



Analyzing Reinforcement Learning Algorithms for Nitrogen Fertilizer Management in Simulated Crop Growth

Michael Vogt, Benjamin Rosman

School of Computer Science and Applied Mathematics
University of the Witwatersrand
Johannesburg, South Africa
mike97vogt@gmail.com, benjamin.rosman1@wits.ac.za

Abstract

Establishing intelligent crop management techniques for preserving the soil, while providing next-generational food supply for an increasing population is critical. Nitrogen fertilizer is used in current farming practice as a way of encouraging crop development; however, its excessive use is found to have disastrous and long-lasting effects on the environment. This can be reduced through the optimization of fertilizer application strategies. In this work, we apply a set of reinforcement learning algorithms – the DQN, Double DQN, Dueling DDQN, and PPO – to learn novel strategies for reducing this application in a simulated crop growth setting. We provide an analysis of each agent’s ability and show that the Dueling DDQN agent can learn favourable strategies for minimizing nitrogen fertilizer application amounts, while maintaining a sufficient yield comparable to standard farming practice.

1 Introduction

The global need for innovative crop management techniques for a changing climate and growing population is paramount. Agricultural sectors now need to produce more food on less arable land, while also keeping the environmental impact to a minimum. The use of nitrogen fertilizer is increasingly used as a method for encouraging crop development, producing more yield while sustaining crop quality as the earth’s resources are depleted. This rapid adoption has however led to nitrogen fertilizers being excessively applied beyond the crops demand, resulting in unfortunate consequences such as soil, water, and air degradation. The long-term application of nitrogen fertilizer increases the acidity in the soil and may render the soil infertile, causing crops to no longer respond to its application. Beyond this, groundwater contamination may cause severe health hazards for humans and livestock alike, while nitrous oxide emissions (N_2O) are considered 300 times more harmful than carbon dioxide emissions (CO_2) in terms of their potential for global warming [9].

There is a subsequent need for establishing next-generation, sustainable application techniques for reducing this impact, as this has the potential to affect humanity on a global scale.

The importance of this lies in the present unknown knowledge of whether current standard farming practice techniques are indeed optimal, or whether they will transfer to new scenarios at all [2]. Performing and evaluating new strategies in practice, however, can be challenging due to the long nature of crop development and the additional resource and labour costs involved.

There has been a significant impact of reinforcement learning applied to many domains with great success, however its application to agriculture has not been adequately researched, with its so far unrecognized potential still to be explored [2] [8]. This research utilizes model-free reinforcement learning techniques, namely the Deep Q-Network (DQN), Double Deep Q-Network (Double DQN), Dueling Double Deep Q-Network (Dueling DDQN), and Proximal Policy Optimization (PPO) algorithms to explore novel nitrogen fertilizer application approaches to this problem in a simulated crop growth setting. We show that the Dueling DDQN agent is in fact able to learn comparable strategies to standard farming practice techniques by reducing the amount of fertilizer applied, thereby limiting environmental impacts while still producing a sufficient yield. The main contributions of this research provide an analysis of the ability of each of these reinforcement learning methods to efficiently manage fertilizer application amounts. A subsequent contribution is providing an initial benchmark of model-free reinforcement learning methods to be extended upon within this domain.

2 Background

Reinforcement learning is a class of methods in which an agent learns, through trial-and-error interaction with its environment, how to behave to maximize a numerical reward signal [11]. Markov Decision Processes (MDP), represented as the tuple (S, A, R, P) , are a way to formalize sequential decision-making processes and provide a simple mathematical framework that defines the interaction between an agent and its environment in terms of its states S , actions A , and rewards $R(s, a)$. $P(s, a, s')$ are the transition dynamics and is the probability that an agent in a state s takes an action a and progresses to state s' . A policy $\pi(a|s)$ is a stochastic rule – mapping states to actions – used by the agent to make decisions to maximize its discounted cumulative future rewards [11]. A policy’s state and action-value functions,

$$v_{\pi}(s) = \mathbb{E}_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s] \quad (1)$$

$$q_{\pi}(s, a) = \mathbb{E}_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s, A_t = a] \quad (2)$$

assign to each state and state-action pair respectively, the expected cumulative discounted return. Here, γ , is the discount factor and is used to instill caution in the agent, due to the uncertainty within the environment and the agent’s ability to gain rewards in the future [11].

It is difficult to draw an accurate taxonomy of reinforcement learning algorithms, however they broadly fall into the groups of model-based and model-free reinforcement learning. Model-free reinforcement learning methods can be further sub-divided into value-based, policy-based, and actor-critic methods. Value-based methods approximate the state or state-action value, from which the optimal policy is then derived. This is known as an off-policy method and differs from policy-based methods, in which the policy is directly approximated and improved. Actor-critic methods are a temporal-difference approach to policy-based methods. Here the learning of the actor is still based on directly optimizing the policy, but a critic is introduced to evaluate the choice of action and provide guided learning [11].

Reinforcement learning has been combined with non-linear function approximators such as neural networks or decision trees, allowing them to learn complex feature representations. The

DQN algorithm is one such algorithm that has shown great success in playing games such as those in the Atari 2600 domain [7].

Recently, within agriculture, the DQN algorithm has been used to determine the optimal irrigation strategy for rice based on short term weather forecasts [3]. It has further shown proven success for the intelligent control of agricultural irrigation in a greenhouse plantation in Hunan [14].

2.1 Deep Q-Network

Q-learning combined with a multi-layered neural network forms the basis of the DQN algorithm and this is used to approximate the optimal action-values, which may be recursively defined as

$$Q_{\pi}^*(s, a) = \sum_{s', r} P(s', r | s, a) [r + \gamma \max_{a'} Q^*(s', a')] \quad (3)$$

Reinforcement learning is known to be unstable when combined with a non-linear function approximator, often resulting in divergence. The cause of this divergence may be due to several factors, such as correlations in sequential observations present in reinforcement learning. Small updates to the non-linear function approximators may cause changes in the policy, thus changing the distribution of the data. Furthermore, within the DQN algorithm, correlations between the action-values and the target values, $r + \gamma \max_{a'} Q(s', a')$, may also cause divergence [7]. To overcome this difficulty, the DQN algorithm has been combined with an experience replay buffer in which experiences of (s, a, r, s') are stored. During training, Q-learning updates are applied to a sample of mini-batches of this experience, which are drawn uniformly at random from the replay buffer. This breaks the correlations experienced by a sequence of observations, resulting in independently and identically distributed data – which assists in convergence. Another issue involved is that the target in the action-value function update is an approximation itself. If this is updated at every iteration, the result is unstable learning as we are moving towards something that is itself an estimate and non-stationary. To overcome this, a separate Q-network is initialized with fixed parameters, which are only periodically updated [7]. The pseudocode for the DQN algorithm is shown in Algorithm 1.

Algorithm 1 Deep Q-Network with Experience Replay [7]

```

Initialize the replay buffer  $\mathcal{D}$  with capacity  $N$ 
Initialize action-value function  $Q$  with random weights  $\theta$ 
Initialize target action-value function  $Q$  with weights  $\theta^- = \theta$ 
for each episode do
  for  $t = 1, T$  do
    Select a random action  $a_t$  with probability  $\epsilon$ 
    Otherwise, select  $a_t = \max_a Q(s, a; \theta)$ 
    Execute action  $a_t$  and observe reward  $r_t$  and next state  $s_{t+1}$ 
    Store transitions  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
    Sample random mini-batches of  $(s_j, a_j, r_j, s_{j+1})$  from  $\mathcal{D}$ 
    Set  $y_j = \begin{cases} r_{j+1}, & \text{if } s_{j+1} \text{ is terminal} \\ r_{j+1} + \gamma \max_{a'} Q(s_{j+1}, a'; \theta^-), & \text{otherwise} \end{cases}$ 
    Perform stochastic gradient descent step on
     $L_{\theta} = (y_j - Q(s_j, a_j; \theta))^2$ 
    Update  $Q_{\theta^-} = Q_{\theta}$  after every  $K$  steps
  end for
end for

```

2.2 Double Deep Q-Network

One lingering problem with DQN is that it suffers from overestimated action-values. Van Hasselt et al. [12] demonstrated that this overestimation causes sub-optimal policies in practice. To overcome this overestimation the authors proposed the Double DQN algorithm, which they show leads to more accurate action-value estimates, as well as better policies. In the standard DQN algorithm, it is indeed the max operator which leads to overoptimistic value estimates, as this tends to prefer overestimated to underestimated action-values. This is because the same parameter values are used both to select and to evaluate the chosen action. Their proposed Double DQN algorithm therefore decomposes the max operator in the target into action selection and action evaluation. Here, the selection of the action is still due to the online weights of the behaviour policy, but the parameters of the target network are used to evaluate the value of this policy. The new target update leading to better estimated action-values is

$$y^{DDQN} = r + \gamma Q(s', \max_a Q(s', a; \theta); \theta^-) \quad (4)$$

2.3 Dueling Architectures

The combination of reinforcement learning with neural networks generally uses conventional deep learning network architectures such as convolutional neural networks or LSTMs. Although showing great success, they may not be optimally suited for reinforcement learning tasks [13]. In order to find architectures more suited to model-free reinforcement learning tasks specifically, Wang et al. [13] proposed the dueling network architecture. The dueling network is composed of a single deep model with two streams for two separate estimators – the value function $V(s; \theta)$, and advantage function $A(s, a; \theta)$. The output of these two streams is then combined to produce the action-value estimate $Q(s, a; \theta)$. The value function is a measure of how good it is to be in a particular state, while the action-value function is a measure of how valuable choosing a particular action is when in this state. The advantage function is a relative measure of the importance of each action, and relates to the value and action-value functions by

$$A_\pi(s, a) = Q_\pi(s, a) - V_\pi(s) \quad (5)$$

In Wang et al. [13] they explain that simply using the definition of the advantage function to produce the action-value estimate is not good idea, as the value and advantage functions cannot be uniquely recovered. They further state that poor performance is achieved when using this expression, emulating this lack of identifiability in practice. During learning, action-value estimates can be produced in the more stable manner proposed, namely

$$Q(s, a) = V(s) + A(s, a) - \frac{1}{|A|} \sum_{a'} A(s, a') \quad (6)$$

The main benefit of the dueling network, as claimed by Wang et al. [13], is that it can generalize learning across its actions without having to change the underlying architecture of the algorithm, thus making it suitable and easy to implement with a variety of model-free reinforcement learning methods.

2.4 Proximal Policy Optimization

The Proximal Policy Optimization algorithm is a policy-based method with an actor-critic style. The motivation behind the design of this algorithm is to maintain the data efficiency and reliability in performance of the Trust Region Policy Optimization algorithm, while improving

sample efficiency by only using first-order optimization methods [10]. One of the main contributions of this method is introducing a new surrogate objective function, with clipped probability ratios, that enables multiple epochs of mini-batch updates without having destructively large policy updates. This objective function has the form

$$L^{CLIP}(\theta) = \mathbb{E}_t[\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)], \quad (7)$$

where $r_t(\theta)$ is the ratio of the log probability of the new policy with the log probability of the old policy.

3 Methodology

Reinforcement learning applied to agriculture is a newly emerging field of interest. Despite its infancy, this has resulted in varying applications and approaches to smart farming such as yield prediction, crop quality monitoring, and weed detection [5]. Despite these successes, the path to sustainable agriculture requires more fine-grained control of the development of the crop itself, in terms of its surrounding environmental parameters [2]. The two focus areas for recent advances in crop management using reinforcement learning are finding optimal strategies for irrigation and fertilizer management. This is because of their large effect on crop growth, as well as water being a limited resource, with fertilizers having a negative effect on soil health and the environment [9]. With more recent focus on learning irrigation strategies, fertilizer management in crop growth has received the least attention, despite its many negative effects. This motivates our investigation into the ability of model-free reinforcement learning to learn new strategies that may acquire similar yields, while limiting the disastrous effects of nitrogen fertilizer. In order to achieve this aim the simulated crop growth environment that Overweg et al. [8] developed is used, termed CropGym¹.

It is a nitrogen-limited crop growth environment in which the underlying crop growth processes involved are modelled using the LINTUL-3 package, provided in the Python Crop Simulation Environment (PCSE)². The parameters of the crop growth model have been set to simulate a winter wheat crop. The state space consists of multidimensional observations of the current state of the crop as output by the LINTUL-3 package, as well as weather data over the past week. The action space consists of discrete doses of nitrogen fertilizer, which may be applied by an agent at intervals of one week. While providing a strong restriction on the limit of nitrogen supply in terms of the quantity observed in the soil, Overweg et al. [8] do not restrict the amount of times fertilizer is allowed to be applied. This allows for the exploration of strategies outside of standard farming practice approaches thus far developed by experts. This is significant, as standard practice knowledge may not transfer well to new scenarios caused by climate change [2].

To contrast the ability of model-free reinforcement learning algorithms to efficiently learn fertilizer management strategies, two baseline agents were implemented. The first baseline is termed the standard practice agent, and mimics current farming standard practice approaches to applying nitrogen fertilizer. This agent applies three discrete doses of nitrogen fertilizer throughout the crop growth season, which spans 8 months – starting on January 1st and ending on August 1st. Together with this, a reactive baseline agent was implemented which applies a high concentration of fertilizer every time the nitrogen content in the soil is depleted.

Along with these baselines, the DQN agent was implemented, as it has been successfully used within the smart farming domain for irrigation-based decision making [14] [3]. The Double

¹<https://github.com/BigDataWUR/crop-gym>

²<https://github.com/ajwdewit/pcse>

DQN agent was also implemented, as it has been shown to achieve more reliable strategies, with more accurate action-value estimates [12]. To investigate the potential performance gains of dueling architectures, and due to the overestimation caused by the DQN, the dueling network architecture was implemented together with the Double DQN algorithm. Thus, the third value-based method implemented was the Dueling DDQN agent. In order to fairly compare and investigate the effectiveness of value-based methods for fertilizer management to previously researched work, the PPO agent was implemented. Overweg et al. [8] implement this agent, and it is the first used for investigating the ability of reinforcement learning to find comparable fertilizer application strategies in this domain. Through the evaluation of these agents and by comparing them to one another along with the baselines, we provide a full analysis of their ability, as well as a more encompassing initial benchmark for model-free reinforcement learning methods to be built upon in the future.

4 Experiments

In order to analyze the performance of the implemented reinforcement learning agents, this section outlines the specifics of the environment, the characteristics of the baseline fertilizer application strategies, the details corresponding to the implementation of the reinforcement learning agents, and the metrics used to evaluate their performance.

4.1 Environment

The crop growth environment is created as a fully observable MDP, meaning the agent can make informed decisions based on all observable state variables. At each time-step during training, the agent observes state variables that are output by the LINTUL-3 winter wheat crop model. This model is calibrated using initial soil, site, and crop conditions, a pre-defined agromanagement strategy, and weather data from the past week. The output of the crop growth model importantly includes observations such as the yield, referred to as the weight of the storage organ (WSO) in LINTUL-3, the total rainfall experienced (TRAIN), and the total nitrogen uptake by the crop (NUPTT). The agent is allowed to apply one of seven discrete doses every week. Included in this action space is the ability to apply no nitrogen fertilizer at all, as this is desirable due to the costs involved and the negative long-term effects of nitrogen on the soil, water, and air. The action space is related to the amount of fertilizer to be applied, by the expression

$$A = \left\{ 20a \frac{\text{kg}}{\text{ha}} \mid a \in \{0, 1, 2, 3, 4, 5, 6\} \right\} \quad (8)$$

The reward function has been modelled in such a way as to punish the agent for excessive fertilizer application amounts, while rewarding it for a large grain yield at the end of the season. It is defined to be the difference between the weight of the storage organ (w_{SO}), achieved with the agent’s fertilizer application strategy, and the weight of the storage organ (w_{SO}^*) achieved if no fertilizer were to be applied. From this, a negative coefficient β , multiplied by the weight of fertilizer (w_{fert}) the agent applies is subtracted. The reward function is

$$r_t = w_{SO,t} - w_{SO,t-1} - (w_{SO,t}^* - w_{SO,t-1}^*) - \beta w_{fert,t} \quad (9)$$

Here, $\beta = 10$ is set to place a strong restriction on the amount of fertilizer an agent would like to apply, thereby ensuring the learning of strategies with minimal application.

This constitutes a complete description of the environment details. No modifications to this environment created by Overweg et al. [8] have been made.

4.2 Agents

The baseline agents are implemented for comparison against the reinforcement learning agents. The standard practice agent applies three doses of $60\text{kg}/\text{ha}$ nitrogen fertilizer. These are applied in the first and third week of March, with the final dose being applied in the first week of May. The reactive baseline agent applies $120\text{kg}/\text{ha}$ every time the soil nitrogen content depletes below $5\text{kg}/\text{ha}$.

Comparing to these baselines are the DQN agent, the Double DQN agent, the Dueling DDQN agent, and the PPO agent. The PPO agent is implemented using the Stable-Baselines3³ package, for which the policy and value function networks are constructed using a multi-layered perceptron (MLP) of two fully connected layers. The activation functions following the output of each layer are tanh non-linearities. The policy network is mapped to the dimension of the action space, after which a SoftMax function is used to assign the probabilities of selecting each available action. Similarly, the Q-network and target network for the DQN, Double DQN, and Dueling DDQN use a basic MLP with two fully connected layers. Following the output of these layers, these agents use a rectified non-linearity.

This simple approach to the neural network design is taken to maintain consistency to investigate the performance of the reinforcement learning agents themselves within the fertilizer application domain. The Dueling DDQN includes two streams for the value and advantage function, as outputs for the network. These streams are not combined via an aggregating layer in the network to produce the action-value estimate Q , as is done by Wang et al. [13], as our agent selects its actions according to the advantage estimate, since the value estimate is simply a scalar valued output and has no relative effect on the advantage of each action.

Many variational improvements exist in terms of more advanced underlying network architectures. These can be easily implemented to further improve the performance following this benchmark analysis.

4.3 Training Details

The reinforcement learning agents were all trained to convergence on the CropGym environment using the calibrated winter wheat crop model and 25 years' worth of weather data from the Netherlands during the period of 1983 to 2016, as can be seen in Figure 1. This environment was normalized with the VecNormalize wrapper to further assist with stable learning. For this, the observations were normalized with an observation clipping of 10, while the rewards remained unchanged.

The fact that the largest growth takes place towards the end of the crop growth season, is why the discount factor for all future rewards provided by the environment, γ , is set to 1 for all agents.

For all the agents, the dimensions of the hidden layers and tunable hyper-parameters are found using the highly scalable Ray[Tune] optimization framework [6]. For this, the Tree-Parzen Estimator hyper-parameter optimization algorithm was used [1].

4.4 Evaluation Procedure

All trained agents were evaluated on a testbed containing the 1984, 1994, 2004, 2014, 2019, and 2020 crop growth seasons, which were naturally excluded from the training set. For each of these testbed years, weather data was taken from $52^\circ N, 52.5^\circ E$ to maintain consistency with Overweg et al. [8]. This range helps provide a full analysis of the agent's ability to

³<https://github.com/DLR-RM/stable-baselines3>

generalize across different decades with varying conditions. The performance of each agent was evaluated in terms of the total reward obtained, yield produced, crop nitrogen uptake, and total fertilizer application amounts. To further test each agent’s ability to generalize, we use agents trained on five random seeds, choosing the best three of each, and evaluate their performance on a fourth random seed for the mentioned evaluation metrics. This is done to maintain generalizability, while ensuring that one destructive policy of one of the seeds does not overshadow the performance of the agents.

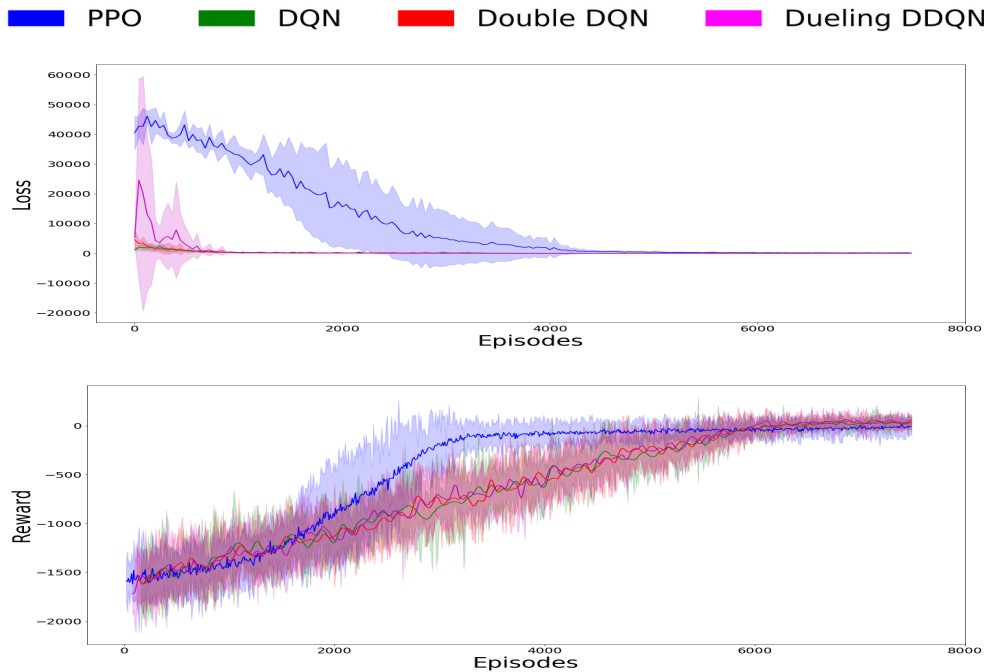


Figure 1: Loss and Reward Curves for all Implemented Agents Averaged over 5 Random Seeds.

5 Results

This section presents the results for the evaluation metrics discussed in the previous section for the testbed years. In doing so, and by further discussing the findings, a full analysis of the ability of some model-free reinforcement learning algorithms applied to the fertilizer crop management domain is provided.

Year	TNSOIL	NUPTT	WSO
1984	1.95	21.1	3.44
1994	1.46	20.9	4.99
2004	1.45	21.1	4.64
2014	1.10	19.6	6.62
2019	1.30	19.8	4.87
2020	1.12	19.7	4.12

Table 1: Crop Growth Variables without any Fertilizer Application.

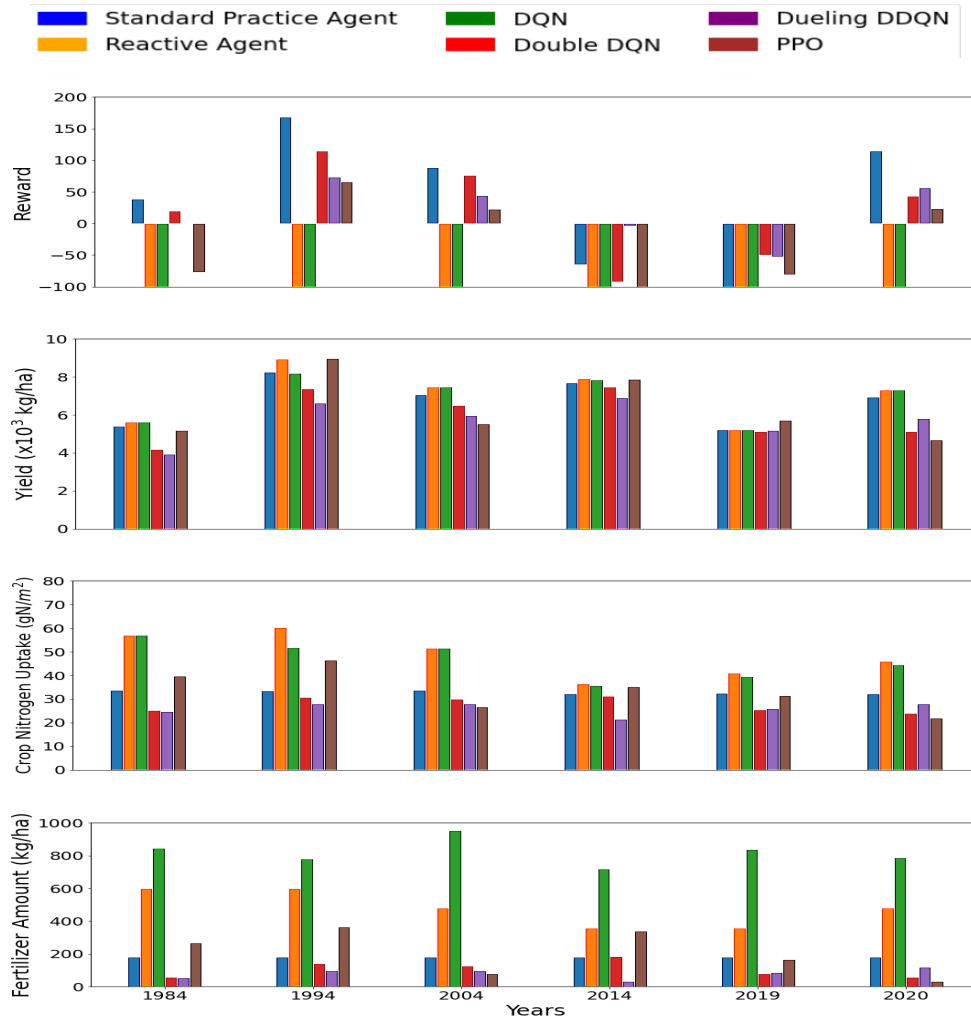


Figure 2: The Performance of each Agent in terms of Reward Obtained, Yield Produced, Crop Nitrogen Uptake, and Fertilizer Application Amount; Averaged over 3 Random Seeds.

As previously discussed in Section 1, the increasingly excessive use of nitrogen fertilizers to maintain yield is causing harm to all facets of the environment in terms of soil, water, and air degradation [9]. There is an apparent need for finding strategies that minimize fertilizer application amounts, while still providing a sufficient yield. The reward function has been modelled in such a way as to enforce this, however the rewards obtained cannot be viewed solely as the agent’s ability to learn optimal strategies. Furthermore, the reward obtained, yield produced, crop nitrogen uptake, and fertilizer application amounts should be analyzed individually per year, but the evaluation of the performance of each agent should take all testbed years into account, with the focus on the latter three evaluation metrics. We can see

this is necessary by analyzing Table 1, which represents some of the important crop growth variables when no fertilizer is applied throughout the crop growth season, and comparing it to the reward plot in Figure 2. Despite having some of the smallest values for the soil nitrogen content (TNSOIL) and the nitrogen uptake by the crop (NUPTT) for the years 2014 and 2019, the yield produced (WSO) in these years are the largest and third largest of all the testbed years. Naturally, nitrogen is not the only factor inducing crop growth, and many other factors such as irrigation, rainfall, and other fertilizers present in the soil are potentially playing a role in this. This growth does however have a relative effect on the reward for those years, due to the modelling of the reward function.

From Figure 2 it can be seen that there is a strong correlation between the amount of fertilizer applied, and the subsequent uptake of nitrogen by the crop. Despite this, there appears to still be a limit as to the amount of yield produced, evident in the year 2019. This is further evidenced when analyzing the DQN agent and the reactive agent. They apply extreme amounts of fertilizer throughout all of the testbed years and subsequently have the largest uptake of nitrogen by the crop, yet only achieve marginally better yields in the years 1984, 2004, 2014, and 2020. Due to these factors, their reward obtained is extremely low, and their application strategies are incompatible with the aims of this research, further showing the need for intelligent application strategies.

Comparing the PPO agent to the DQN and reactive agents, we can see that it successfully learns to apply less nitrogen fertilizer in the years 1984, 1994, 2014, and 2019, yet still achieves comparable yields, even surpassing the yield produced in 2019. The PPO agent can also be seen to apply the least amount of fertilizer in 2004 and 2020, and subsequently achieves the least yield in these years. When comparing it to the current baseline benchmark in farming, the standard practice agent, the PPO agent applied a fair amount more nitrogen in 1984, 1994, and 2004, with only slightly better yields. The PPO agent is therefore deemed to learn better policies than the DQN agent and reactive agent, however, does not optimally generalize across all the testbed years for favourable results against the standard practice agent.

The standard practice agent applies the same fixed amount of nitrogen fertilizer for every testbed year, thus allowing it to receive the best reward in four of the six years as the remainder of the reward is largely due to the growth incurred after this application. Despite this being a good strategy in general, for the years 2014 and 2019 it appears that little nitrogen fertilizer is needed for maximum growth. This shows the inability of this strategy to adapt to changing climate conditions. For this the Double DQN agent and the Dueling DDQN agent show favourable results.

The Double DQN agent learns an application strategy in which it applies less fertilizer than the standard practice agent in every single testbed season. It subsequently achieves less yield in every year, however only marginally in the years 1994, 2004, 2014, and 2019, with a fair amount less fertilizer applied in 2004 and 2019. This shows the Double DQNs ability to substantially reduce fertilizer application amounts, while still achieving a sufficient yield for many of the years. When comparing the Double DQN and DQN agents, it can evidently be seen that the overestimated action-values result in sub-optimal policies when using the DQN algorithm.

The Dueling DDQN agent, although taking a cautious approach to fertilizer application, shows the best generalized application strategy across the decades represented by the testbed years. This can be particularly seen in 2014 and 2020. When looking at the reward obtained, yield produced, and fertilizer application amount for 2014, and comparing to all other agents, no other agent was able to learn that this year required truly little fertilizer to still achieve a sufficient yield. It applies far less fertilizer, achieves comparable yield, and subsequently has the best reward, while all other agents achieve a terribly negative reward due to their fertilizer

application amounts. This is impressive when analyzing the 2020 crop growth season as well. For 2020, contrary to the other reinforcement learning agents, the Dueling DDQN agent has learnt that more nitrogen fertilizer is required to achieve a sufficient yield. This shows the agent's ability to generalize learning across its actions for different decades, with vastly varying conditions.

The Dueling DDQN agent is concluded to learn the best generalizable fertilizer application strategies. It takes a far more cautious approach than all its competitors, however it can learn favourable strategies for the extreme cases of 2014 and 2020.

6 Related Work

Within the domain of crop management using reinforcement learning, recent advances have focused on irrigation and fertilizer management, with more focus on irrigation management.

Within the fertilizer management domain, Garcia [4] provides some of the earliest work, building a reinforcement learning model to achieve a satisfying crop yield, while maintaining below the prescribed limit of nitrogen found in drinking water. However, this implementation only used R-learning to predict the three best fertilizer application dates and is limited in its approach and is now outdated.

After this there was limited interest in the use of reinforcement learning for fertilizer crop management. However, the combination of deep learning with reinforcement learning spurred increased interest in its applications. This allowed Overweg et al. [8] to provide the CropGym framework, with a focus on nitrogen fertilizer management, for the exploration of the application of reinforcement learning to sustainable agriculture. They, however, only implement a PPO agent without any hyper-parameter tuning. Our research builds on their initial analysis by providing a more encompassing variety of model-free reinforcement learning algorithms applied to this domain.

Prior to our research, Overweg et al. [8] and Chen et al. [3] have been the most recent successful implementations using reinforcement learning for fertilizer and irrigation management, respectively. To our knowledge, there are no publications that have yet implemented a combination of these crop development factors.

7 Conclusion

The aim of this research was to provide a comprehensive analysis and an initial benchmark for the ability of some model-free reinforcement learning algorithms to learn novel application strategies that reduce the environmental impacts of nitrogen, while sustaining a sufficient yield. This investigation has shown that agents can learn application strategies comparable to standard farming practice approaches. The combination of the dueling network architecture with the Double DQN algorithm was shown to find the most favourable strategies, sustaining a sufficient yield across the varying soil and weather conditions experienced throughout the testbed years. The Dueling DDQN agent was successfully able to differentiate between seasons that require more fertilizer application amounts to maintain yield, and seasons that require very little. This approach was limited to only learning nitrogen fertilizer application strategies, and more research is required to apply this in the real world. Future work may incorporate the use of more detailed process-based models, further combining other crop growth factors such as irrigation, and phosphorous and potassium fertilizers, for a more realistic management of the crop.

References

- [1] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyperparameter optimization. *Advances in neural information processing systems*, 24, 2011.
- [2] Jonathan Binas, Leonie Luginbuehl, and Yoshua Bengio. Reinforcement learning for sustainable agriculture. In *ICML 2019 Workshop Climate Change: How Can AI Help*, 2019.
- [3] Mengting Chen, Yuanlai Cui, Xiaonan Wang, Hengwang Xie, Fangping Liu, Tongyuan Luo, Shizong Zheng, and Yufeng Luo. A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. *Agricultural Water Management*, 250: 106838, 2021.
- [4] Frédéric Garcia. Use of reinforcement learning and simulation to optimize wheat crop technical management. In *Proceedings of the International Congress on Modelling and Simulation (MODSIM'99) Hamilton, New-Zealand*, pages 801–806, 1999.
- [5] Konstantinos G Liakos, Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. Machine learning in agriculture: A review. *Sensors*, 18(8):2674, 2018.
- [6] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.
- [7] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [8] Hiske Overweg, Herman NC Berghuijs, and Ioannis N Athanasiadis. Cropgym: a reinforcement learning environment for crop management. *arXiv preprint arXiv:2104.04326*, 2021.
- [9] Upendra M Sainju, Rajan Ghimire, and Gautam P Pradhan. Nitrogen fertilization i: Impact on crop, soil, and environment. In *Nitrogen Fixation*. IntechOpen, 2019.
- [10] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [11] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [12] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- [13] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pages 1995–2003. PMLR, 2016.
- [14] Ni Zhou. Intelligent control of agricultural irrigation based on reinforcement learning. In *Journal of Physics: Conference Series*, volume 1601, page 052031. IOP Publishing, 2020.