CHALMERS

Investigating Irrigation Management Strategies with AQUACROP

Master's thesis in Systems, Control and Mechatronics

Tilde Bengtsson
Tommy Sy

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2021
www.chalmers.se

MASTER'S THESIS 2021

Investigating Irrigation Management Strategies with AQUACROP

Tilde Bengtsson
Tommy Sy

Department of Electrical Engineering
Division of System and Control
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2021

Investigating Irrigation Management Strategies with AQUACROP

Tilde Bengtsson
Tommy Sy

Supervisor: Tobias Olsson, Combine AB
Examiner: Torsten Wik, Department of Electric Engineering

Master's Thesis 2021
Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2021

Investigating Irrigation Management Strategies with AQUACROP
Tilde Bengtsson, Tommy Sy
Department of Electrical Engineering
Chalmers University of Technology

Abstract

The world's population is expected to increase from 7.9 to between 9.4 to 10.2 billion people in the coming thirty years and an increasing water demand is thereby inevitable. Globally, the largest domain of water demand is agriculture,consuming 70 % of the world's water usage. More effective usage of water within the agriculture sector is thereby required. AquaCrop is a crop model software program used for calculating, e.g., crop yield after specifying input parameters such as crop cultivar, irrigation management system and climate. Using this software, two reinforcement learning algorithms were applied to find the optimal policy representing two irrigation strategies, i.e. net amount irrigation and soil moisture targets. The goal was to find a policy minimizing irrigation amount while maintaining a yield above a certain percentage of the maximizing yield, using Q-learning. In addition, two grid searches were created for comparison. The reinforcement algorithms were able to find optimal policies if the percentage of max yield to be maintained is sufficiently low, around 90 %, for most of the time.

Keywords: Irrigation Management Strategy, Precision Irrigation, Reinforcement Learning, AquaCrop.

Acknowledgements

We would like to thank Tobias Olsson, our supervisor at Combine, for his great support and encouragement and Torsten Wik for his ideas and guidance during this project. In addition, we would like to give gratitude to our colleagues at Combine for their support and encouragement as well.

Tilde Bengtsson and Tommy Sy, Gothenburg, June 2021

Contents

List of Figures

List of Tables

1
Introduction

The world's population is expected to increase from 7.9 to between 9.4 to 10.2 billion in the coming 30 years [1]. This along with economic growth and changing consumption behaviors result in an inevitable increase in food and water demand[2] [3]. Yet, the world's resources of clean water are shrinking due to climate change, depletion of groundwater, and pollution. As known, climate change has the potential of disrupting weather patterns which cause unpredictable water availability and an increased risk of water contamination [4].

In addition, the rapid depletion of groundwater basins without knowing its remaining quantity has placed more than 30 percent of our world's largest groundwater systems in distress. Furthermore, pollution of water which is correlated with population density and economic growth has increased in the previous decades [3]. Globally, the largest domains for water demand are agriculture, industry, and domestic usage. Agriculture stands for 70 percent of the world's water demand, where the majority of the water is used for irrigation purposes and is expected to increase by 60 percent remaining the largest category of consumption [3]. Since water usage cannot exceed water availability, higher effective usage of water in agriculture is needed.

1.1 Background

The most crucial parameter for plant growth is sufficient soil moisture. Within agriculture, farmers must supply crops with enough water while preventing over-watering. If available water in the plant root zone becomes limited, the plant experiences water stress and when no available water exists for the plant to extract, the permanent wilting point, WP, is reached. At this point, the soil still contains water, but the crop is unable to extract it [5]. During wilting, stem and leaves lose their rigidness and plant growth rate stagnate. Depending on the intensity and duration of water shortage, a plant can revive to its original state once re-watered[6]. However, severe damages could be irreversible [6].

To avoid this supplying the crops with more water than necessary reduces the risk of plant stress and is an economically safer choice . The field capacity, FC, is the maximum soil moisture content where water will not drain due to gravity. If irrigation occurs at FC, the soil becomes over-watered which can cause deep percolation and surface run-off. Percolation occurs as water moves downward from the surface to the groundwater and washes the soil from its pesticides, fertilizer, and other nutrients important for the plant, resulting in lower crop yield. Surface-runoff occurs when the quantity of supplied water is too large resulting in it being unable to infiltrate the soil [7][8]. The volume of water between the FC and WP, is referenced as the Total Available Water, TAW, and depends on the soil profile. There exist thresholds of TAW for stomatal closure, canopy expansion, and early canopy senescene, i.e., aging [9]. These thresholds can be utilized in order to create an irrigation management strategy.

Irrigation controllers can operate with and without feedback from the field [2]. In fact, several constructions and management systems exist for creating an irrigation controller. Open-loop strategies include time- and volume-based irrigation while closed-loop strategies are based on feedback from the field, usually soil moisture. Due to open-loop control's low-cost properties, open-loop controls are more used worldwide, but the utilization of closed-loop controls is increasing [10]. An example of a closed-loop controller, could be a closed-loop controller fitted with a soil moisture sensor while trying to keep the soil moisture at a constant level, for example at FC or WP, or somewhere in between. What is known, is that the higher percentage of water in the soil, the larger amount of water will be lost to the surroundings due to evapotranspiration.

Evapotranspiration is the sum of water losses due to soil moisture evaporating into air and plants exchanging gas with its atmosphere in a process called transpiration. The equation for evapotranspiration, ET, can be can be written

$$ET = E + Tr, \tag{1.1}$$

where E and Tr are the evaporation and transpiration. Since these processes are hard to distinguish, the terms are normally combined into the single term. The factors for evaporation include but are not limited to temperature, air humidity, wind conditions, soil moisture and porosity, whereas transpiration rate is, among other factors, correlated with vapor pressure, plant species and water potential [8]. Due to the complex nature of the soil water dynamics, finding the optimal soil moisture target compose a challenging model problem.

In order to test control strategies, it is beneficial to simulate these in advance. This enables strategy testing in minutes instead of conducting real experiments, lasting up to a year. Also, the crops take damage if the controller is erroneous, which can be extremely costly [8]. AquaCrop is a crop model simulation tool which enables scientist and engineers to develop irrigation controllers before the real experiment. With this simulation tool, it is possible to develop irrigation strategies without conducting costly field experiments [9]. Examples of commonly used irrigation strategies are scheduled irrigation, where irrigation is conducted after a certain number of days, and irrigation as the moisture content crosses a given boundary.

As empirical relationships and heuristics are relied on for determining water loss through ET, choosing the best irrigation strategy could be a challenge. Much data is required since the best irrigation strategy is dependent on the field's geographical position and climate. The data ranges from temperature and rainfall to soil type. Data on water accessibility and crop type are also as important in choosing a suitable irrigation strategy. For example, it would not be considered wise to choose daily irrigation in places with water shortages, but still necessary if one wants to maximize the crop yield. For finding a suitable irrigation strategy from data, machine learning could be of interest, since it is capable of learning from data. Machine learning is defined as an algorithm capable of improving itself using data and collecting experience. An area within machine learning called reinforcement learning where an agent must accomplish a specific goal, is of particular interest. The basic idea within reinforcement learning can be explained as that an agent is interacting with an environment in order to acquire a reward for the new perceived state of the environment. The agent will maximize the rewards it can get as it learns the actions it should take to accomplish its goal. The benefit of this algorithm is that reinforcement learning is model-free, thereby making it suitable for applications where empirical relationships and heuristics are relied on [11].

1.2 Related Work

Several studies have validated the accuracy of AquaCrop [12], [13], [14], [15], [16]. Why AquaCrop was utilized in these papers was due to its easy-to-use implementation and low amount of required input parameters. Other models, such as the CERES-Maize Model, Muchow-Sinclair-Benett Model, EPICphase, CropSyst and the Hybrid-Maize model, require an advanced calibration process and more hard-to-measure input parameters [12]. These studies included several different crop cultivars, namely maize, rice, cotton and wheat. The field data were gathered from different types of climate i.e. arid, semiarid and rainy climate. Soil profiles included loamy and sandy soil. Irrigation strategies included both full and restricted irrigation causing water stress of various degrees.

These strategies included rainfed irrigation, irrigation where applied water was restricted to a certain depth, and withholdment of irrigation during a given amount of days or total withholdment after flowering stage [12] [13] [14]. The measured and simulated properties for comparison included canopy cover, final grain yield, above-ground biomass and soil water content. The predictions were evaluated with tools such as standard- and normalized root mean square error and a coefficient of efficiency. In order to evaluate the predictions from the different papers, the following thresholds for the normalized root mean square error was created: below 10 %, 10 - 20 %, 20 - 30 % and above 30 %, corresponding to the quality rating, excellent, good, fair and poor. The result from these studies determined that AquaCrop was able to accurately predict these crop properties when irrigation was abundant. Sometimes though, the quality of the predictions during water stress were only fair [12], [13], [14].

Multiple attempts have been conducted to optimize irrigation strategies with AquaCrop [15], [16], [17]. These studies investigated which factors impacted the final yield, irrigation and water use efficiency performance and proposed alterations of irrigation management using AquaCrop. For instance, one study evaluated the yield predictions for a varying climate, planting date and soil type by calculating a relative root mean square error [15]. The optimal irrigation was determined as optimal water productivity, i.e., the highest quotient between yield and water availability. The authors found that planting date had the largest impact on the yield outcome and created three irrigation strategies given early, normal and late planting [15].

Another study investigated cotton in North China and found an irrigation strategy that included irrigation amount, frequency and period. In terms of water productivity, a good strategy was to irrigate once for a rainy season, twice if the season contains moderate rain and three times when the given season is dry. However, this applies only to certain growth stages of cotton namely the seedling for a wet season, seedling and squaring for a normal season and lastly seedling, squaring and flowering for a dry season [16]. A third study used AquaCrop in order to find an optimal irrigation schedule given different irrigation amounts, initial soil moisture and irrigation events [17].

However, the amount of studies regarding investigations of when irrigation should be triggered given feedback from the soil moisture or which soil moisture level to remain constant at is limited. It is therefore of interest to investigate irrigation management strategies keeping the soil moisture at desired levels while reducing the water amount required.

1.3 Scope

The aim of this report was to investigate irrigation management strategies and reduce water demand within agriculture. The goal was to minimize seasonal irrigation while only obtaining a small deviation from the maximum possible final yield value. To obtain this, research regarding irrigation management in practice and the software AquaCrop was performed. Later on, for the two most prominent irrigation strategies, reinforcement learning implementation was conducted within the software AquaCrop. The goal with this implementation was to find which irrigation would provide sufficiently large yield, but with as little water as possible. To complement the results from this, two grid searches were performed which then could be used for comparison.

The calculations focused on the cultivar maize in an arid climate. No other method than grid search and reinforcement learning was performed. The data were solely simulated and not validated against any field experiment.

1.4 Outline

Chapter 2 gives a brief insight into the software program and its different inputs and output properties as well as some insight about how irrigation is performed in practice. The theory of reinforcement learning is summarized in Chapter 3. Chapter 4 summarizes how this machine learning algorithm was implemented in AquaCrop. Lastly, Chapter 5 and Chapter 6 presents the results, conclusion and proposals for future work.

2

AquaCrop

The Land and Water Division at Food and Agriculture Organization of the United States, FAO, has developed a crop model simulation software which allows its users to estimate crop yield given certain environmental conditions. The application of this program varies between performing a gap yield analysis, meaning comparing actual to potential yield, understand how environmental changes will affect biomass production in future, and investigate irrigation strategies in locations with water scarcity. It was intended to be used as an assisting tool for decision making in irrigation management systems, but it has also been used by engineers and scientist for analysis purposes. The software takes rainfall, crop system, irrigation system, climate temperate and field properties as input and return biomass production, fertility-, salinity and water stress values as output, among other properties. The software only calculates input and output in form of irrigation, rainfall, capillary raise, evapotranspiration and percolation. The crop development and transpiration is assumed to be uniform [18].

An open-source version of AquaCrop called AquaCrop-OS has also been developed by researchers from University of Manchester, Water for Food Global Institute, Imperial College London and FAO. This version of AquaCrop was written in Matlab, unlike the version developed by FAO, which was written in Delphi and distributed for Windows only. The advantage with AquaCrop-OS is its compatibility with other operating systems namely Macintosh and Linux-operating systems. Also, it is possible to integrate a water resource management since the code can be linked with additional models [19]. Besides AquaCrop-OS, there is a Python version of AquaCrop-OS called AquaCrop-OSpy and it works exactly like AquaCrop-OS. By utlizing AquaCrop in Matlab, or Python, it is possible to combine the features in AquaCrop with different machine learning algorithms.

2.1 Crop Characteristics

AquaCrop distinguishes between different types of crop: Fruit, grains, leafy vegetables, roots and tubers. This is due to the fact that grains need to consider a flowering stage, while root and tuber crop need to consider a specific formation unlike vegetables [9]. Different from various simulation tools which uses the index Leaf Area Index, LAI, AquaCrop uses Canopy Cover in order to measure plant growth [20]. Canopy cover (CC) is defined as the fraction of surface covered by a green canopy shadow, i.e.,

$$CC = \frac{A_s}{A_{tot}}, \tag{2.1}$$

where A_s is the area covered by green canopy and A_{tot} is the total area. An illustration of three plants with different values of canopy cover can be seen in Figure 2.1. One may notice how the shadow sizes differentiate from each other depending on the size of the canopy.

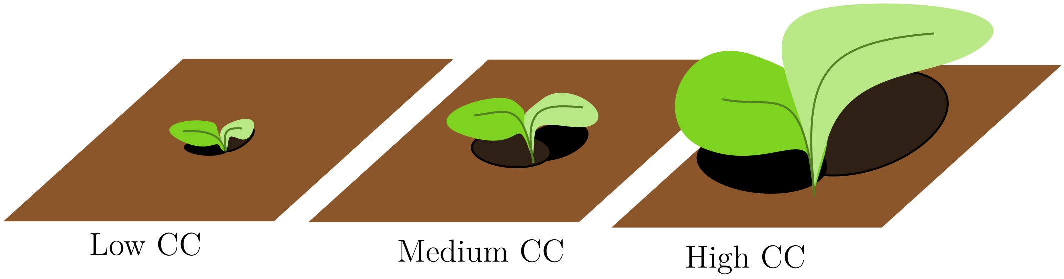

Figure 2.1: One crop with low, medium and high canopy cover.

The life cycle of a crop can be divided into four stages: emergence, canopy expansion, maximum canopy cover and canopy senescence. Emergence lasts from the moment a seed is planted until the seedling starts to sprout. It is during this stage the root system begins to develop. The Initial Canopy Cover, CC_0 has been formed as 90 % of the seedling has begun to sprout. In the canopy expansion stage, the canopy cover is increasing with the Canopy Growing Coefficient, CGC until a maximum has been established, CC_x. The next stage occurs until the crop has matured and canopy cover starts decreasing. It is also during this stage flowering occurs. When the crop has reached maturity, then the final stage, canopy senescence, begins. Canopy Cover decreases during this stage with canopy decline coefficient, CDC. During non-stress conditions, these four coefficients is all that is required to describe to crop cycle during a non-stress life cycle. An illustration of an life cycle during non-stress can be seen in Figure 2.2.

Multiple factors can cause plant stress. The most obvious, water stress, can trigger early senescence and prohibit a crop from reaching the maximum canopy cover. For this reason, the Early Senescence Stress Coefficient, Ks_{sen}, is used to modify the parameter CDC, i.e.,

$$CDC_{adj} = (1 - Ks_{sen}^8)CDC, \tag{2.2}$$

where CDC_{ajd} is the adjusted CDC. As a consequence of the modified decline coefficient, the crops life expectancy is in the risk of getting reduced. Furthermore, a reduced concentration of vital minerals inhibits plant growth and is modelled in AquaCrop by a set of Soil Fertility Stress Coefficients, Ks_{salt}, i.e.,

$$B_{rel} = 100Ks_{salt}, \tag{2.3}$$

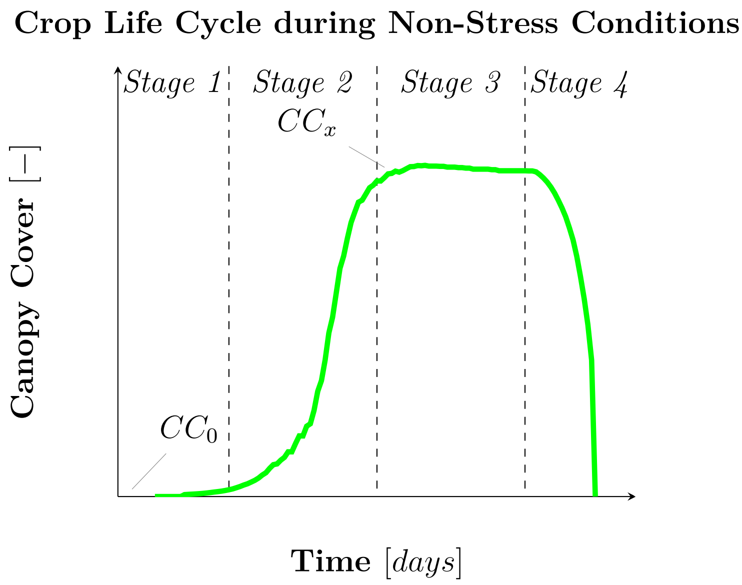

Figure 2.2: A crop life cycle with non-stress conditions expressed with the parameters Inital Canopy Cover, CC_0, Canopy Growing Coeffiecient, CGD, Maximum Canopy Cover, CC_x and Canopy Decline Coefficient, CDC.

where B_{rel} is the relative biomass. Salt increases water stress by reducing water available water in the root zone. An illustration of the life cycle of a crop experiencing stress during its life cycle can be seen in Figure 2.3. Air temperature stress affects the crop since growth stagnates below a certain temperature, the base temperature, T_b. The knowledge of low temperature stress can then be used in calculations with so called Growing Degree Day, GDD.

In order for AquaCrop to model a crop the cultivar needs to be calibrated first. In short, AquaCrop takes multiple input parameters in form of field properties, geographical location, climate and dates, i.e., planting date and date when new stages have been reached. The base temperature is the lowest temperature where the plant can grow. Given the base temperature, the mean of the daily minimum and maximum temperatures, $\bar{T}$, and dates of planting and the stage transitions, it is possible to calculate how many growing degrees is required for a plant to grow from one stage to another. A growing degree day is the difference between mean and base if the mean is above the base temperature, else it is zero, i.e.,

$$GDD = \begin{cases} \bar{T}_i - T_b, & \text{if } \bar{T}_i > T_b \\ 0, & \text{otherwise} \end{cases} \tag{2.4}$$

An illustration of growing degrees days can also be seen in Figure 2.4.
One could then for example obtain that it requires 1000 °C accumulated GDD for a certain cultivar to grow from one stage to another. This would give that a cultivar planted in the cold winter where the temperature is below or close to the base temperature requires more days to accumulate the sufficient amount of growing degrees days, when compared to the same cultivar planted in the spring where the mean

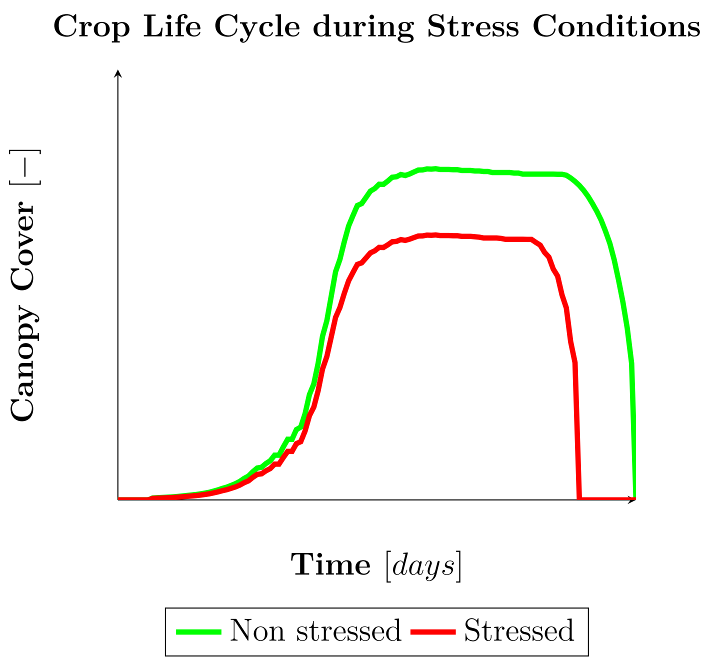

Figure 2.3: A crop life cycle with stress conditions expressed with the parameters Inital Canopy Cover, CC_0, Canopy Growing Coeffiecient, CGD, Maximum Canopy Cover, CC_x and Canopy Decline Coefficient, CDC.

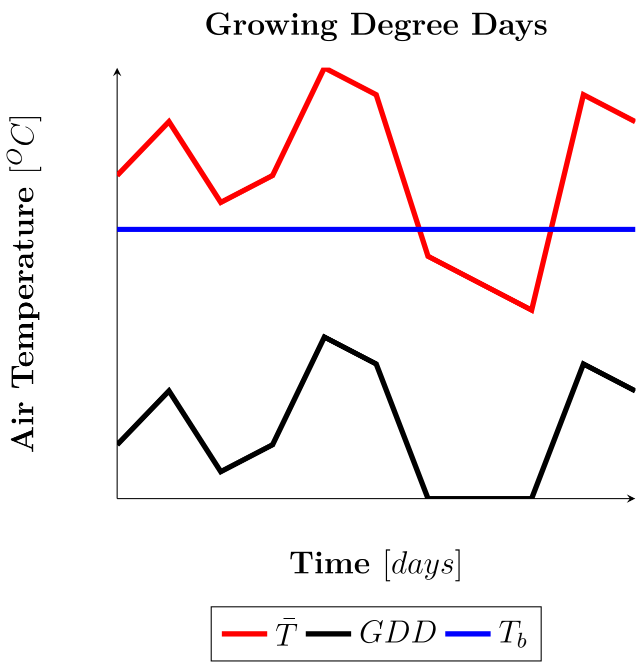

Figure 2.4: An illustration showing how mean temperature, $\bar{T}$, base temperature, T_b and growing degree days, GDD, are related.

temperature is usually much higher than the base temperature. The moment that a crops life cycle has been calibrated in growing degrees days, it is possible to simulate the growth process based on the mean temperature in the new climate. Obviously, the plant experience stress when the temperature is too high, but this is not considered in Equation (2.4).

A crop contains parameters which are both consistent and non-consistent. The consistent parameters do not change with geographical location, field management properties, time, cultivar and climate, while the non-consistent parameters do. Examples of consistent parameters are different coefficients for stress and normalized water productivity. Some non-consistent cultivar does not change with cultivar for example seed size and some non-consistent parameters change both with cultivar and with their surroundings, for example maximum canopy cover depends on plant density and maximum root depth depends on soil type.

2.2 Soil Types

The crop's growth is also affected by the soil. The thresholds for Field Capacity and Wilting Point are affected by the soil profile. For example, the presence of gravel in the soil reduces the soil's available water and the root growth can be limited if the penetrability of the soil type is low. Minerals are divided into three types of soil given the size of the particle: sand, slit, clay. Sand and clay have the largest and smallest particles, respectively. Soil have different compositions of sand, slit and clay and there are also twelve different textures, which are described in the so called textural triangle (see Figure 2.5).

A user of AquaCrop can select a soil profile after the soil types existing in the textural triangle. There is also a possibility of creating a custom soil profile, but this is usually not necessary due to AquaCrop sufficient yet wide range of profile options.

The soil can be seen as a reservoir of water which is described by the Water Balance Equation,

$$W_{ack} = I + R - P - S - ET_0, \tag{2.5}$$

where W_{ack} is accumulated water, I is irrigation, R is rainfed irrigation, P is percolation, S is surface runoff and ET_0 is the reference evapotranspiration. It is with (2.5) AquaCrop monitors the water fluxes within the crop model.

2.3 Climate Properties

AquaCrop needs daily measurements of weather data consisting of the minimum- and maximum temperature, rainfall and reference evapotranspiration. In addition, the annual concentration of carbon dioxide is required. The maximum and minimum is the highest and lowest temperature measured during a 24 hour period starting at midnight. The temperature is, as mentioned in the Section 2.1, used for calculating in form of growing degree days, whereas rainfall and evapotranspiration is used in

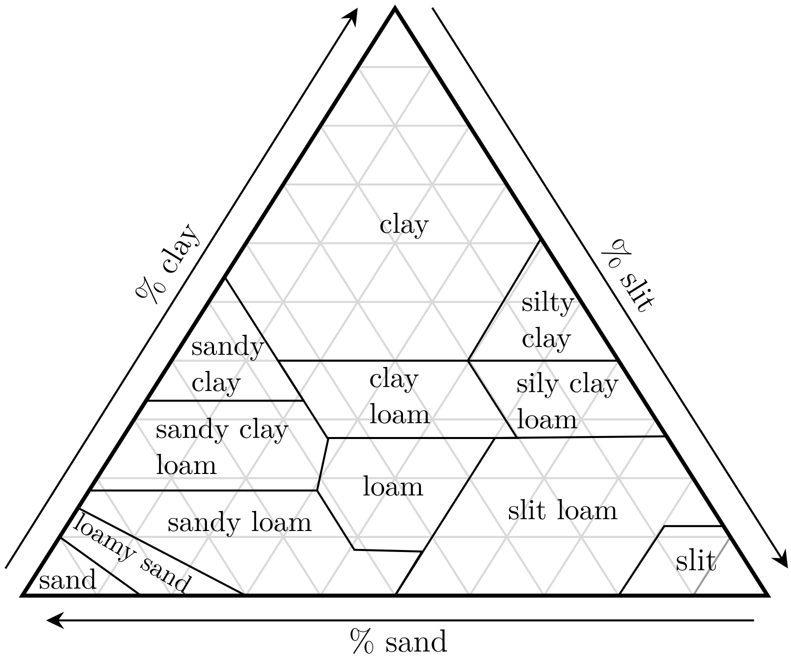

Figure 2.5: The different textural soil classes given different percentage of slit, loam and sand.

(2.5). Rain is measured in milimeters and the rain is assumed to be homogeneous over the field. AquaCrop performs calculations with a reference carbondioxide air concentration measured in year 2000 at Mauna Loa Observatory in Hawaii [9].

Since the evapotranspiration is not as easily measured as the other inputs, it is usually derived from weather data with the so called Penman-Monteith,

$$ET_0 = \frac{\Delta(R_n - G) + \rho_a c_p(\delta e) g_a}{(\Delta + \gamma(1 + \frac{g_a}{g_s}))L_v}, \tag{2.6}$$

where L_v is the energy required per water volume vaporized, Δ is the rate of change of saturation specific humidity with air temperature, R_n is the net irradiance, G is the ground heat flux, c_p is the specific heat capacity of air, ρ_a is the dry air density, δe is the vapor pressure deficit, or specific humidity, g_a is the conductivity of air, atmospheric conductance, g_s is the conductivity of stoma, surface conductance and γ is the psychrometric constant. These constants are experimentally determined e.g. the psychrometric constant is determined while using a psychrometer consisting of two thermometers with the purpose of measuring air humidity.

2.4 Field Management Properties

AquaCrop lets the user specify and manage some properties for the field. Examples of properties include the presence of mulches, affecting the soil evaporation, and soil fertility, having a negative impact on biomass production if it is limited. AquaCrop allows it user to specify if the field has soil bunds and mulches. Soil bunds decreases surface run-off since the bunds acts as a barrier preventing the water from leaving the field. Mulches decreases evaporation. How much water which can be prevented from being evaporated depends on the type of mulches, i.e. if they are made of an organic or plastic material, and how large fraction of the fields area is covered in mulches. Weed infestation is a different aspect to consider, affecting biomass production, and is measured using relative cover of weeds, RC, defined as

$$RC = \frac{WC}{CC_{tot}}, \tag{2.7}$$

where WC is weed covered area and CC_{tot} is the total area covered by weed and canopy. A large RC indicates that weed offers a stronger competition to the crop since solar energy is more absorbed by the weed. The extent of weed infestation described by RC may also be specified by choosing classes for weed management ranging from perfect to very poor.

2.5 Irrigation Management

In practice, different irrigation methods have been developed in order to satisfy the crop's water demand. These methods can be divided into traditional and modern methods depending on their proficiency in saving water, as well as their possibility of being measured, scheduled and controlled more precisely. Modern techniques can be divided into two categories, surface and subsurface irrigation, depending on whether water is applied from above or below [2].

Examples of traditional methods are manual irrigation, rain-fed farming and flooding. Manual irrigation relies primarily on the workforce when water is applied to the field while rain-fed farming relies on a hydraulic cycle. However, reliance on the hydraulic cycle as with rain-fed farming is unreliable and the quantity of water applied can vary vastly with respect to time and quantity. Flooding is the type of irrigation where an extensive quantity of water is dispersed over a field due to gravity. Water dispersion over a field means that no pump is utilized, which could reduce the cost of maintenance. Disadvantages of these techniques include uneven distribution of water in some areas, leaving some areas dryer compared to others. A difference between crop water demand and applied water over a field could also result in an increased risk for over-watering or drought [2].

Modern surface irrigation systems include sprinkler and drip techniques. The sprinkler technique involves water application in a similar manner to natural rainfall using spray-heads, pipes and pumps. Here, the spray-head disperses the water into the air before the water breaks up into small falling drops. The drip technique involves water application with narrow tubes directly or close to the plant root zone is termed drip irrigation. One reason for using this technique is to reduce water loss due to evaporation since water application is concentrated in areas where water demand is large. The result of sprinkler and drip techniques is a more uniform, precise and reliable application[2]. In contrast to surface irrigation systems, subsurface drip irrigation implies that water is applied from below to satisfy the crop water demand. Water could be applied through emitters buried underground. Also, subsurface drip irrigation has the capability of lowering water consumption, when compared to a sprinkler- and drip techniques [21].

AquaCrop simulates these application methods by utilizing different values for percentage wetted soil surface. These values can be seen in Table 2.1.

Table 2.1: The percentage soil surface wetted for irrigation strategies sprinkler, surface drip, subsurface drip and furrow.

Application Method	**Soil surface wetted** [%]
Sprinkler	100
Drip	15 - 40
Drip, subsurface	0
Flooding	100

These techniques would not be able to schedule and control irrigation without a monitoring system. Determination of irrigation need is done by observing the state of the plant or the state of the surrounding mediums. Three main types of irrigation monitoring exist: plant-based, soil-based and climate-based [2].

Soil-based monitoring means irrigation is determined by measuring certain values of the soil - usually the soil moisture, but salinity, pH and other concentration, can also serve the purpose [22]. Several studies have been conducted using soil-based monitoring where soil moisture sensors are using capacitance of sensor probes in soil for sensing the volumetric water content of the soil [23]. This works by sending an electromagnetic pulse from the sensor rod and the volumetric water content is determined from the conductance through soil and back reflection to the soil's surface. The sensor used was a low-cost one, but with sufficient accuracy, connected to an Arduino Uno micro-controller as an analog input. Gathered data was then sent to cloud storage using a WiFi-module. The cloud storage was part of Internet of Things, IoT, technology to ease the monitoring. In addition, studies have shown that it is possible to utilize pH-sensors for measuring the acidity in soil by sensing the hydrogen ion concentration, or various sensor to measure salinity in order to monitor irrigation need [24], [25].

Furthermore, plant-based monitoring is based upon visual observations. Here, optical sensors are used to judge health of the crop and for identifying causes behind low crop yield e.g. attacks from pest and lack of nutrition but most importantly water deficit. In practical sense, optical sensors such as high-definition cameras may be mounted on drones for determining the Leaf Area Index, LAI, which is a parameter used for optimal irrigation [26]. To emphasize, LAI is defined as half the total leaf area per unit ground surface area [20]. As a plant experience water stress, the leaves loose their rigidness due to wilting which can be used as monitoring option and observed with optical sensor. This proved plant-based sensing to be functional but due to practical difficulties, this type of monitoring is limited in commercial usage

[27].

In climate-based monitoring, measurement of parameters as air humidity and temperature, solar radiation and wind velocity are used for estimating the reference evapotranspiration. Evapotranspiration is the sum of evaporation from the ground surface area and transpiration i.e. the water transportation within the plant and the evaporation from the plant. The estimated reference evapotranspiration is then used to estimate water losses from the plant and soil for determining water losses from plant and soil which will be accounted for when applying water [2]. However, a study showed that this estimation of evapotranspiration requires a lot of data of high quality to avoid large estimation errors [28]. Since AquaCrop is a simulation tool, values regarding the soil, plant and climate can be read directly from AquaCrop with no measurement error. This is an advantage over real scenarios where the presence of measurement error exists.

Three types of irrigation strategies to be selected in AquaCrop include rainfed, conditioned and scheduled irrigation. Rainfed irrigation relies on rain as a water source whereas conditioned triggers irrigation after passing a threshold e.g. humidity. Scheduled irrigation on the other hand, implies that a periodic timestamp specifies the occurrence of irrigation. Alternatively, scheduled irrigation may be chosen to not be periodic. The final type of strategy is net irrigation i.e. the soil-water level shall remain constant and irrigation everyday by applying a constant volume. From these three types, six irrigation strategies exist within AquaCrop: rainfed irrigation, irrigation with soil moisture target, periodic irrigation, a predefined irrigation schedule, net irrigation and constant depth applied each day. Rain-fed irrigation implies that no artificial irrigation occurs which makes this method uncontrollable, non-consistent and risk of long time periods without water. This increases the risk for mild and severe water stress and should therefore be avoided.

AquaCrop contains three strategies that do not utilize feedback from neither the soil or the plant and can therefore be considered open loop. The first one scheduled irrigation where the user specifies an integer N, which means that irrigation is triggered every N:th day until field capacity is reached, unless a maximum irrigation amount has been specified. It is also possible to create a predefined irrigation schedule giving an overview on when water should be added and how much. In the second open loop strategy, the beginning on a predefined irrigation schedule could as for example be described: Day 1 irrigate 10 mm, day 3 irrigate 40 mm and so on. It is also possible to specify a constant depth applied each day. This is where a specified, constant irrigation amount is applied daily. These three strategies are easier to implement in practise due to their low-cost properties, though irrigation may not occur when necessary. The downside of not being able to irrigate when necessary is an irrigation strategy which is unable to handle disturbances as extreme drought or heavy rainfall.

It is therefore usually better to select an irrigation method utilizing feedback within AquaCrop, which applies water after crop water demand. Irrigation with soil moisture target means that irrigation is triggered once the water content in the soil moisture zone drops below four certain threshold - one threshold for each crop growth stage. The thresholds constitute four percentage values of the TAW in the soil reservoir. Once irrigation is triggered, water is added to the field until FC has been reached unless a maximum irrigation amount has been specified. Furthermore, the second closed loop controller within AquaCrop irrigates with respect to a given net irrigation amount. This net amount is a value of a certain fraction of TAW. Once the soil moisture lies below this value, irrigation is triggered, but adds only a small portion of water to be above the threshold. In addition, it is possible to specify how effective the irrigation is, i.e. how much of the water reaches the soil. It is also possible to specify a maximum irrigation amount in total for the season. Given all this information, an irrigation strategy can be specified for the simulation.

2.6 Calculation Scheme

The calculation can be summarized in four steps:

1. Crop Development
2. Crop Transpiration
3. Aboveground Biomass Production
4. Yield Formation.

Crop development is where the canopy cover and root zone increases before reaching their maximum in between the daily steps. Stress have a negative impact on potential maximum cover CC_{pot} and may prevent the crop from reaching it. In the next step, called crop transpiration, a crop coefficient, Kc_{Tr}, is calculated,

$$Kc_{Tr} = Kc_{Tr,x} \cdot CC, \tag{2.8}$$

where $Kc_{Tr,x}$ is a coefficient. By using $Kc_{Tr,x}$, the transpiration can be determined,

$$Tr = Ks \cdot Ks_{Tr}(Kc_{Tr,x}CC*)ET_0, \tag{2.9}$$

where Ks denote the combined water- and soil salinity stress, while Ks_{Tr} and ET_0 represent the temperature stress and reference evapotranspiration.

Following the transpiration given by (2.9), the biomass, can be calculated as,

$$B = WP \cdot \sum Tr, \tag{2.10}$$

where WP represents water productivity. Water productivity, in turn, depends on the carbon dioxide concentration in the air, soil fertility, crop and yield. A standard value for WP is used for a reference air concentration and if the carbon dioxide concentration is higher or lower than the reference concentration, then the WP value is multiplied with a correction factor. The cultivars, if rich in lipids, has an impact on the water productivity since they require more energy per dry unit biomass in the synthesis of carbohydrates. This decreases water production during the yield formation step. A substantial low soil fertility amount impacts the water productivity as well, by multiplying water productivity with a coefficient for soil fertility stress, Ks_{WP}. However, if the soil fertility is sufficient to not provoke stress, then WP is not altered. Finally, the yield, Y, can be calculated,

$$Y = HI \cdot B, \tag{2.11}$$

where HI is the harvest index. This index depends on which cultivar is used i.e. if its a grain, fruit, leafy vegetable or roots and tubers. After simulation the performance, i.e., water use efficiency, or performance, P, is calculated,

$$P = \frac{Y}{I}, \tag{2.12}$$

where I is the total seasonal irrigation [9].

3
Reinforcement Learning

When one thinks of learning, one may think of learning by acquiring knowledge through books. Learning means that new skills are acquired, e.g., learning how to drive a car or conversating with people, while being aware of the response from the environment. It is through actions that one tries to affect the events within the environment and one could state that learning by interacting with the environment is fundamental in all theories of learning. Therefore, a computational way of interacting with the environment for the sake of learning is called reinforcement learning [11].

3.1 Fundamentals

The basic idea of reinforcement learning is the agent interacting with the environment in order to acquire a reward for the new perceived state of the environment. A reward for venturing into a given state is chosen arbitrarily. For instance, when learning a humanoid robot how to walk, the upright position in conjunction with forward motion is considered a good state, whereas a lying position after falling over is considered a bad state. The agent wants to maximize its accumulative reward over time and will succeed if it learns its task [32]. Both the agent and the environment may be part of a Markov Decision Process, MDP, i.e. a mathematical description of the reinforcement learning problem where theoretical statements can be made. In short, the MDP is a Markov Reward Process, MRP, with decisions, where a MRP is seen as a tuple containing a finite set of states, a state probability transition matrix, a reward function and a discount factor for considering future states. It is commonly used for describing decision-making in sequence. Actions influence the reward and the states that follow. States that follow then affect the future reward. Also, the states are Markov, with the meaning that useful information from history is contained within the states [11] [32] [33]. An illustration of the reinforcement learning algorithm can be seen in Figure 3.1.

3.2 Expected Return

The agent's objective is to maximize its total reward over time. A term that was coined for defining the total reward after some time instance t is the expected return

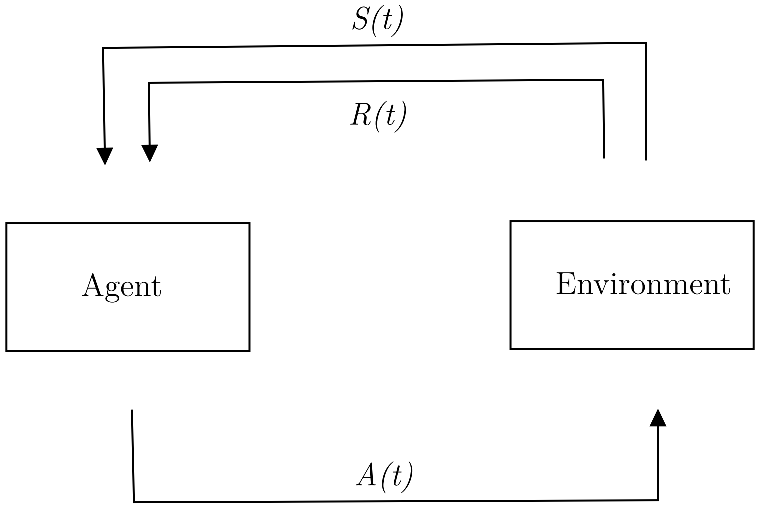

Figure 3.1: Illustration of a reinforcement learning algorithm with elements: agent, environment, action, $A(t)$, state, $S(t)$ and reward, $R(t)$.

G_t. The return at the final time t_f may be defined as

$$G_t = R_{t+1} + R_{t+2} + ... + R_{t_f}, \tag{3.1}$$

where the t_f is the time instance at which the agent ends in a terminal state and where R_t denotes the reward at time t. During an episode, transitions occur from one initial state to a terminal state. However, as the length of an episode increases the return also increases simultaneously, possibly resulting in an infinite expected return. To prevent that, a discount factor γ is introduced in (3.1) leading to discounted return defined as

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + ... = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}, \tag{3.2}$$

The discount factor γ denotes the importance of future rewards for the agent to consider and assumes a value between $0 \leq \gamma \leq 1$. Only the immediate reward R_{t+1} is considered if $\gamma = 0$, whereas future rewards become increasingly important for the agent as γ approaches 1 [11], [33].

3.3 Value-Functions

Value functions of an MDP give an estimate of whether states are good or bad. This is dependent on the expected return. The expected return is dependent on the actions the agent is about to choose. The choice of actions is associated with a policy, which is a mapping of states to future actions. Given a policy π and a state s the value-function is the expected return, i.e., [11], [33]

$$v_\pi(s) = E_\pi[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1}]. \tag{3.3}$$

Similarly, an action-value function of an MDP, given policy π, state, s, and action, a, is defined as

$$q_\pi(s,a) = E_\pi[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1}] \tag{3.4}$$

where the difference between v_π and q_π is the additional dependency on action a in q_π compared to v_π.

3.4 Bellman Equation

An alternative expression of v_π is

$$v_\pi(s) = E_\pi[R_{t+1} + \gamma G_{t+1}], \tag{3.5}$$

which is also referred to as the Bellman equation. This equation gives the average of values possible when at a particular state. Thus, an optimal value function defined as

$$v_*(s) = \max_a v_\pi(s) \tag{3.6}$$

is used for finding an optimal policy, which can be rewritten as

$$v_*(s) = \max_a E[R_{t+1} + \gamma v_*(s')] \tag{3.7}$$

, referred to as the Bellman optimality equation, for a next state $s' \in S$. $v_*(s)$ is giving the highest value as the best policy has been found. Then, the best action can be taken for a particular state in order to maximize future rewards. Similarly, a Bellman optimality equation for q_* could be expressed as

$$q_*(s,a) = E[R_{t+1} + \gamma \max_{a'} q_*(s', a')] \tag{3.8}$$

by including a next action $a' \in A$. From the Bellman optimality equation, an optimal policy is found by choosing an action at a given state to maximize $v_*(s)$ or $q_*(s,a)$. In terms of difficulty in finding the actions, the advantage with action value-functions over state value-functions is that actions are easily found at a given state. There is no need to do a one-step-ahead search for finding an action maximizing the action value function [11], [33].

3.5 Q-learning

Q-learning is a reinforcement learning algorithm where an action-value function Q estimates the optimal action-value function q_*. This action-value function is learned and converges to the optimal action-value function after updates according to

$$Q(s_t, a_t)_{new} = Q(s_t, a_t) + \alpha(R_{t+1} + \gamma(\max_a Q(s_{t+1}, a_t) - Q(s_t, a_t)), \tag{3.9}$$

where α is the learning rate. α is viewed as a measure of how fast the update should become. If the learning rate is enough low, each state-action pair will be visited and updated in order to achieve convergence . Q-learning is an example of an off-policy learning algorithm , meaning that two different policies are used, where one is used for data collection through action selection, while the other is improved as a result of the data connection. On the other hand, on-policy learning means that the policy is evaluated and improved upon while being used for action selection [11], [35].

3.6 Greediness Factor

As the agent chooses an action to perform, it may choose an action through either exploration or exploitation. Exploitation means that an agent chose the greedy action in order to maximize the reward it can get from its current action-value estimates. On the other hand, exploration means that a random action is selected with the purpose of acquiring new knowledge about actions, which may lead to greater rewards than exploitation. A way of alternating between exploitation and exploration while choosing actions is by the so-called ϵ - greedy action selection. Given a probability ϵ, the agent will perform a random action with probability ϵ. Otherwise, a greedy action will be chosen [11], [32], [35]. The ϵ-greedy action is expressed mathematically as

$$a_t = \begin{cases} \max\limits_{a} Q(s_t), & \text{with probability 1-}\epsilon \\ \text{any action,} & \text{with probability } \epsilon \end{cases} \tag{3.10}$$

This gives that a balance between exploration and exploitation is achieved for finding an optimal policy.

4
Method

The following chapter explains the setup of AquaCrop and how different irrigation management strategies were evaluated and optimized with reinforcement learning.

4.1 AquaCrop Setup

AquaCropOS-Py was available on pypi.org for download [29]. Version 0.1.5 was the most recent release on March 2, 2021. In order to perform a simulation in AquaCrop, a handful of input parameters need to be specified. AquaCrop contains already existing climate data files consisting of the minimum and maximum temperate, daily precipitation and reference evapotranspiration from a number of worldwide locations. One of these locations where data was gathered from was Tunis, characterized by its Mediterranean climate [30]. Since the majority of the world's agriculture is positioned in a hot climate where water stress is prevalent, Tunis was considered to be a suitable location for this study [31]. This climate file contained data from January 1979 to May 2002. It was decided that the number of years was sufficient.

In addition, FAO has calibrated and validated a handful of crop cultivars that already exist within the software program. These crops contain some consistent parameters not required to be re-tuned, and some non-consistent parameters which need to be re-tuned if the climate and field properties differ from where calibration has been performed. One of the most common crops in the world is maize and, since this crop had been calibrated in a Mediterranean climate, it was decided that this crop would provide sufficiently accurate results for our thesis without re-tuning [31].

Since maize had been calibrated in loam soil, it was decided that this soil type would be used for this paper. Also, the default field properties in AquaCrop were used, i.e., no mulches and neither bunds nor surface-runoff, to the limit the number of parameters that had to be adjusted.

4.2 Irrigation Strategies

AquaCrop contains six different irrigation strategies. Alternative irrigation strategies may be created in AquaCrop-OS, but the choice of irrigation strategy fell on the two options where irrigation is triggered once a threshold has been crossed, i.e., net amount irrigation or soil moisture targets. Since maize had been calibrated with a sprinkler, it was decided to use this irrigation application [9]. To simulate the sprinkler method, the percentage of wetted soil surface in Table 2.1 was used. It was assumed that the irrigation application efficiency was 100 % and no maximum irrigation amount was specified. It was also assumed that the maximum yield would be calculated with soil moisture at the FC, i.e., the soil moisture target is set to the thresehold 100 %

4.3 Reinforcement Learning Implementation

Next, two reinforcement learning algorithms were created for each irrigation strategy.

At first, the environment was created. The environment contained information regarding which crop species, initial soil moisture content, geographical location and field properties. Once the environment was set, a random year from the Tunis climate file was selected with simulation start and end, on Janaury 1st and December 31st that year. The planting date was selected to be a random day in May. This was done in order to create a more dynamic environment and reduce bias throughout the episodes.

States were defined differently for the two irrigation strategies. For the net amount irrigation, without several stages to consider, a state was specified to be an integer multiple of 5% TAW. Meanwhile, for the soil moisture target irrigation, a state was specified to be an integer multiple of 10% for each of the four stages. Actions for the irrigation strategies were defined as the target TAW to be assumed. Thus, the actions for the net amount irrigation and soil moisture target irrigation were specified to be integer multiples of 5% and 10% TAW respectively.

The reward system was designed to have the following structure: at first, the user needs to specify what yield constraint should apply, i.e., above which percentage of the maximum yield must be obtained. The reward was set to be zero for the initial state at the beginning of each episode. Once a new stage has been reached, the irrigation amount added during this stage was subtracted from the reward. Once the final state had been reached, then the final yield could be compared against the yield constraint. If the final yield was higher than the yield constraint, then a bonus reward was given of 1000. This gave that the highest reward would be given as the yield constraint was satisfied, but it should also be obtained with as little water as possible.

In order to explore finding information about the environment while exploiting known information for the purpose of maximizing the reward, an epsilon greedy function was implemented according to (3.10). The epsilon greedy function was chosen because of it being a simple way to guarantee lasting exploration, where ε denotes the probability of choosing a random action. However, given a large number of states and actions with no prior information about the environment from start, exploration was determined to be significantly more important than exploitation. Hence, the ε was chosen to decay in a logarithmic way according to

$$\varepsilon = \varepsilon - (1 - \varepsilon)e^{-c*k} \tag{4.1}$$

over the number of episodes k, where c is a constant specifying the speed of ε decay over the episodes. It was chosen to be 10 since this value provided a sufficient decay for 100 000 episodes. A limit on how small ε could become was selected to be 0.1, since continual exploration is important if a seemingly optimal irrigation strategy has been found. The learning rate was chosen to be 0.1 and gamma to be 0.9. A high value of gamma would give that the next and final reward would impact the Q-value in a previous state.

At last, the Q-learning algorithm could be written as described in Chapter 3. The result from this was one optimal policy for every yield constraint and irrigation strategy.

4.4 Grid Search

Next, validation data was created with grid searches for the two chosen irrigation strategies. This validation data would then be utilized for comparisons with the optimal irrigation strategies found by the reinforcement learning algorithm. For the net amount irrigation strategy, a discretized range from five to 100 percent of TAW steps of five was used as input into AquaCrop, i.e.,

$$N_i = i \cdot 5, \qquad \text{where } i = 1, 2, ..., 20, \tag{4.2}$$

and N_i denotes the net amount input. For the soil moisture target irrigation strategy, AquaCrop needed four different percentage values of TAW for each growing stage. Therefore, a grid search was created with values from ten to 100 with a step size of 10. The grid can be expressed:

$$SMT_{ijkl} = [i, j, k, l], \qquad \text{where } i, j, k, l = 10, 20, ..., 100, \tag{4.3}$$

where SMT is the soil moisture targets. From the grid searches the expected seasonal irrigation, yield and performance were calculated. It was noted which irrigation strategy provided a yield higher than 90 %, 95 % and 99 % of the expected maximum yield.

The data from the grid search would then be compared with the data from the reinforcement learning implementation. This was done by comparing the magnitude of the deviations i.e. absolute error.

5

Results

The following chapter displays the result from the grid search and the reinforcement learning implementation. At first, the maximum yield for maize growing in Tunis was calculated as the expected yield with full irrigation, i.e. the irrigation strategy was soil moisture target with threshold [100,100,100,100] for the 22 years of data available. The maximum yield could then be pinpointed to be 11.41 tonne/ha. The yield representing 90, 95 and 99 % of maximum yield was calculated and used as the yield constraints. Given this, it was found, with help of the reinforcement learning implementation and grid search, which thresholds for the irrigation strategies would provide a yield above the yield constraints, with a minimal amount of irrigation water. Further, the expected yield, seasonal irrigation, performance for the two irrigation strategies and implementations were found. The results from this can be seen in Table 5.1.

Table 5.1: The found thresholds for the net amount, N, and soil moisture target, SMT, and the exptected yield, Y, irrigation amount, I, and performance, P, for the different yield constrains, Y_{max}, calculated with the grid search and with reinforcement learning implementation, RL.

RL								
	NET				**SMT**			
$\mathbf{Y}_{max}$ [%]	**N** [%]	**Y** $[\frac{tonne}{ha}]$	**I** $[mm]$	**P** $[\frac{tonne}{ha \cdot mm}]$	**SMT** [%]	**Y** $[\frac{tonne}{ha}]$	**I** $[mm]$	**P** $[\frac{tonne}{ha \cdot mm}]$
90	80	11.3	572	0.0198	[40,50,20,20]	11.1	546	0.0203
95	90	11.4	583	0.0196	[60,60,30,50]	11.4	598	0.0191
99	5	0.00	3.65	0.00	[20,10,60,60]	8.45	528	0.0160
GRID								
	NET				**SMT**			
$\mathbf{Y}_{max}$ [%]	**N** [%]	**Y** $[\frac{tonne}{ha}]$	**I** $[mm]$	**P** $[\frac{tonne}{ha \cdot mm}]$	**SMT** [%]	**Y** $[\frac{tonne}{ha}]$	**I** $[mm]$	**P** $[\frac{tonne}{ha \cdot mm}]$
90	55	10.4	515	0.0202	[40,30,10,10]	10.4	515	0.0202
95	65	11.0	552	0.0199	[50,30,20,10]	10.9	530	0.0205
99	80	11.3	572	0.0120	[50,60,30,20]	11.3	562	0.0201

Figure 5.1 displays the result from the reinforcement learning implementation for the net amount irrigation strategy. This includes the resulting reward, yield, total seasonal irrigation and performance throughout the episodes.

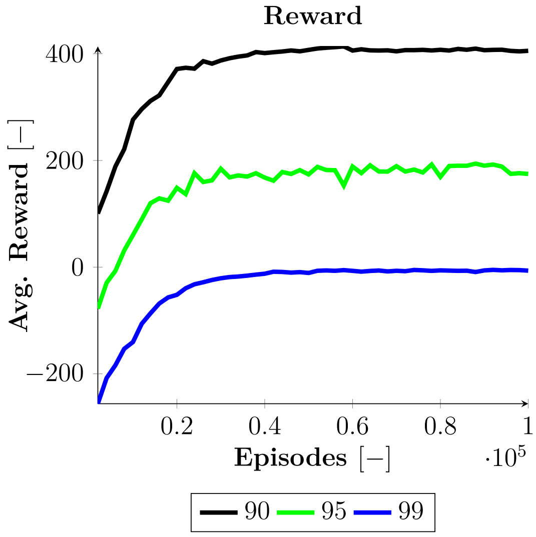

(a) Reward

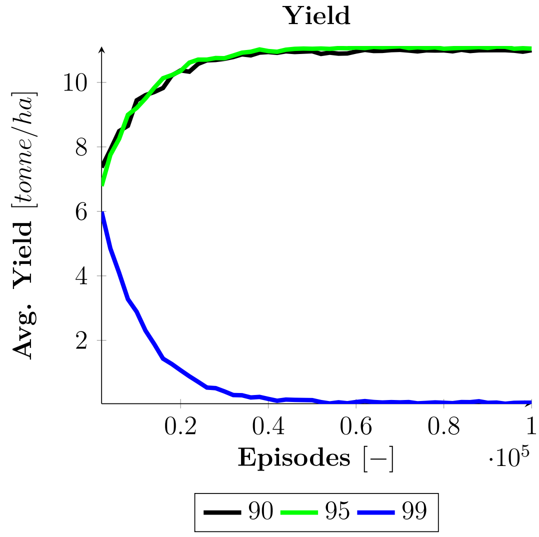

(b) Yield

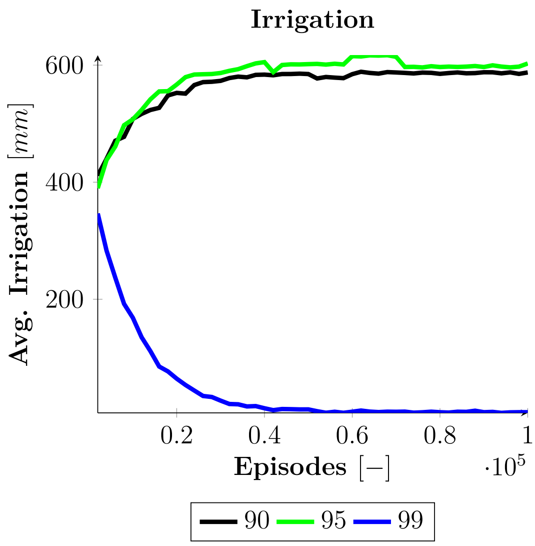

(c) Irrigation

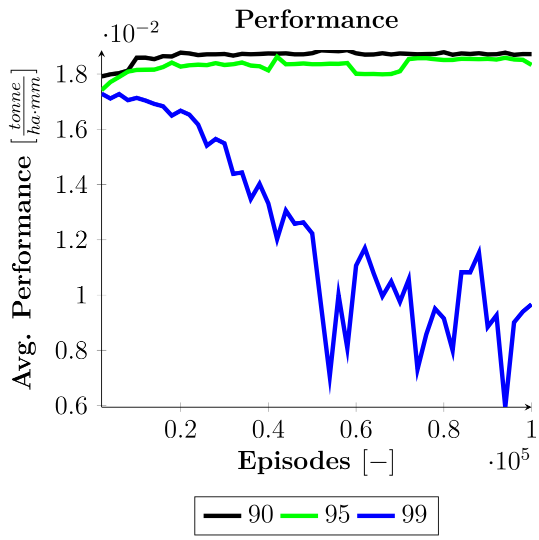

(d) Performance

Figure 5.1: The average reward, yield, total seasonal irrigation and performance during 10000 episodes with the reinforcement learning implementation on the net amount seasonal irrigation where the constraint at least 90, 95 or 99 % of max yield.

Figure 5.2 displays the reward, yield, irrigation and performance from the reinforcement learning implementation for the soil moisture target irrigation strategy.

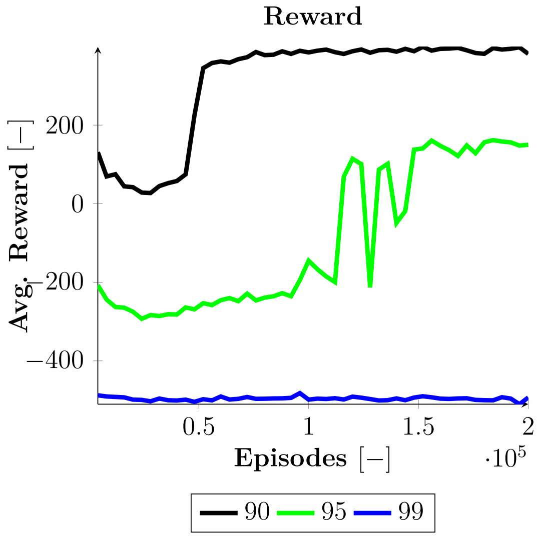

(a) Reward

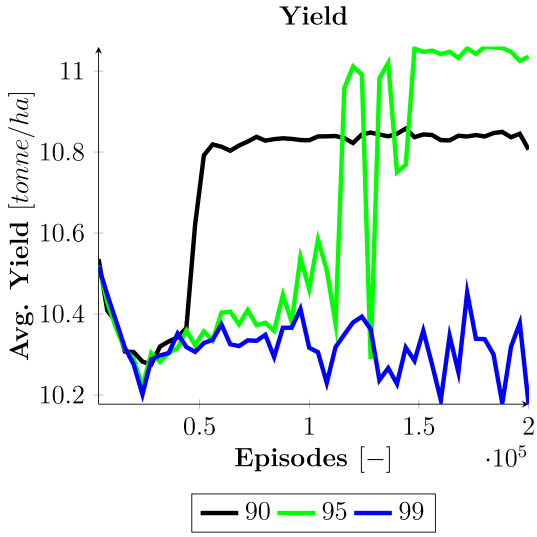

(b) Yield

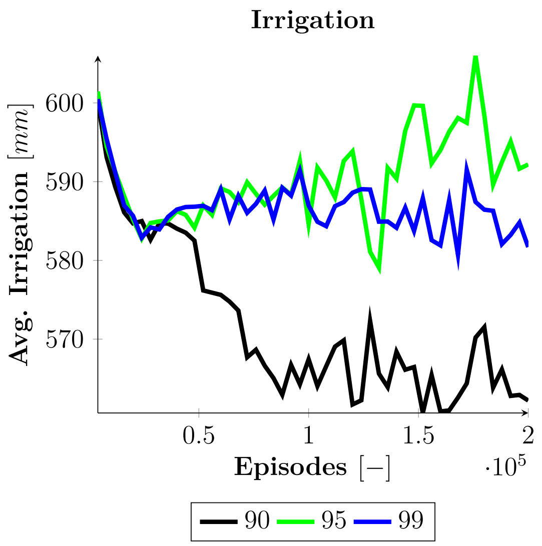

(c) Irrigation

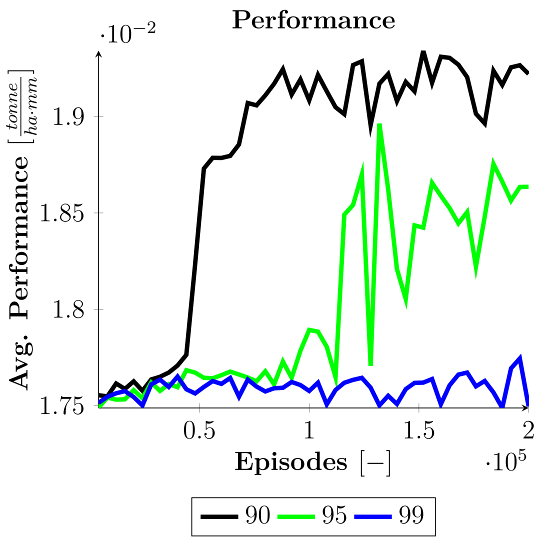

(d) Performance

Figure 5.2: The average reward, yield, seasonal irrigation and performance during 10000 episodes with the reinforcement learning implementation on the net amount seasonal irrigation where the constraint at least 90, 95 or 99 % of max yield.

Figure 5.3 displays the result from the grid search for the net amount irrigation strategy.

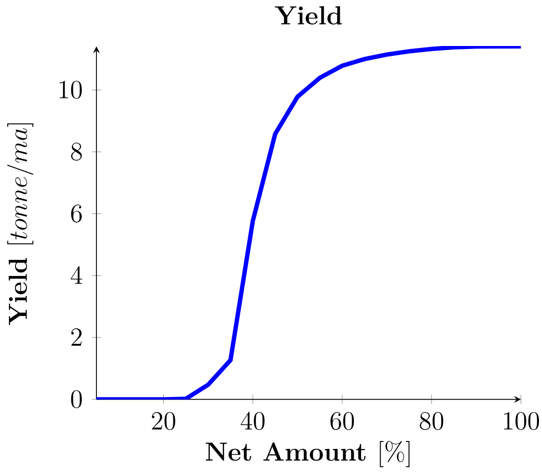

(a) Yield

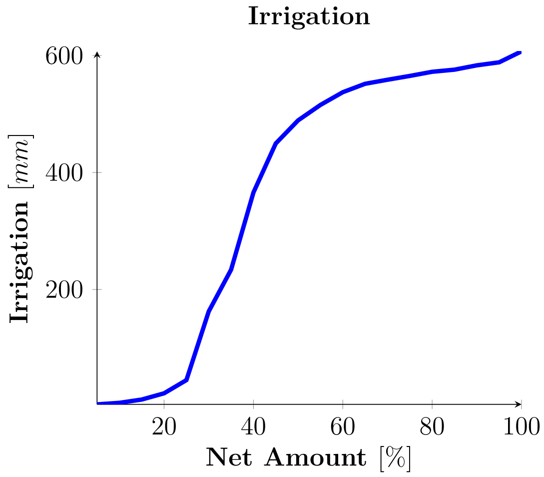

(b) Irrigation

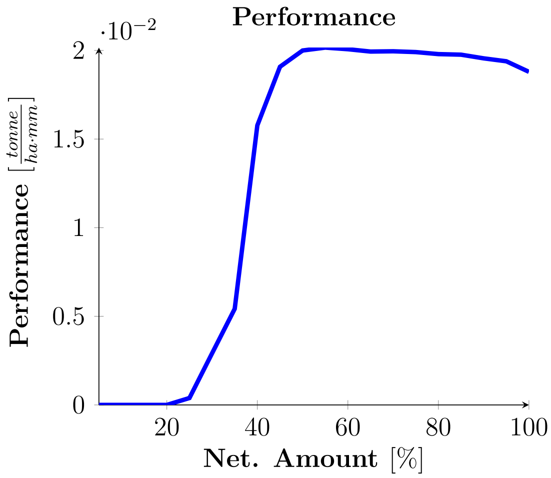

(c) Performance

Figure 5.3: The resulting yield, irrigation and performance from the grid search with net amount irrigation strategy.

Figure 5.4 displays the result from the grid search for the soil moisture target irrigation amount.

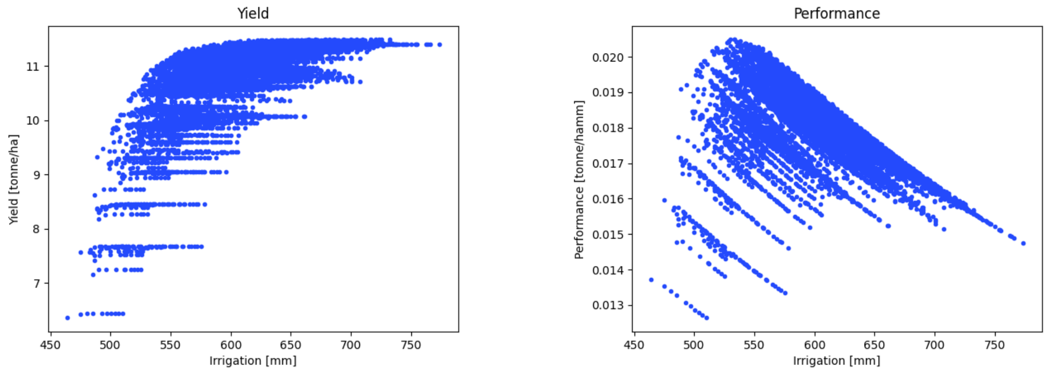

(a) Yield

(b) Performance

Figure 5.4: The yield, irrigaiton and performance from the grid search with soil moisture target irrigation strategy.

6

Discussion

The following chapter discusses the result, choice of method and how this correlates with the purpose of the thesis, as well as irrigation management in practice. Finally, the work is concluded and suggestions for future work are presented.

6.1 Results

By observing the plots in Figure 5.1 (a), it can be seen that the agent finds an optimal policy for the net amount irrigation strategy as the reward plot converges. If an optimal policy was never found the reward plot would not converge, and because the action space is limited and discretized, it is interpreted as if the agent is choosing the same action sequences repeatedly, since there are no better sequences of actions to choose at a particular environmental state. This guarantees to find an optimal policy if the action space is small as in 5.1 for 100000 episodes. Interesting to note is that the reward systems was giving a bonus as final yield was above 90 and 95 % of maximum yield, since the reward converges to a positive value which implies it was able to capture the bonus reward when exploiting. This means that an optimal policy has been found simply because the yield constraints was surpassed. As can be seen in Table 5.1, the result from the grid search implied that a sufficient yield could be obtained when net amount was 55 and 65 % of TAW for yield constraints 90 and 95 %. In terms of saving water, this gives that the result from the grid search is better than the result from the agent, because 55 % and 65 % from the grid search is lower than 80 % and 90 % in order to achieve sufficient yield according to the agent.

When observing Table 5.1, it is of interest to explain why the result was found to be 80 and 90 % of TAW for the agent since it is not very obvious to why this is the case. There are many causes of this, such as the different action space between the grid search and the net amount irrigation but besides that, an explanation could be due to the yield varying throughout the years because yield data in AquaCrop was collected from real sites where crop yield may have huge deviations from a mean yield among all years . The maximum yield in this case is the expected maximum among all years. This gives that the yield constraint might be higher than expected, and as result the agent will find an optimal policy with more water with the purpose of surpassing the yield constraint.

From Figure 5.1 (a), it can be seen that the reward when yield constraint is set at

99 % converges towards a negative value. An explanation could be that the agent was rarely above 99 % of the yield maximum for the irrigation strategies it has found. This implies that the agent rarely finds an irrigation strategy which provides a sufficient yield and the policy can therefore be regarded as unsuccessful. As a consequence of the agents inability to find a successful strategy, the agent exploit water saving strategies as this provide a higher reward with this reward system. This explains why the policy for yield constraint 99 % is 5 % of TAW i.e. the most water saving method in Table 5.1. This alone explains why yield and irrigation amount converges towards zero which is illustrated in Figure 5.1 (b) and (c).

Furthermore, for the soil moisture strategy it can be seen from Figure 5.2 a) that the reward system with constraint 90 % and 95% was able to find an optimal solution with a yield above the constraint. The rewards converge at around 50 000 and 150 000 episodes because the reward is given at the final stage if the constraint is low enough. 90 % and 95% were revealed to be low enough given the looks for the reward, yield, irrigation and performance in Figure 5.2.

For the constraint 99 %, the agent was unable to find a policy satisfying the yield constraint, which explains why the reward in Figure 5.2 (a) is negative. The reason for this is believed to be because of the yield attained in the final stage not being high enough to acquire the positive reward. Much watering applied while the reward is never attained, makes the cumulative reward negative. Thus, the cumulative reward with constraint 99 % is smaller than 90 % and 95 %.

A comparison between net amount and soil moisture target irrigation can also be made from Table 5.1. By comparing the grid searches it is illustrated that the SMT strategy obtained a slightly higher performance. It can be seen that the yield was roughly the same whereas the irrigation amount was a bit lower for the SMT strategy, which explains the increased performance. The reason why SMT had a lower irrigation amount is believed to be because this strategy takes into consideration the different crop cycles. In the net amount strategy, the irrigation is the same all the time, which gives that the water uptake in each crop stage is not considered. This could lead to over-watering in the first stage where roots are not fully developed and under-watering in later stages where more water is required. SMT is therefore better than net amount irrigation in the sense that water usage is more optimized with respect to the different stages. The performance for SMT with Q-learning was lower than for the grid search was, which is believed to be due to the large number of states, making it hard for the agent to find an optimal strategy. Also, the performance for net amount irrigation was lower for the Q-learning implementation then for the grid search.

The computation for net amount irrigation and soil moisture target was more expensive than the grid search due to exploration and the large action space for the agent. A decrease in computation time is achieved if the action space is decreased. The downside using this approach is the additional decrease in accuracy for the agent to find a good irrigation strategy. Thus, the computation for SMT was more expensive than net amount irrigation and much more expensive than grid search, further explaining its low performance as not sufficient time has passed.

6.2 Irrigation Strategy

For this experiment, it had been decided to use a sprinkler system with a 100 % wetted surface. By changing the irrigation method to a drip or surface drip system, it is believed that the thresholds values from the result would obtain lower values. However, a comparison between these methods lied outside the scope, but could be worth investigating in future work.

6.3 Reinforcement Learning

The designed reward system was functional when the constraint for reward was low, but ineffective for higher constraints as the agent was unable to receive the bonus. In order to avoid this, the design of the reward system could be altered. This reward system could be either performance-based which implies the agent receives a bonus when it performs an action that increases the performance of the irrigation method. However, this could potentially result in a yield closer to half of the maximum possible small amounts of only if water is added. This implies that the crop is only half-grown and the models are not calibrated for this. Instead, the reward could be the sum of irrigation and yield multiplied with different weights.

The structure of the states and actions could also be designed differently. One alternative would have been to have states represented by canopy cover, soil moisture and air temperature where the agent would take an action, i.e. irrigation amount in millimeters instead of soil water percentage. It would be preferred to not discretize the states, since true optimal irrigation strategy is not likely to exactly match the values, i.e. several runs could give different irrigation strategies close to each other. This implies that an optimal irrigation strategy is somewhere in between the discretization step. The reward system has been utilized for a hot environment where the evapotranspiration is high, causing the soil moisture content to decrease fast. A temperate environment would cause the soil moisture content to decrease slower, and it would be of interest to see how different the results the current reward system would be.

It could also be discussed if reinforcement learning and Q-learning was the most suitable option for finding irrigation strategies. The agent was able to find an irrigation strategy, but was restrained by the stage and action space. If the optimal value for the net amount irrigation would lie between, e.g, 60 and 65, it would not be possible for the agent to find it. This could be achieved be selecting a larger state space, but it would then be harder for the agent to find the optimal value. Instead it could be more promising to investigate optimization methods such as the gradient descendant method or Quasi-Newton for finding optimal irrigation strategies.

6.4 In Practice

The resulting thresholds from the net irrigation and soil moisture target could be used by agriculturists as reference points for their closed loop controllers. This is done by measuring the field capacity and wilting point of the soil and the threshold can thereby be obtained.

Irrigation methods for water application may not be as accurate in reality as water losses are non-existent within AquaCrop and one could ask whether AquaCrop is sufficiently accurate in representing irrigation in practice. Unfortunately, AquaCrop does not contain a feature which predicts the uncertainty of the result. It can be assumed that if the crop cultivar is simulated for an environment similar to that has been calibrated, then AquaCrop should provide results with enough accuracy as discussed in Chapter 2. In order to validate if these irritations method are beneficial in reality, field experiments have to be performed. Then, a yield gap analysis can be done, which measures the gap between the simulated and actual yield analysis.

6.5 Conclusion

The conclusion of this thesis was that it is possible to deploy irrigation management strategies with reinforcement learning and AquaCrop, given a reward system with sufficiently low yield constraints. The Q-learning implementation was however not able to find the same optimal strategies as the grid search. By comparing the irrigation strategy with net amount irrigation and soil moisture target, it could be seen that in general, the irrigation with soil moisture target could generate a higher performance by finding strategies using less water. This is due to the net amount irrigation not considering the different water uptake in each crop stage.

Furthermore, the choice of method, i.e., reinforcement learning could be discussed since this method was not able to find an optimal policy outside the state and action space. While trying to find the optimal threshold for irrigation, Q-learning was discovered to not be the most suitable option due to the vast number of states and actions even though a quite rough discretization of those states and actions have been made. An alternative route instead of RL could be an optimization method based on gradient descent.

6.6 Future Work

For future work, suggestions include investigating and tuning a weight-based reward system, investigating how different irrigation method such as drip and sprinkler irrigation affect the resulting optimal policy and complementing the result with field studies. Field studies are also important for validating the result. Measurements regarding soil moisture could be collected from sensors in varying depth as too determine whether SMT and net amount irrigation is the most suitable. Measurement data could be relayed to an irrigation system that also takes weather and climate data into account. Therefore, a future irrigation system utilizing RL could become very effective in saving water if the weather or climate changes rapidly in the upcoming future. It is also of interest to make predictions of whether the plant will grow given the irrigation amount. However, for a farmer to simply make the decision on irrigation amount without even planting the crop, the farmer could make measurements on soil moisture and weather parameters such as temperature on a specific location. If there is an obvious irrigation amount corresponding to a pattern of the data, one could be for sure that the agent in this thesis has been validated to find an optimal irrigation strategy in the sense that yield is maximized while irrigation is minimized as represented by performance.

Bibliography

[1] UNESCO World Water Assessment Programme, *The United Nations world water development report 2018: nature-based solutions for water*, Paris: UNESCO, 2018, pp. 10

[2] E. A. Abioyea, M. S. Z. Abidina, M. S. A. Mahmuda, S. Buyamina, M. H. I. Ishaka, M. K I. A. Rahmana, A. O. Otuozec, P. Onotub, M. S. A Ramlid, "A review on monitoring and advanced control strategies for precision irrigation", *Computers and Electronics in Agriculture*, vol. 173, pp. 1 - 22, 2020. Available doi: 10.1016/j.compag.2020.105441

[3] A. Boretti, L. Rosa, "Reassessing the projections of the World Water Development Report", *npj Clean Water*, vol. 2, no. 15, pp. 1 - 2, 2019. Available doi: 10.1038/s41545-019-0039-9

[4] M. M. Q. Mizra, "Climate change and extreme weather events: can developing countries adapt?", *Climate Policy*, vol. 3, no. 3, pp. 233 - 248, 2003. Available doi: 10.3763/cpol.2003.0330

[5] J. A. Tolk, "Soils. Permanent Wilting Points", United States Department of Agriculture, 2003.

[6] X. Zu, G. Zhoung, H. Shimizu, "Plant responses to drought and rewatering", *Plant Signaling and Behavoir*, vol. 5, no.6, pp. 649 - 654, 2010. Available doi: 10.4161/psb.5.6.11398

[7] Meeradevi, Supreetha M.A., Mundada M.R., Pooja J.N. (2019) Design of a Smart Water-Saving Irrigation System for Agriculture Based on a Wireless Sensor Network for Better Crop Yield. In: Kumar A., Mozar S. (eds) ICCCE 2018. ICCCE 2018. Lecture Notes in Electrical Engineering, vol 500. pp. 93 - 95, Springer, Singapore. Available doi: 10.1007/978-981-13-0212-1_11

[8] P. D. Colaizzi, "Overwatering controller for landscape irrigation systems", M.S. Thesis, Dpt. of Agricultural Biosystems Engineering, University of Arizona, pp. 11 - 27, 1997. Accessed on: Mar. 15, 2021. Available doi: http://hdl.handle.net/10150/278585

[9] *AquaCrop Reference Manual*, 6th ed, Food and Agriculture Organisation of the United Nations, Rome, 2018.

[10] Y. Mao, S.Liu, J. Nahar, J.Liu, F.Ding, "Regulation of soil moisture using zone model predictive control," *IFAC (International Federation of Automatic Control)*, vol. 51. no. 18. pp. 762 - 767, 2018. Available doi: 10.1016/j.compag.2020.105441

[11] R. S. Sutton, A. G. Barto, *Reinforcement Learning: An introduction*, 2nd ed, Cambridge, MA: MIT Press, 2020

[12] L. K. Heng, T. Hsiao, S. Everett, T. Howell, P. Steduto, "Validating the FAO Aquacrop Model for Irrigated and Water Deficient Field Maize", *Agronomy Journal*, vol. 101, no. 3, pp. 488 - 489, 2009. Available doi: 10.2134/agronj2008.0029xs

[13] B. Andarzian, M. Bannayan, P. Steduto, H. Mazraeh, M. E. Barati, M.A. Barati, A. Rahnama, "Validation and testing of the AquaCrop model under full and deficit wheat production in Iran", *Agriculture Water Management*, vol. 100, no. 1, pp. 1 - 8, 2011. Available doi: 10.1016/j.agwat.2011.08.023

[14] M.Maniruzzaman, M.S.U. Talukder, M.H. Khan, J.C. Biswas, A. Nemes, "Validation of the AquaCrop model for irrigated rice production under varied water regimes in Bangladesh", *Agriculture Water Management*, vol. 159, pp. 331 - 340, 2015. Accessed on: May 29, Available: https://doi.org/10.1016/j.agwat.2015.06.022

[15] A. Araya, I. Kisekka, J. Holman, "Evaluating deficit irrigation management strategies for grain sorghum using AquaCrop", *Irrigation Science*, vol. 34. no. 6, pp. 465 - 481, 2016. Accessed on: May 3, Available: 10.1007/s00271-016-0515-7

[16] F. Li, D. Yu, Y. Zhao, "Irrigation Scheduling Optimization for Cotton Based on the AquaCrop Model", *Water Resorce Management*, vol. 33, pp. 39 - 55, 2019. Available: 10.1007/s11269-018-2087-1

[17] J. Li, J. Song, M. Li, S. Shang, X. Mao, J. Yang, "Optimization of irrigation scheduling for spring wheat based on simulation-optimization model under uncertainty", *Agricultural Water Management*, vol. 208, pp. 245 - 260, 2018. Available: 10.1016/j.agwat.2018.06.029

[18] E. Vanuytrecht, D. Raes, P. Steduto, T. C. Hsiao, E. Fereres, L. K. Heng, M. G. Vila, P. M. Moreno, "AquaCrop: FAO's crop water productivity and yield response model", *Environmental Modelling Software*, vol. 62, pp. 351 - 360, 2014. Available: https://doi.org/10.1016/j.envsoft.2014.08.005

[19] T. Foster, N. Brozovic, A. P. Butler, C. M. U. Neale, D. Raes, P. Steduto, E. Fereres, "AquaCrop-OS: An open source version of FAO's crop water productivity model", *Agricultural Water Management*, vol. 181, pp. 18 - 22, 2018. Available: doi.org/10.1016/j.agwat.2016.11.015

[20] J.M. Chen, T.A. Black, "Measuring leaf area index of plant canopies with branch architecture," *Agricultural and Forest Meteorology*, vol. 57. no. 1-3, pp. 1 - 12. 1991. Available doi: 10.1016/0168-1923(91)90074-Z

[21] E. Nazari, S. Bezharat, K. Zeinalsade, A. Mohamadi, "Measurement and simulation of the water flow and root uptake in soil under subsurface drip irrigation of apple tree", *Agriculture Water Management*, vol. 255, pp. 1 - 10, 2021. Available doi: 10.1016/j.agwat.2021.106972

[22] I.A. Lakhiar, G. Jianmin, T.N. Syed, F.A. Chandio, N.A. Buttar, W.A. Qureshi, "Monitoring and Control Systems in Agriculture Using Intelligent Sensor Techniques : A Review of the Aeroponic System," *Journal of Sensors*, vol. 2018, pp. 1 - 18, 2018. Available doi: https://doi.org/10.1155/2018/8672769

[23] S. Ryo, K. Yoshihiro, G. Divya Goud, B. Balaji Naik, "Capacitive-Touch-Based Soil Monitoring Device with Exchangeable Sensor Probe", *2018 IEEE SENSORS*, pp. 1 - 4, 2018. Available doi: 10.1109/ICSENS.2018.8589698

[24] J. Divya, M. Divya, V. Janani, "IoT based Smart Soil monitoring System for Agricultural Production," *IEEE Interntional Conference on Technological Innovations in ICT for Agriculture and Rural Development (ITAR 2017)*, pp. 209 - 214, 2017. Available doi: 10.1109/PDGC50313.2020.9315836

[25] G.I. Metternicht, J.A. Zinck, "Remote sensing of soil salinity: potentials and constraints", *Remote Sensing of Environment*, vol. 85, no. 1, pp. 1 - 20, 2005. Available doi: 10.1016/S0034-4257(02)00188-8

[26] M. El Adnani, S. Belaqziz, A. Chehbouni, S. Er-Raki, S. Khabba, M.H. Kharrou, L. Jarlan, M. le Page, "A new irrigation priority index based on remote sensing data for assessing the networks irrigation scheduling", *Agricultural Water Management*, vol. 119, pp. 1 - 9, 2013. Available: doi.org/10.1016/j.agwat.2005.02.013

[27] Hamlyn G. Jones, "Irrigation scheduling: advantages and pitfalls of plant-based methods", *Journal of Experimental Botany*, vol. 55, no. 407, pp. 2427 - 2436. Available doi: 10.1093/jxb/erh213

[28] S. Trajkovic, "Temperature-Based Approaches for Estimating Reference Evapotranspiration", *Journal of Irrigation and Drainage Engineering*, vol. 131, no. 4, pp. 316 - 323, 2005. Available: 10.1061/(ASCE)0733-9437(2005)131:4(316)

[29] https://pypi.org/project/aquacrop/

[30] S. Achour- Younsi, F. Kharrat, "Outdoor thermal comfort: Impact of the geometry of an urban street canyon in a Mediterranean subtropical climate – Case study Tunis, Tunisia", *Procedia - Social and Behavioral Sciences*, vol. 216, pp. 689 - 700, 2016. Available: 10.1016/j.sbspro.2015.12.062

[31] "Value of Agriculture Produktion", Food and Agriculture Organization of the United States, 18 March 2021. Accesses: May 05 2021. Available: http://www.fao.org/faostat/en/data/QV

[32] D. Silver. (2015). Introduction to Reinforcement Learning [PowerPoint slides]. Available: https://www.davidsilver.uk/wp-content/uploads/2020/03/intro_RL.pdf

[33] D. Silver. (2015). Markov Decision Processes [PowerPoint slides]. Available: https://www.davidsilver.uk/wp-content/uploads/2020/03/MDP.pdf

[34] D. Silver. (2015). Model-Free Prediction. Available: https://www.davidsilver.uk/wp-content/uploads/2020/03/MC-TD.pdf

[35] D. Silver. (2015). Model-Free Control. Available: https://www.davidsilver.uk/wp-content/uploads/2020/03/control.pdf

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY