version of the EPIC crop model (Williams et al. 1989) to compute the crop-related variables.

The SWATGym Environment

The main contribution of our paper is SWATGym, an RL environment based on SWAT that simulates a crop’s phenological development, growth, and yield on a daily basis by taking into account the effects of factors such as nutrient cycling, water availability, and temperature.

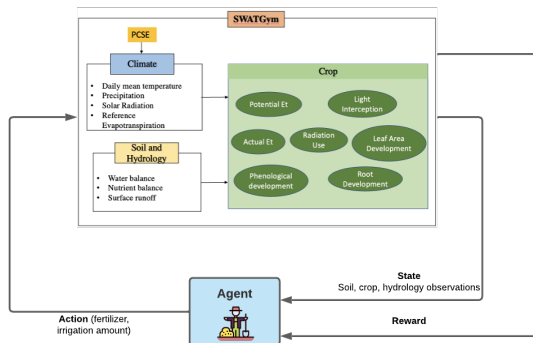


Figure 1: The SWATGym environment is the first SWAT-based RL environment. Agents interact with the environment to learn crop management strategies in simulation.

SWATGym is the first Python-based implementation of SWAT, which is written in FORTRAN and is not readily accessible to RL applications. Our environment is built on top of the OpenAI Gym framework, the gold standard framework for developing RL environments (Brockman et al. 2016). SWATGym primarily considers corn production and incorporates dynamical representations of corn yields and corresponding economic costs of input resources as well as their environmental impact. It has a continuous state space comprising of 14 state variables describing various processes related to weather, soil, crop, and hydrology dynamics (see Table 2). It also has a continuous multidimensional

action space. At time t , the action is given by $a_t = [F_t, I_t]$, where F and I represent fertilizer and irrigation amounts applied at that time. SWATGym can be run in either stochastic or deterministic modes. The former is enabled by default and is applied primarily through randomness added to climate inputs. The latter mode ensures the same actions lead to the same states.

Observation	Unit
Mean air temperature	°C
Precipitation	mm
Reference Evapotranspiration	mm
Solar Radiation	MJ/mm^2
Mean Vapor Pressure	hPa
Actual Evapotranspiration	mm
Water balance	mm
Daily runoff curve number	-
Leaf area index	-
Nitrogen Uptake	kg/ha
Denitrification	kg/ha
Nitrogen stress factor	-
Temperature stress factor	-
Water stress factor	-

Table 2: Environment’s observations include weather inputs, plant state variables, and soil state variables.

Reward Function

SWATGym produces a reward that characterizes the effect of different choices of actions on crop production. The reward at each time step is computed as

$$r_t = yld_t - \alpha F_t - \beta I_t, \quad (1)$$

where yld is the estimated crop yield on a particular day and α and β are penalty terms associated with the estimated cost of applying fertilizer F and irrigation I . In our case, $\alpha = 2.43$ and $\beta = 0.16$.

Crop Growth Simulation

SWATGym simulates crop growth for a full growing season (about 120 days for corn). Users have the option to specify the location and simulation start date as well as duration. The simulation starts with crop emergence and ends with the harvest. The environment operates on a daily time step. An episode begins on the specified/default simulation date or when the environment is reset and ends after the specified/default duration of the growing season (harvest day). The environment also has the option to save all relevant data

about the current growing season, including weather observations, crop states, soil and hydrology balances, actions taken, and yield achieved to date. This feature enables the collection of expert data, which can be used for other tasks such as offline learning.

State-space dynamics

SWATGym has several major modules such as hydrology, weather, soil temperature, crop growth, and agricultural management (Arnold et al. 1998). It includes state variables describing the time evolution of hydrological, soil, and crop variables. At time (day) t , the plant state variables are described by

$$z_p(t) = [LAI, BIO, E_a, N_{strs}, W_{strs}, T_{strs}], \quad (2)$$

and the soil state variables by

$$z_s(t) = [SW, RCN, DN, N_{up}]. \quad (3)$$

LAI is the leaf area index, BIO is the cumulative biomass, E_a is the actual evapotranspiration rate, and N_{strs} , W_{strs} , and T_{strs} are growth factors related to stress caused by nitrogen, water, and temperature on the plant. For soil state z_s , SW is the available soil water content, RCN is the daily surface runoff curve number, DN is the denitrification rate, and N_{up} is the nitrogen uptake.

We also define a vector of climatic inputs,

$$\xi(t) = [P, Et, Ta, Rd], \quad (4)$$

where P is the precipitation received on that day, Et is the reference evapotranspiration rate, Ta is average daily air temperature, Rd is daily solar radiation. Therefore, the plant state at time $t + 1$ is given explicitly as a nonlinear function of the state and the input climatic variables at time t :

$$z_p(t + 1) = f_p(z_p(t), z_s(t), \xi(t)), \quad t \in [0, T], \quad z_p(0) = x \quad (5)$$

where, for a given initial state x , $z_p(t)$ is the vector with the plant state variables at time t , $z_s(t)$ is the vector with the soil state variables; and $\xi(t)$ is the vector of the climatic inputs provided by PCSE (de Wit 2018) for a specific location and day.

Excluding daily weather data, most of the state variables are propagated by equations provided in the SWAT2009 Theory Documentation (Neitsch et al. 2011). Below we highlight a few of the crop-related state variables:

1. **Phenology:** Similar to the original SWAT model, we express crop growth/phenological development in terms of heat units, which are driven by daily mean temperature. Growth is accelerated at or above the optimal temperature for the crop, T_{opt} , and is slowed or stopped at or below the base temperature T_{base} . A crop's phenological development is based on daily heat unit accumulation, given by

$$HU_i = \bar{T}_i - T_{base}, \quad \forall i \in 1, \dots, T, \quad T > 1 \quad (6)$$

where HU is the value of heat units and \bar{T}_i is the average air temperature in $^{\circ}\text{C}$ on day i .

The fraction of potential heat units accumulated for a given day d is given by

$$fr_{PHU} = \frac{\sum_{i=1}^d HU}{PHU} \quad (7)$$

where PHU is the total number of heat units required for a plant to reach maturity. This is often calculated from the planting date to the harvest date (on the last day T) if not known beforehand.

Temperature is a key driving factor for crop growth in SWAT. Below we show a plot of observed daily mean temperatures during one simulation run. We also indicate the base and optimal temperatures for corn production. In

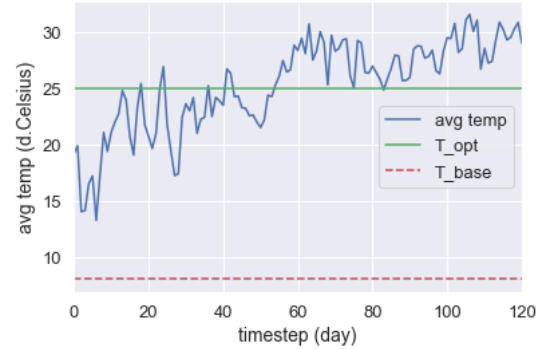


Figure 2: Temperature variation during one growing season, along with the crop-specific base temperature and optimal temperatures for corn production. $T_{base} = 8^{\circ}\text{C}$ and $T_{opt} = 25^{\circ}\text{C}$

this example, growth will be limited early on due to sub-optimal temperatures. However, other factors such as nitrogen and water balance may encourage growth despite temperature stress.

2. **Potential Growth:** Other factors related to plant growth that is modeled by SWATGym include leaf area development, light interception, and a plant-specific radiation use efficiency metric (which measures the conversion of intercepted light into biomass).

(a) *Leaf Area Development:* The leaf area index (LAI) is the area of green leaf per unit area of land. It is computed as a function of crop canopy height by

$$\Delta LAI_i = K_f \left(1 - e^{5*(LAI_{i-1} - LAI_{max})}\right) \quad (8)$$

where $K_f = LAI_{max}(fr_{LAI_{max},i} - fr_{LAI_{max},i-1})$ and $LAI_0 = 0$.

$$LAI_i = LAI_{i-1} + \Delta LAI_i \quad (9)$$

where h_c is the canopy height. For corn, $LAI_{max} = 3$ and $fr_{PHU, sen} = 0.9$ (Arnold et al. 2012).

(b) *Light Interception:* Using Beer's law, the amount of daily solar radiation intercepted by the plant is computed as

$$H_{phosyn} = 0.5H_{day}(1 - exp(-k_l LAI)) \quad (10)$$

where H_{photosyn} is the amount of intercepted photosynthetically active solar radiation (MJ/m^2), H_{day} is incident total solar, k_ℓ is light extinction coefficient and LAI is leaf area index.

- (c) **Biomass Production:** Radiation Use Efficiency (RUE) is defined for each plant species and is independent of the plant's growth stage. The potential increase in total plant biomass on a given day is given by

$$\Delta bio = RUE \cdot H_{\text{photosyn}} \quad (11)$$

The total plant biomass on a given day d is subsequently given by

$$bio = \sum_{i=1}^d \Delta bio_i, \quad d \leq T \quad (12)$$

3. **Crop Yield:** SWATGym computes crop yield as the product of a plant's above-ground biomass (and its roots if they are a harvest-able product) and harvest index, which is defined as the fraction of above-ground plant dry biomass removed as dry economic yield (with values typically between 0 and 1). HI , the potential harvest index for a given day in the plant's growing season is computed using the following relationship:

$$HI = HI_{\text{opt}} \frac{100 fr_{\text{PHU}}}{100 fr_{\text{PHU}} + \exp(11.1 - 10 fr_{\text{PHU}})}, \quad (13)$$

where HI_{opt} is the potential harvest index at the time of maturity, and fr_{PHU} is the fraction of potential heat units accumulated for the plant on a given day in the growing season. If needed, the actual harvest index can be derived from Equation 13 by taking water deficiency into account.




Figure 3: Potential harvest index for corn production during one growing season

Overall, the estimated yield is given by

$$yld = \begin{cases} bio_{\text{ag}} HI, & \text{for } HI \leq 1 \\ bio \left(\frac{HI}{HI+1} \right), & \text{otherwise,} \end{cases} \quad (14)$$

where yld is the crop yield (kg/ha), bio_{ag} is the above-ground biomass (kg/ha), HI is the harvest index, and bio

is the total plant biomass on the day of harvest. bio_{ag} is computed as follows,

$$bio_{\text{ag}} = (1 - fr_{\text{root}}) bio, \quad (15)$$

where $fr_{\text{root}} = 0.4 - 0.2 fr_{\text{PHU}}$ is the fraction of total biomass in the roots on harvest day and fr_{PHU} is the fraction of potential heat units accumulated for the plant on a given day in the growing season.

4. **Growth Constraints:** Plant growth may be affected by insufficient or excess water, nutrients, and extreme temperatures. Stress factors are typically 0 under normal conditions and approach 1 as growth conditions become suboptimal. Equations that describe the propagation of these stress factors are provided in (Neitsch et al. 2011).

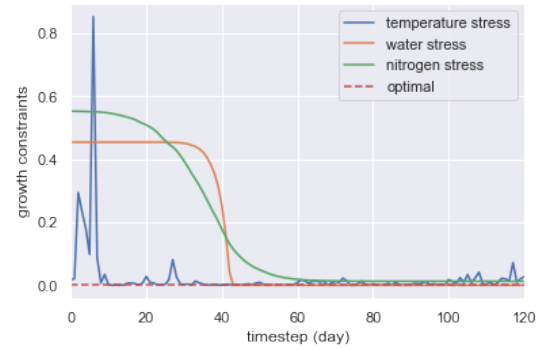


Figure 4: Stress factors during a growing season.

Source Code

SWATGym is released as a free and open-source environment under the terms of the Apache License 2.0. The code is publicly available on GitHub at <https://github.com/IBM/SWATgym>. While SWAT accounts for a variety of crop species and environmental processes, we focus on corn production in this paper. We further limit ourselves to key processes such as surface runoff, water balance, and denitrification. For a detailed description of the physiological processes, refer to (Neitsch et al. 2011).

We provide a simplified implementation of SWAT that makes the following assumptions:

- all climatic and agro-management inputs are applied uniformly and daily, over a single growing season (typically 120 days for corn).
- all soil layers (except the surface layer, top 10mm) have largely the same characteristics i.e. homogeneous soil profile.
- no/negligible percolation and bypass flow exiting the soil profile at the bottom; no lateral and base flow (which are typically included when computing soil water content).
- no growth-reducing factors such as weeds and pests and negligible growth impact of all other nutrients besides nitrogen.

Experiments

In addition to providing the SWATGym environment, we also evaluate a selection of crop management strategies on the environment to demonstrate that RL agents can learn useful crop management strategies. We benchmark the following agents:

- **DDPG.** Deep Deterministic Policy Gradient is a state-of-the-art RL algorithm for continuous control tasks (Lillicrap et al. 2015) that has been widely successful in data-rich applications.
- **TD3.** Twin Delayed Deep Deterministic policy gradient algorithm (Fujimoto, Hoof, and Meger 2018) builds on DDPG and applies various modifications to improve its stability and learning performance.

For both algorithms, we use the publicly available implementations provided in (Fujimoto, Hoof, and Meger 2018) along with the default hyperparameters. We also provide performance measures of three baseline agents:

- **Random Agent** - a dynamics-agnostic agent which selects random amounts of fertilizer and irrigation to apply at each time step.
- **Standard Practice Agent** - applies predetermined amounts of inputs on scheduled days during the early, mid, and late stages of the growing season. This corresponds to traditional farming methods of applying inputs during different phases of the crop’s growth.
- **Reactive Agent** - applies high concentrations of fertilizer and irrigation whenever nitrogen levels are depleted and the soil water content is below a certain threshold.

The implementation of the Standard and Reactive agents is based on the descriptions provided in (Overweg, Berghuijs, and Athanasiadis 2021).




Figure 5: Performance of baseline methods on one full growing season.

Evaluation

We measure the performance of the various agents on SWATGym. We conduct two sets of evaluations for each agent, (1) one that covers a full growing season (equivalent to one episode with 120 time steps/days) and (2) another set that is applied over multiple seasons spanning 200 episodes.

The latter is repeated 5 times, each with a different seed. In both cases, evaluations are performed every 7 days and the results averaged over 10 runs.

Single growing season. We train each agent for a full corn growing season, equivalent to 120 days. Figure 5 shows the outcomes of the three baseline agents. The standard approach obtains the best performance overall compared to that of the Random agent and Reactive agent, which is an expected outcome given this strategy is widely used in practice. This comparison also highlights one way in which SWATGym can be used to evaluate management strategies. With the Standard agent as the best baseline, we also evaluate its strategy against the RL-based strategies of the DDPG and TD3 agent (Figure 6).




Figure 6: Performance of best baseline method against RL methods on one full growing season

We observe that the DDPG algorithm performed better than both the Standard Agent and the TD3 agent on this task. This is interesting in itself because TD3 is an improvement of DDPG. We provide the learning curves of both algorithms in Figure 7.




Figure 7: Learning curves of the RL algorithms on SWAT-Gym. The shaded area represents half a standard deviation of the average evaluation over 10 runs of the experiment, each with a different seed.

Multiple growing seasons. We also compute the average return (\pm one standard deviation) across 5 runs of the exper-

iment, each with a different seed. Table 3 shows the maximum average return performance of the agents on SWATGym and Figure 8) shows the average reward for all agents across 5 repetitions of the experiment. In comparison to the baseline agents, the DDPG and TD3 results illustrate the potential of RL in revolutionizing crop management strategies.

Method	Performance
Random	-4547.152 ± 146.386
Standard	4396.402 ± 13.116
Reactive	279.121 ± 32.234
DDPG	4169.202 ± 2249.513
TD3	3344.010 ± 4519.612

Table 3: Maximum average return \pm one standard deviation for all agents across 5 repetitions of the experiment.




Figure 8: Average reward across 5 runs for all agents.

Discussion and Research Directions

Beyond highlighting the potential of RL strategies in facilitating sustainable crop management, our preliminary results demonstrated the value of SWATGym as a benchmark framework. In this section, we outline key challenges in modeling crop growth and briefly describe social challenges such as the adoption of RL-based crop management strategies.

Challenges in crop modeling. Accurately modeling crop growth is a key and active research problem. SWATGym has lots of room for improvement to increase its fidelity. Yet, one must also balance simplicity and generalizability to different regions, climates, crop types, and soil characteristics. Future releases of the framework will extend the dynamics to include more processes such as groundwater seepage and other applications such as offline RL.

Seasonal spatial-temporal variability. Most crops that are simulated are short-season crops. Regardless of the season duration, it is important to capture the spatial-temporal variability of the season. For example, evapotranspiration is higher in the summer than in winter, so the crop management algorithm should apply more irrigation to match water loss during that period. Likewise, surface runoff could be

higher in the spring or rainy season (e.g., when the snow melts/after it rains). As such, the best strategy might be to apply less fertilizer to avoid leaching and waste or to apply it early in the growing season before runoff affects it.

Reality Gap. SWATGym simulates crop growth daily and allows agents to select inputs on any given day. In practice, such operations are done over a week or more, depending on the size of the field and the type of equipment used to irrigate or fertilize. Furthermore, the environment can be further constrained to limit the total amounts of inputs applied throughout the growing season and terminate whenever this threshold is reached. RL developments must be done in conjunction with domain experts to ensure the applicability of ML recommendations.

Social Impact. Through IBM’s Sustainability Accelerator program, a pro bono social impact program that leverages technologies such as hybrid cloud and AI to enhance and scale non-profit and government organizations, plans are underway to deploy this framework in small-holding farms in arid regions of the United States. The goal is to help farmers make better decisions about crop management and offer a platform where they can easily evaluate different strategies to optimize crop production. This highlights a challenging aspect of this work, RL, and digital solutions to real-world problems: adoption. Working with small-holding farmers will present unique opportunities to learn how to translate research ideas into useful real-world products.

Climate Change. The negative impacts of climate change are already being felt, in the form of increasing temperatures, weather variability, shifting agroecosystem boundaries, invasive crops, and pests, and more frequent extreme weather events (Calzadilla et al. 2013). ML-informed decision-making must be responsive to both shifting spatial and temporal patterns that influence decision-making, as also the increased frequency of extreme weather events. Approaches that improve robustness to these events and allow farmers to better understand the implications of different actions in a more volatile environmental system are critical to improving adaptation to climate change

Conclusion

Evaluating and comparing different crop management strategies in the real world is a costly and time-consuming task. Simulated environments offer a more efficient solution, allowing for the simultaneous benchmarking of multiple strategies with minimal cost. To this end, we introduced SWATGym, a reinforcement learning environment designed to make it easy to simulate crop growth and evaluate crop management strategies. SWATGym models crop growth processes from emergence to harvest and can be used to evaluate crop management strategies, which can inform decision-makers and promote sustainable agriculture practices. We hope that SWATGym will facilitate follow-up work and encourage collaboration between researchers and practitioners in both reinforcement learning and agriculture to develop better crop management strategies and promote sustainable agriculture.

References

- Alexandratos, N.; and Bruinsma, J. 2012. World agriculture towards 2030/2050: the 2012 revision. *ESA Working Papers 12-03, Food and Agriculture Organization of the United Nations (FAO)*.
- Arnold, J.; Kiniry, J.; Srinivasan, R.; Williams, J.; Haney, E.; and Neitsch, S. 2012. Soil and Water Assessment Tool Input/Output file documentation version 2012. Technical report, Texas Water Resources Institute.
- Arnold, J. G.; Srinivasan, R.; Muttiah, R. S.; and Williams, J. R. 1998. Large area hydrologic modeling and assessment part I: model development 1. *JAWRA Journal of the American Water Resources Association*, 34(1): 73–89.
- Ashcraft, C.; and Karra, K. 2021. Machine Learning aided Crop Yield Optimization. *arXiv preprint arXiv:2111.00963*.
- Binas, J.; Luginbuehl, L.; and Bengio, Y. 2019. Reinforcement learning for sustainable agriculture. In *ICML 2019 Workshop Climate Change: How Can AI Help*.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Calzadilla, A.; Rehdanz, K.; Betts, R.; Falloon, P.; Wiltshire, A.; and Tol, R. S. 2013. Climate change impacts on global agriculture. *Climatic change*, 120(1): 357–374.
- Chen, M.; Cui, Y.; Wang, X.; Xie, H.; Liu, F.; Luo, T.; Zheng, S.; and Luo, Y. 2021. A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. *Agricultural Water Management*, 250: 106838.
- Chlingaryan, A.; Sukkarieh, S.; and Whelan, B. 2018. Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review. *Computers and electronics in agriculture*, 151: 61–69.
- de Wit, A. 2018. PCSE: Python Crop Simulation Environment, <http://pcse.readthedocs.io>.
- Dutia, S. 2014. Agtech: Challenges and opportunities for sustainable growth. Available at SSRN 2431316.
- Elavarasan, D.; and Vincent, P. D. 2020. Crop yield prediction using deep reinforcement learning model for sustainable agrarian applications. *IEEE access*, 8: 86886–86901.
- Follett, R. F.; and Hatfield, J. L. 2001. Nitrogen in the environment: sources, problems, and management. *TheScientificWorldJournal*, 1: 920–926.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, 1587–1596. PMLR.
- Gandhi, R. 2022. Deep Reinforcement Learning for Agriculture: Principles and Use Cases. In *Data Science in Agriculture and Natural Resource Management*, 75–94. Springer.
- Gassman, P. W.; Reyes, M. R.; Green, C. H.; and Arnold, J. G. 2007. The soil and water assessment tool: historical development, applications, and future research directions. *Transactions of the ASABE*, 50(4): 1211–1250.
- Liakos, K. G.; Busato, P.; Moshou, D.; Pearson, S.; and Bochtis, D. 2018. Machine learning in agriculture: A review. *Sensors*, 18(8): 2674.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Muller, A.; Schader, C.; El-Hage Scialabba, N.; Brüggemann, J.; Isensee, A.; Erb, K.-H.; Smith, P.; Klocke, P.; Leiber, F.; Stolze, M.; et al. 2017. Strategies for feeding the world more sustainably with organic agriculture. *Nature communications*, 8(1): 1–13.
- Neitsch, S. L.; Arnold, J. G.; Kiniry, J. R.; and Williams, J. R. 2011. Soil and Water Assessment Tool theoretical documentation version 2009. Technical report, Texas Water Resources Institute.
- Overweg, H.; Berghuijs, H. N.; and Athanasiadis, I. N. 2021. CropGym: a Reinforcement Learning Environment for Crop Management. *arXiv preprint arXiv:2104.04326*.
- Ranganathan, J.; Waite, R.; Searchinger, T.; and Hanson, C. 2018. How to sustainably feed 10 billion people by 2050, in 21 charts. *Insights, World Resources Institute*.
- Rimski-Korsakov, H.; Rubio, G.; and Lavado, R. S. 2009. Effect of water stress in maize crop production and nitrogen fertilizer fate. *Journal of plant nutrition*, 32(4): 565–578.
- Sharpley, A. N.; Smith, S.; and Naney, J. 1987. Environmental impact of agricultural nitrogen and phosphorus use. *Journal of Agricultural and Food Chemistry*, 35(5): 812–817.
- Shibu, M.; Leffelaar, P.; Van Keulen, H.; and Aggarwal, P. 2010. LINTUL3, a simulation model for nitrogen-limited situations: Application to rice. *European Journal of Agronomy*, 32(4): 255–271.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Van Diepen, C. v.; Wolf, J. v.; Van Keulen, H.; and Rappoldt, C. 1989. WOFOST: a simulation model of crop production. *Soil use and management*, 5(1): 16–24.
- Williams, J.; Jones, C.; Kiniry, J.; and Spanel, D. A. 1989. The EPIC crop growth model. *Transactions of the ASAE*, 32(2): 497–0511.
- Wu, J.; Zhao, P.; Tao, R.; Hovakimyan, N.; Marcillo, G.; Martin, N.; Ferreira, C.; Kalantari, Z.; and Hobbs, J. 2021. Optimization of Agricultural Management for Soil Carbon Sequestration based on Deep Reinforcement Learning and Large-Scale Simulations. In *NeurIPS 2021 Workshop on Tackling Climate Change with Machine Learning*.
- Zhao, C.; Liu, B.; Xiao, L.; Hoogenboom, G.; Boote, K. J.; Kassie, B. T.; Pavan, W.; Shelia, V.; Kim, K. S.; Hernandez-Ochoa, I. M.; et al. 2019. A SIMPLE crop model. *European Journal of Agronomy*, 104: 97–106.