Climate-Resilient Water Management via Reinforcement Learning

Impact of varying climate conditions on water management of the Nile River Basin using Reinforcement Learning

Tadas Lukavičius[1]
Supervisor(s): Zuzanna Osika[1]**, Pradeep Murukannaiah**[1]
[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 26, 2024

Name of the student: Tadas Lukavičius
Final project course: CSE3000 Research Project
Thesis committee: Zuzzana Osika, Pradeep Murukannaiah, Luis Miranda da Cruz

Abstract

This project aimed to investigate reinforcement learning (RL) algorithms to improve water management policy development in the Nile Basin, with a focus on the Multi-Objective Natural Evolution Strategies (MONES) and Evolutionary Multi-Objective Direct Policy Search (EMODPS) algorithms. This project intended to refactor a Nile Basin simulation to be compatible with the MONES algorithm, which continues the exploration of different machine learning algorithms in water resource management. Additionally, the RL algorithms were aimed at training using two climate data sets: human-favourable and climate-varying conditions, and then evaluating on the satisfaction and regret metrics. The successful integration of the MONES framework shows the feasibility of utilizing advanced RL algorithms for water management problems. Initial results indicate that the MONES algorithm underperforms compared to the EMODPS algorithm according to hypervolume and diversity of solutions, however, further research is needed to test whether this claim holds. The EMODPS algorithm faced challenges in finding optimal solutions when dealing with variable climate conditions scenarios, accentuating the need for robust solutions, which consider a variety of possible climate outcomes. The observed sensitivity to variable climate conditions underlines the crucial importance of accurate and recent data, as well as the need to consider the climate change effects on water management. The study concluded with suggestions that future simulations of water management strategies may be improved if a broader set of external factors and a more realistic representation of objectives are included in the simulation model. These improvements stand to positively impact the accuracy, applicability and reliability of future simulations.

1 Introduction

In today's rapidly advancing world, the challenges posed by climate change, specifically the need for effective water management, become increasingly prominent. Recent climate patterns and weather events are leading to periodic droughts and floods [15]. Since water resources are becoming more scarce and less predictable-such as the expected rainfall or groundwater levels at specific times of the year- it is of crucial importance to manage them effectively[8]. Additionally, projects in this field have to take into account multiple concerns that are not only based on resources themselves but also those of the environment and society [21]. Fortunately, machine learning algorithms can be of help when addressing such multi-faceted issues[6]. It is possible to cast the problem of devising effective policies for the management of water resources as an optimization problem. The problem is subject to constraints imposed by environmental and societal concerns in addition to the scarcity of resources. A machine learning algorithm is a kind of tool which is capable of generating solutions to this type of multi-faceted constrained optimization problem.

Specifically for this project, the point of focus is on the pressing issue of the Nile Basin: Ethiopia, being at the highest point of the river, has constructed a dam - Grand Ethiopian Renaissance Dam (GERD), which generates a lot of crucial energy for the country. Sudan, positioned at the river's upper course as well, has multiple dams so they can control the inflow for their agricultural demand. This leaves Egypt with a water deficiency if its water supply is not controlled accordingly[3]. It has been demonstrated that an agreement between all three countries could result in a more optimal solution, where all countries benefit [19].

Machine learning can improve the ability to derive policies for water resources more effectively in the face of climate change. Reinforcement learning (RL) is suitable for this challenge, as the algorithm is capable of learning optimal policies through iterative interactions with the environment, making it particularly effective for dynamic and complex problems in water management. This overarching project aims to model the problem as a reinforcement learning problem to see if any improvements in efficiency and accuracy can be achieved.

The significance lies not only within this particular problem but could potentially be applied to broader water management crises. The inspiration for this work is from a Master Thesis done by Yasin Sari [12]. The thesis has implemented a model that simulates the Nile Basin conflict. Using the model, this project implements a MONES reinforcement learning algorithm, instead of EMODPS, which has been a key choice algorithm for similar water management problems [21].

The Nile Basin model simulates the flow of water through the Blue Nile, the Main Nile, and ending at the High Aswan Dam (HAD), incorporating key dams like the GERD and HAD, Roseires and Sennar dams. The simulation makes use of predefined data for 8 water catchments and 6 irrigation districts, which represent water inflows and depletions, respectively. It requires both static data, such as water storage descriptions in the reservoir, and dynamic data, namely river inflow rates, requested depletions by countries, and water evaporation rates.

The impact of climate change on this dynamic data is significant[5]. Climate conditions affect river inflows and evaporation rates, which influences water availability. This makes the discussion of climate change-impacted data crucial, as it underlines the need for adaptive water management strategies. It is through the understanding of how climate change may affect water resources we can develop more efficient and robust water management solutions. Thus, this project will use datasets with varying climate conditions, including droughts and floods, to illustrate the extreme climate scenarios that could happen within 20 years.

Given that climate change significantly impacts water management problems, it is crucial to take into account different climate change scenarios when developing solutions for these challenges[9]. The project aims to use two data sets to illustrate plausible scenarios of human-favourable and varying climate change conditions for the Nile Basin's future. These data sets were then used to assess how the reinforcement learning algorithms performed on the Nile Basin simulation. The results showed that both algorithms faced more difficulties when dealing with varying climate scenarios, producing models that posed greater challenges for the countries involved. The results of cross-evaluation of the RL models on the normal and varying climate scenarios highlight the need to take climate change into account, as this would provide a more accurate representation of the problem and improve the efficiency of the solutions that are developed.

2 Methodology

The methodology of this research consists of six key components: gathering the necessary background information needed for this experiment, refactoring of the Nile Basin simulation, data analysis and modelling for the algorithms, integration of the model into MONES RL algorithm, experimental setup and evaluation methods used for assessing solutions' performance.

2.1 Background

EMODPS

EMODPS is a well-established algorithm commonly used in the water management community[6]. This algorithm combines direct policy search and multi-objective evolutionary algorithms to find the optimal set of solutions, otherwise called the Pareto front. By using non-linear approximating networks EMODPS generates policies that can then be used to address water management challenges. While this algorithm performs well on complex reservoir systems, further research is needed to better explore alternative algorithms that could outperform EMODPS[21].

MONES

MONES uses evolutionary processes to a population of neural networks (NNs), aiming to find a set of optimal solutions. In MONES, each network acts as an "agent", which interacts with the environment. The weights of each NN are randomly sampled from a Gaussian distribution. When an agent completes an action or takes a step in the environment, MONES evaluates their performance based on predefined objectives. These evaluations are then used to make adjustments on the NNs' weight distributions. This way the algorithm iteratively improves the population and produces a diverse set of solutions[7].

Gymnasium

Gymnasium is an open-source project that provides an environment to build reinforcement learning algorithms[1]. Within the Gymnasium environment, users define actions that the reinforcement learning agent can take and the corresponding observation spaces. The core interaction is the "step" function, which takes an action as input and produces an observation and a reward based on the objectives. This framework provides users with standardized environments, which allow for reproducibility and comparability. Additionally, there is an equivalent framework built for Multi-Objective reinforcement learning (MORL) algorithms - MO-Gymnasium[2].

MORDM vs. MORO

In the research of Yasin Sari, the posterior multi-objective robust decision making(MORDM) was used[12, 14]. This method involves using the same data for every simulation and then testing with different data that represent climate scenarios. For this project, however, a different climate scenario is chosen randomly before every simulation of the training. This approach is called many-objective robust optimization (MORO)[4]. By training on different possible outcomes of data tailored to extreme climate conditions, it is expected the algorithm to better perform on future scenarios that may impose extreme droughts or floods in the basin.

The Nile Basin simulation

The simulation only includes a single thread of the river (see appendix A), which starts in Ethiopia, passes through Sudan, and ends in Egypt. There are four dams included in the model: GERD in Ethipoia, Roseires and Sennar in Sudan, and HAD in Egypt. Key points of the river are the water inflows: the Blue Nile River, water catchments from GERD to Roseires dams, water catchments from Roseires to Abu Naama dams, water catchments from Suki to Sennar, the Dinder and Rahad rivers, followed by the larger White Nile and Atbara rivers. Additionally, water depletions in the simulation are represented by irrigation districts: Downstream and Upstream Sennar, Gezira, Hassanab, and Taminiat, all located in Sudan, and Egypt as a single irrigation district. Egypt is considered as a single irrigation district as the simulation only includes the Nile until the HAD dam, but Egypt's water resources are a crucial objective. The simulation considers a time step of one month, as this is enough to consider varying seasonal climate conditions, but is not too small as to highly increase the complexity of the simulation[17].

Simulation objectives

Since managing water resources in the Nile Basin simulation consists of balancing a few competing objectives, it is considered an optimization problem. The objectives that are being optimized within this simulation are:

- **Minimize Egypt's average yearly water deficiency**. The lack of water in the irrigation farms throughout the 20 years of the simulation should be minimized.
- **Minimize HAD months below the hydro-power plant threshold**. The amount of months the HAD reservoir level is below the level the hydro-power plant can generate power should be minimized.
- **Minimize Sudan's average yearly water deficiency**. The lack of water in the irrigation farms throughout the whole simulation should be minimized.
- **Maximize Ethiopian yearly hydro-energy**. Ethiopian-produced hydro-power plant energy should be maximized.

2.2 Refactoring the Nile Basin simulation model

The initial Nile Basin simulation model is tightly integrated with the EMODPS algorithm, with the algorithm's code being embedded within the simulation. Additionally, EMODPS generates policies for the whole duration of a simulation, whereas MONES does it a single step at a time. Thus, to simulate the MONES algorithm, the functionality of the existing simulation model had to be extracted. This refactoring was essential, as the original model was tailored specifically for the EMODPS algorithm.

The new version of the Simulation model is designed to fit a bigger variety of water management problems. The generalization of the framework allows specialists to adjust multiple water system models with varying qualities and entities. For instance, the model now accommodates adding water catchments, irrigation systems and power plants, each of which can be configured with varying water demand, supply and energy generations.

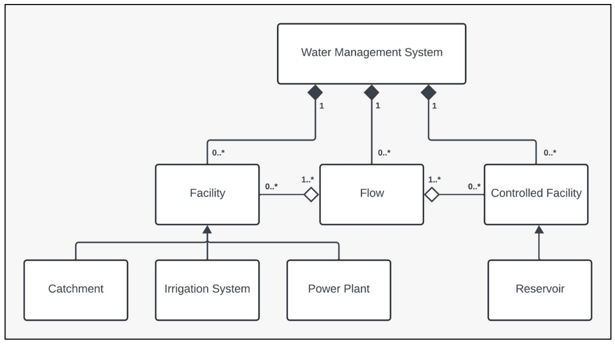

Figure 1: Water management system simulation UML diagram.

The figure above (see fig. 1) displays the general structure of the model. The simulation consists of these modular classes:

- Controlled Facility: this abstraction is used for classes for which there are decisions to be made in the water management systems. There are possibilities for other classes, for example, a controlled water depletion point.
  - Dam: this class is responsible for controlling how much water passes is stored in the reservoir and how much passes downstream.
- Facility: this abstraction is used for classes that work without external control from the outside world.
  - Catchment: this object is responsible for indicating how much water is collected through natural inflows. It is usually indicated statically through a list over a period of time.
  - Irrigation District: this object represents a district which consumes water. In this implementation, it is indicated through a static list.
  - Power Plant: this object calculates the amount of energy the power plant produces given the flow from the dam.
- Flow: this object represents an edge in the network of the water management system, it contains a number which represents the water flow from one facility to another.
- Water Management System: it is the main connecting class that connects the system. It contains the functionality of connecting the flow through all the elements of the system.

Simulation validation

It was necessary to make sure the new model is correct and behaves the same as the reference simulation model. The simulation itself is deterministic, thus with the same context data and actions, the same results will be seen. Firstly, all data was copied from the original simulation model to the new model. Then, using the same machine a set of random actions was generated for 20 years. Throughout the simulation, at every time step Reservoirs' inflow and outflow as well as GERD power-plant energy production were logged. These checkpoints throughout the simulation served as markers, as all other processes depend on this data. Then upon comparing the results of the old and new simulations, we found that they yielded identical results.

2.3 Data analysis and modelling

The data used in this simulation can be separated into two parts, namely one that is affected by the climate and socio-economic factors and one that is independent. The former includes lake evaporation factors, irrigation and civilization water demands, and inflows from catchments in the upstream parts of the basin. The latter data includes water surface area, level, and volume relativity of the dams' reservoirs, and minimum and maximum release policies dependent on the reservoirs' water levels. This subsection further describes the modelling process.

A simplified overview of how the data influences the model is illustrated in Figure 2. The simulation dynamically receives generated data before each run. Based on this data, the simulation and the RL

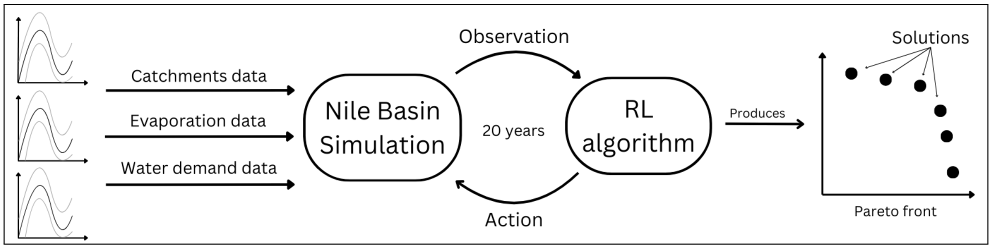

Figure 2: Abstract representation of the algorithm.

algorithm start an iterative process where the simulation provides an observation to the RL agent, and the agent responds with an action to the Nile Basin simulation. Once the iterative process is completed, the RL algorithm produces the Pareto front - the solution set.

Static data

Each of the three countries' authorities imposes minimum and maximum water release policies that allow for control of the dam. These policies control extreme situations, for example, so that the valley does not completely dry out or overflow in case of a flood, and provide a sufficient amount of energy. Additionally, detailed information about the water storage inside the dam's reservoir is needed. Specifically:

- Water surface area: used to calculate the evaporation rates of the water in the reservoir.
- Water level: used to determine energy productions from the hydropower plant and implement the minimum and maximum water release policies.
- Water volume: crucial for the river's water flow process throughout the simulation.

This data was adopted from the reference Nile Basin simulation.

Climate and Socio-Economic Dependent Data

Evaporation rates, water demands and inflows from the smaller parts of the river are highly uncertain, and with climate change, increasingly difficult to control. This subsection describes the modelling of the data that is influenced by external factors and explains how the uncertainty was applied to the dataset.

Evaporation factor

Evaporation rates are taken from Kevin G. Wheeler's simulation of the Nile Basin [19]. To account for increasing temperatures due to climate change, the relation found between the increasing temperature and the evaporation factors in Lake Qarun in Egypt has been taken[13]. It describes the relationship as a linear function between the temperature of the surface and evaporation. This acted as a guideline when creating the climate models for the simulation. It described a mean increase in evaporation rates as 0.3% yearly from 1980 to 2019. As temperature increase is forecasted to increase in the future, the increase rate yearly was taken as 0.4%. Uncertainty was injected into this dataset by getting a value from a normal distribution for each year, and applying a log transformation with it, which allows for randomness and non-linear distortions. Since Wheeler described the data in 2018, The same transformation has been applied to start the new data from 2024.

Water inflows

Kevin G. Wheeler et al. have described and created 3 artificial datasets with different climate scenarios: baseline, with 15% ias and 20% Hurst coefficients for uncertainty [19]. 15% ias coefficient implies 15 % interannual standard deviation between the years 2018 and 2060, whereas 20% Hurst coefficient implies a standard deviation increase at each inflow. Each dataset has 100 traces describing different climate outcomes, with each trace having 50 years of data for 162 catchments. For simplicity of this project, this dataset has been reduced from 162 catchments to 8 and we act as there are 8 big catchments in the Nile Basin (Blue Nile, Atbara, White Nile, Rahad, Dinder rivers, Gerd to Roseires, Roseires to Abu Naama, Suki to Sennar river catchments). These traces are going to be used for the training of the reinforcement learning models, to increase robustness, according to the MORO. For this project, the baseline and 20% Hurst coefficient datasets are used to represent favourable climate conditions as well as the more extremes, which are similar to those described by Strzepek et al.[16].

Table 1: Standard deviation and mean difference of Hurst and Baseline datasets, measured in %.

Region	$\frac{(\text{std}_{\text{Hurst}} - \text{std}_{\text{Baseline}}) \times 100\%}{\text{std}_{\text{Baseline}}}$	$\frac{(\text{mean}_{\text{Hurst}} - \text{mean}_{\text{Baseline}}) \times 100\%}{\text{mean}_{\text{Baseline}}}$
Atbara	7.58	0.02
Blue Nile	11.34	-1.67
Dinder	12.66	-4.14
Gerd To Roseires	31.15	-3.56
Roseires to Abu Naama	0.95	-3.62
Suki To Sennar	0.00	0.00
White Nile	2.63	8.12

In Table 2, the differences between the Hurst and Baseline datasets are displayed in terms of standard deviation and mean differences. These metrics were calculated by determining the standard deviation and mean for every month across the 8 catchment districts and then averaging these values over the 20-year simulation period. The standard deviation metrics were consistently higher of the Hurst data set. This indicated that the inflows were more variable, presenting an increased possibility for flood and drought conditions. A negative mean difference implies reduced water resources in the simulation in the Hurst dataset, whereas a positive mean difference indicates an increase.

Notably, in the case of the Roseires to Abu Naama catchment, both the standard deviation and mean are 0.0. This indicates the datasets were the same. This anomaly could be due to a computational error or a faulty initial dataset. Given that the Roseires to Abu Naama catchments contribute only 0.196% of the total water resources on average throughout the simulation, the cause of this anomaly was not further investigated due to limited time resources.

Water demands

To account for water demands we have taken 8 different irrigation districts and their water demands as Kevin G. Wheeler et. al. model has outlined [18]. This data contains 8 different irrigation districts in Sudan and one as a whole in Egypt. Since 2 water depletion districts were outside of the river line, the water demand was subtracted from the inflow data of that river line. This data is highly dependent on population, thus we will project it according to the residents' planned increase. Uncertainty was added by generating noise randomly from a normal distribution, with a standard deviation specified by the user. This noise is applied annually, this introduces variability. Population growth and increased water needs were represented by a 2% mean increase every year, taking into account the 0.1–4.5% annual growth economic range of scenarios considered[10].

2.4 Integration into MONES

To make the integration of the simulation easier, we used the MO-Gymnasium framework. To use our model with MO-Gymnasium environment, our model had to inherit from the Environment class of the said framework.

Our implementation integrates the inheritance by all facilities and the main water management system having 'step' and 'reset' methods. This ensures that all actions can be executed through the main structure of the simulation. The 'step' method represents one iteration of the simulation and allows it to progress. The step action produces rewards on which the algorithm evaluates and produces actions for the next step. The 'reset' method re-initializes the simulation and sets the states to their starting state.

The 'step' function takes an action as input and returns a reward, an observation and booleans whether the simulation is finished or has to be truncated. The action is determined by the machine learning algorithm. The reward can be determined by any class, for example, for a dam it might be a minimum level of water, or irrigation farms the satisfaction of a demand. Classes like catchments usually do not have objectives, so no object can be returned. Observation is data on which the machine learning algorithm decides its next action, in other words - the state of the simulation. Similarly, it can be tailored to various cases, in ours specifically it is the water storage in the dam. While we don't use termination and truncation in our model, it can be determined by for example a dam overflowing.

2.5 Experimental setup

The EMODPS algorithm was used to analyze the impact of climate scenarios on water demand, evaporation and water inflows to the basin. The algorithm was run twice: using baseline, human-favourable data and using extreme climate conditions (further referred to as baseline and Hurst runs). For water catchments, pre-generated datasets were used with Baseline conditions and 20% Hurst coefficient uncertainty. For both cases, it was assumed that water demand would grow 2% annually, with a standard deviation of 1% to account for outside factors. Lastly, evaporation rates were set to have a yearly increase of 0.4%. For baseline conditions, the evaporation rates had a noise that followed a normal distribution of mean 0.5 and standard deviation of 0.5. For varying climate conditions the numbers were 1 and 5 respectively, introducing significant variability.

To ensure convergence of the algorithm, each run had 70,000 iterations. Taking into account the computational complexity of the training, the runs were done on the DelftBlue supercomputer with 48 CPU cores. The process for each run took approximately 21 hours, compared to an estimated 90 hours on a standard 6-CPU core laptop.

2.6 Solutions' evaluation

To evaluate the overall performance of the solutions, their efficiency was assessed on all the climate scenarios used in this project. This approach gave insight into how robust the solutions were.

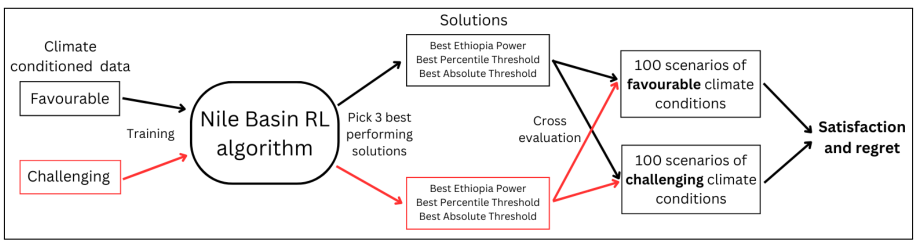

Figure 3: Evaluation process using satisfaction and regrets metrics.

The evaluation, outlined in Figure 3, began by picking out solutions that seemed best from the training, based on their return results of the algorithm. This step was inserted to reduce the space of solutions to be evaluated. This reduction was necessary due to a constraint on available computational resources. For each of the two models, three solutions were picked:

- **Best Ethiopia Power** This solution performed best in terms of energy production for Ethiopia, which is crucial as Ethiopia's dam purpose is to maximize energy production for the country.
- **Best Percentile Threshold** This solution has all the objectives exceed a certain percentile threshold, where the 100th percentile is the best. This percentile threshold in this project varied from the 40th to the 50th percentile.
- **Best Absolute Threshold** This solution is superior when all the objectives are normalized and their combined sum exceeds a specific threshold. The threshold values in this project were between 0.8 and 0.9.

Next, these three solutions from the models trained on human-favourable and climate-varying data—totalling six solutions— were evaluated on all 100 scenarios of the two datasets separately. This evaluation gave insight into satisficing and regret metrics [20]. The satisficing-based robustness metric represents how well the model performs on baseline data, while regret-based robustness tries to minimize the deviations in performance under more challenging climate conditions.

3 Results

This section presents the analysis of the results obtained from the EMODPS and MONES algorithms under different scenarios. It includes the convergence behaviour, trade-offs in the solutions, and the cross-evaluation of EMODPS between the baseline and Hurst datasets.

To know that the machine learning algorithm has stabilized and the results are reliable for decision-making, convergence behaviour is necessary. MORL convergence can be quantified using metrics like ε-progress and hypervolume[11, 7]. The ε-progress is a metric for evaluating the improvement of the solutions over iterations, whereas the hypervolume metric assesses the volume of the objective space created by the Pareto front. Beyond convergence, it is necessary to determine which solutions are optimal according to the metrics defined.

In MORL, understanding trade-offs between solutions is essential, as real-world problems often include conflicting objectives. The decision-makers have to understand the connection between these different objectives to make an informed decision. Examining trade-offs allows for understanding how an improvement of one objective may lead to a deterioration of another. This allows for a more balanced and prioritized solution selection.

3.1 EMODPS results' analysis

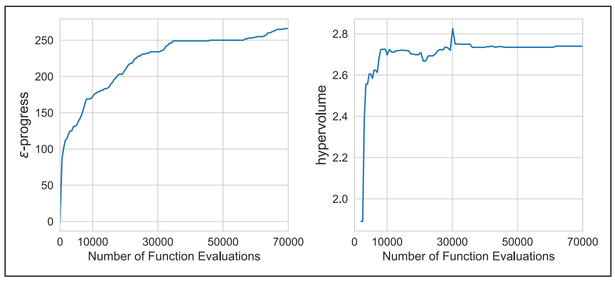

Figure 4: EMODPS convergence graphs for the baseline run.

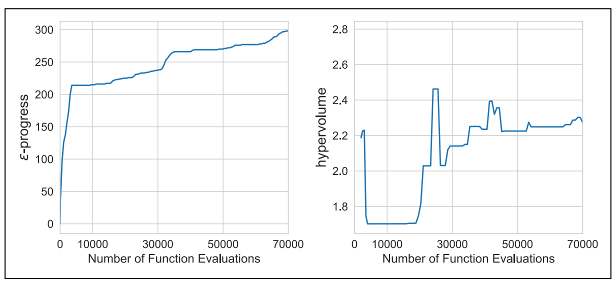

Figure 5: EMODPS convergence graphs for the Hurst run.

In the figures above (see fig. 4, 5) epsilon progress and hypervolume metrics were used to check the convergence of the algorithms. For the case of the baseline Run, both ε-progress and hypervolume indicated quick convergence, and the algorithm could have been terminated after 50,000 iterations without a significant loss of accuracy. The Hurst run, however, has displayed more challenges converging. The ε-progress displayed a slow increase after rocketing in the beginning, while the hypervolume metric did not reach convergence, but rather showed an upward trend. This shows, that the algorithm struggled with the Hurst run which includes the increased variability and randomness of the data.

In Figure 6 the scores of all the considered solutions are plotted for four categories of scores: "Egypt Irr. Deficit" - water deficit in Egypt measured in billion cubic meters(BCM) per year, "Egypt Low HAD" - HAD reservoir being below energy making threshold, measured in %, "Sudan Irr. Deficit" - water deficit in Sudan measured in BCM/year and " Ethiopia Hydropower" - the amount of energy Ethiopia makes per year, measured in measured in terawatt-hours(TWh) per year. Of all the solutions, 6 of the solutions are highlighted and assigned a description and a colour. The first of these four solutions are "Best Egypt Irr", "Best Egypt HAD", "Best Sudan Irr" and "Best Ethiopia Hydropower". These four solutions each correspond to a solution that would be optimal for maximizing only one relevant objective. The other two highlighted solutions display compromise between the relevant objectives. These solutions are "Compromise Percentile Threshold" and "Compromise Absolute Threshold". The former solution displays a policy that makes all objectives better than the 40th

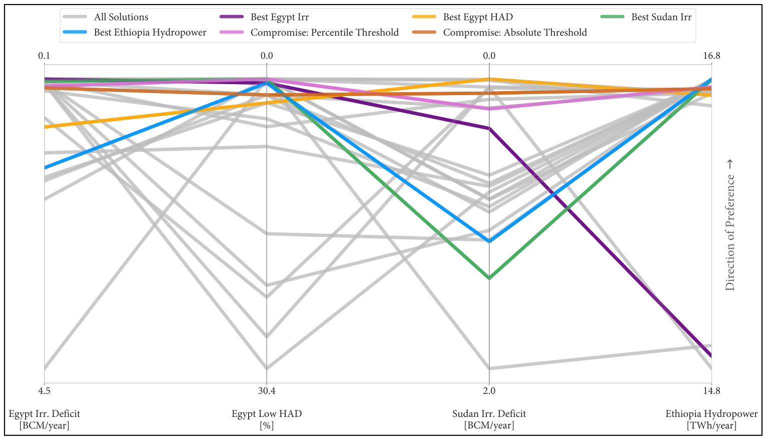

Figure 6: EMODPS solutions for the baseline run.

percentile, where the 100th percentile is the most desirable. In contrast, the latter solution gets a value that is bigger than 0.88 when the outcomes for the objectives are normalized.

The EMODPS run on baseline scenarios found in a total of 24 solutions. A quick breakdown of the notable solutions is as follows:

- **Best Ethiopia Power** This is the best case found for Ethiopia power generation, making 16.81 TWh of energy yearly. This is important, as it maximizes the power made for Ethiopia. Note that a need for larger power generation in Ethiopia was the initial cause of the water use conflict under consideration. It leaves Egypt with approximately 1.44 billion cubic meters (BCM) and Sudan with around 1.12 BCM of water per year. Additionally, over a 20-year period, there would be one month when Egypt would have the water level in the HAD reservoir under the energy generation threshold for the hydro-power plant.
- **Best Percentile Threshold** In this solution, Ethiopia loses around 0.6 TWh, still making 16.75TWh per year compared to the solution of maximizing Ethiopia power. However, the reallocation of water to Egypt and Sudan allowed by this solution leads to decreased water shortages of only 0.2 BCM and 0.23 BCM respectively. HAD levels stay above the threshold throughout 20 years.
- **Best Absolute Threshold** This solution proposes 16.75 TWh energy production for Ethiopia and 4 months for HAD under the threshold. The trade-off seems to happen for Egypt and Sudan water depletions, producing 0.23 and 0.12 BCM water shortages.

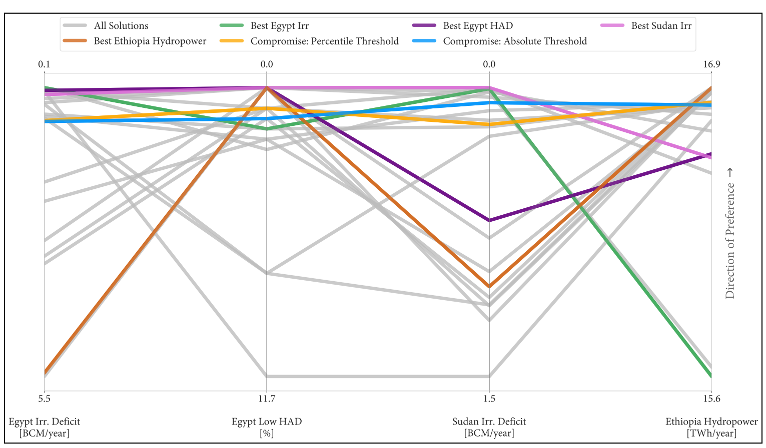

Figure 7: EMODPS solutions for the Hurst run.

Similarly to the baseline run, the Hurst run produced 21 solutions, which are displayed in Figure 7. These are the three solutions picked the same way as the baseline run:

- **Best Ethiopia Power** The best case scenario for Ethiopia makes 16.93 TWh of energy yearly, as well as keeping HAD levels above the threshold throughout the year. However, this leaves Egypt and water with high shortages, being 5.44 BCM and 1.04 BCM yearly respectively.
- **Best Percentile Threshold** This scenario makes almost the same amount of energy for Ethiopia, being 16.86TWh but leaves HAD levels below the threshold two months in 20 years. The water depletion lack in Egypt and Sudan is 0.71 and 0.21 BCM.
- **Best Absolute Threshold** Absolute threshold scenario offers 16.85 TWh of energy and 3 months of HAD being below the threshold. Egypt has a 0.72 BCM water shortage and Sudan has 0.095 BCM of water per year.

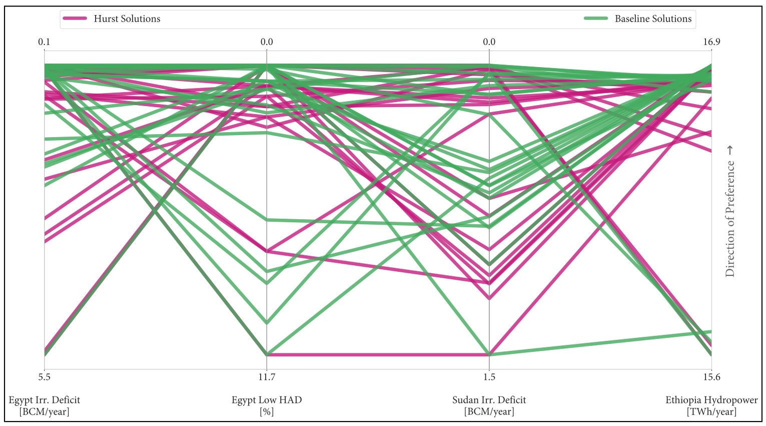

Figure 8: Solutions of Baseline and Hurst runs combined.

In Figure 8 all solutions of both algorithms are displayed. Both runs have high variability in their solutions. Significantly better performance of one run is not seen over another, however, green lines representing baseline solutions seem to assume the top positions of the graph. This shows that the solution of the Baseline run better fulfils the objectives. Taking the hypervolumes of the runs, Baseline has 2.78, whereas Hurst has 2.47.

3.2 MONES results' analysis

Due to the lack of computational resources, only the performance of the Hurst run on the MONES model was executed.

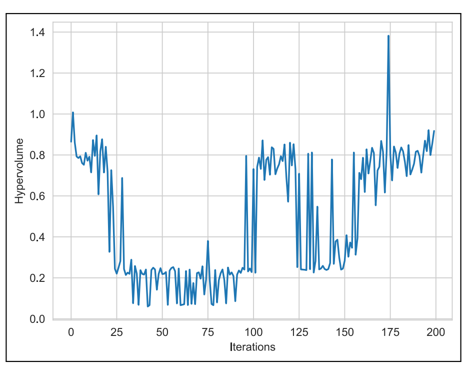

Figure 9: MONES hypervolume trend over iterations for the Hurst run.

Figure 9 displays the hypervolume trends over iterations of the MONES algorithm using Hurst data, which represent varying climate conditions. The graph indicates sporadic jumps in the hypervolume.

A slight overall trend upward can be observed, however, it cannot be said that the graph converged. This indicates that the model did not fully learn, or there were not enough iterations for the algorithm to fully converge. The overall hypervolume was on average lower than that of the EMODPS run.

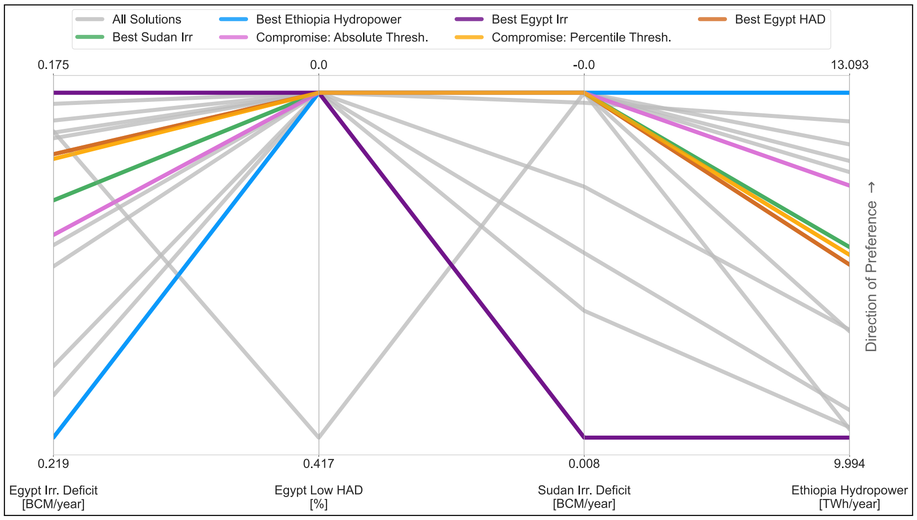

Figure 10: MONES solutions for the Hurst run.

As for EMODPS training, a parallel coordinates plot (see Figure 10 was used to depict 15 solutions that MONES found. Similarly to EMODPS, the same logic was applied when picking the three notable solutions:

- **Best Ethiopia Power** The best case scenario for Ethiopia produces 13.09 TWh energy per year, leaving HAD dams above the energy production threshold throughout the simulation. This solution minimizes the water deficit in Sudan to 0, whereas for Egypt there is a 0.22 BCM water deficit per year.

- **Best Percentile Threshold** This scenario decreases the energy production for Ethiopia to 12.26 TWh per year, still leaving HAD reservoir levels above the threshold of energy production. In this scenario, Sudan is provided with a sufficient amount of water, however, Egypt has a deficit of 0.19 BCM per year.

- **Best Absolute Threshold** Absolute threshold scenario offers 11.64 TWh energy for Ethiopia per year, also keeping the HAD levels above the threshold. The water deficit for Sudan is again 0, with Egypt being 0.18 BCM per year.

3.3 EMODPS and MONES comparison

Comparing the performance of the EMODPS and MONES algorithms, some differences emerge when it comes to their ability to handle varying data scenarios, particularly on the Hurst dataset. EMODPS produced a robust solution set, displaying clear trade-offs between objectives, on both the baseline and Hurst runs, which achieved high energy production for Ethiopia while having low water deficit levels for Sudan and Egypt. Additionally, EMODPS kept the hypervolume trend upward throughout the training. In contrast, MONES struggled more visibly with the Hurst dataset, showing sporadic hypervolume jumps and failing to converge. MONES achieved a notable water deficit reduction for Sudan and Egypt, simultaneously keeping HAD water levels above the threshold throughout the whole simulation, albeit with lower energy outputs for Ethiopia. The low hypervolume together with the overall solution set objectives shows a narrow Pareto front, with a less trade-off compromises. This comparison highlights EMODPS's adaptability and efficiency in optimizing multiple objectives, producing a wide Pareto front.

3.4 EMODPS solutions' evaluation

Three chosen solutions, namely Best Ethiopia Power, Best Percentile Threshold and Best Absolute Threshold, from each training of baseline and Hurst datasets, were cross-evaluated on both datasets in terms of all 100 scenarios.

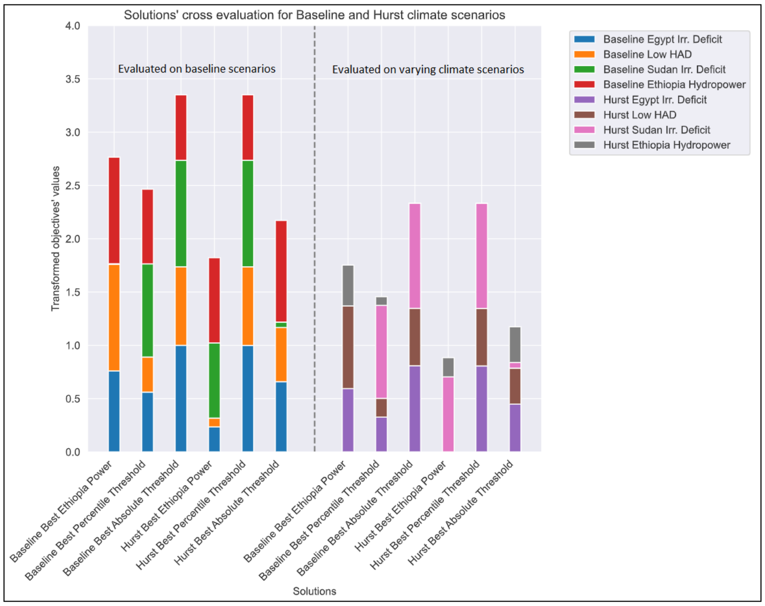

Figure 11: EMODPS solutions' cross evaluation results.

The graph in figure 11 presents the evaluations of the 6 solutions produced by the EMODPS algorithm. All the values were normalized based on the minimum and maximum values of the objectives. Additionally, the objectives for Egypt's water deficit, HAD level and Sudan's water deficit were inverted so that all objectives are on the maximization scale, making the tallest bar present the best solution. The left side of the chart shows the average performance of the solutions across 100 baseline scenarios, while the right side shows the average performance across 100 challenging climate scenarios.

For both sets of tests, the absolute best solution was tied between those trained on Hurst and Baseline data. In the baseline evaluations, the mean performance of solutions trained on baseline scenarios was higher than that of solutions from Hurst training. The same pattern can be observed in evaluations of the Hurst climate scenarios. Importantly, the overall optimization of the objectives in the Hurst evaluation, regardless of the training data, was lower than that of baseline scenarios.

Examining the overall performance of the models in the evaluation (see Appendix B), a significant decrease in performance is noticeable compared to the training results. The best two solutions evaluations from both baseline and Hurst climate conditions compared with their respective training outputs:

- **Trained on Baseline, tested on Baseline - Best Absolute Threshold** compared to Best Absolute Threshold solution from the Baseline training: A decrease of 1.6 TWh in energy generated by GERD, and an increase in water deficit for Egypt and Sudan by 1.51 BCM and 0.2 BCM, respectively. There was a major increase of 44 months for HAD being below the threshold.
- **Trained on Hurst, tested on Baseline - Best Percentile Threshold** compared to Best Percentile Threshold solution from the Baseline training: The values changed by almost the same amount as in the previously discussed solution. The models trained similar policies that ended up with similar results.
- **Trained on Baseline, tested on Hurst - Best Absolute Threshold** compared to Best Absolute Threshold solution from the Hurst training: A reduction of 1.8 TWh in energy

for Ethiopia, and an increase in water deficit for Egypt and Sudan by 1.41 BCM and 1.13 BCM, respectively. There was an increase of 54 months in which HAD levels were below the threshold.

- **Trained on Hurst, tested on Hurst - Best Percentile Threshold** compared to Best Percentile Threshold solution from the Hurst training: A reduction of 1.81 TWh in energy for Ethiopia, and an increase in water deficit for Egypt and Sudan by 1.42 BCM and 0.12 BCM, respectively. There was an increase of 52 months for HAD being below the threshold.

4 Responsible Research

This project was completed adhering to the principle of responsible research. The focus was on transparency, reproducibility, and ethical considerations throughout the whole process. All data used within this project is sourced and taken from reputable papers, which are aimed at achieving transparency and maintaining a standard of scientific integrity. This allows others to verify the sources and methods independently. To make the project reproducible, a detailed process is described within this document, along with providing the code that was used for obtaining the results. This enables other researchers to replicate this study. Additionally, the author has taken measures to clearly outline potential biases in the process of this project as well as the data used, ensuring that any influencing factors are clearly defined and accounted for. Lastly, the author declares no financial or personal affiliations with any parties that could have influenced the project. This is aimed to produce unbiased results.

All the data and code can be accessed via `https://github.com/lukavicius/cse3000_bachelor_thesis`.

5 Discussion

This research aimed to provide insights into the impact of data and its representation for future scenarios in the context of using different reinforcement learning algorithms for water management, specifically in the Nile Basin. Some constraints and limitations need to be acknowledged to comprehend the constraints of the results.

Firstly, only a single line of the river was taken to represent the whole Basin in the simulation. A larger scale of the simulation, taking into account the broader context of the problem, could offer a more accurate solution. The current simulation lacked accuracy in terms of other irrigation districts that were on the other lines of the river. Additionally, several dams on the considered line of the river were not considered, which could have influenced the results. Meteorological factors, such as evaporation factors in the rivers were not taken into account, potentially affecting the water flow and general water resources availability. Furthermore, this study focused on four main objectives, which can be insufficient to precisely represent the multifaceted discussions countries have to partake in when managing real-life water management challenges.

The training of the MONES algorithm showed results that were significantly lower than those of EMODPS. This could have been due to the fewer iterations used to train the model. The hypervolume metric did not converge by the end of the training, indicating that the algorithm could have learnt more with additional iterations. Another potential improvement for the MONES algorithm is to explore a wider solution set, thus allowing for bigger trade-offs for the optimization problem, and thereby providing decision-makers with a better idea of the possibilities.

Moreover, the data that was used for this research was mostly taken from the previous century to project future scenarios, which is likely to decrease the accuracy of the algorithm in terms of rapidly changing climate conditions. Although transformations which were applied to account for the time difference tried to mitigate that effect, the transformations may not directly represent the climate. Similarly, the rough estimates of water depletions in the future do not exactly represent the actual increase in demand.

The EMODPS algorithm used in this project performed better in terms of hypervolume than that of Yasin Sari, likely due to the change of the data in the catchments [12]. In Sari's research, the same 8 catchments outlined by Wheeler were used, but only the largest points of Wheeler's grouping were considered[19]. This project mapped all of the points into 8 catchments, effectively making the water resources larger. This allowed the algorithm to perform better on the objectives, showcasing an overall higher hypervolume.

Another improvement that made a significant impact on this project was the use of different simulation data for each run. This allowed the algorithm to learn from different scenarios the simulation could encounter, which enhanced its robustness. However, this approach simultaneously increased the complexity of the algorithm. While the baseline run of the algorithm converged quickly, the Hurst run with varying climate conditions did not reach full convergence. This issue could potentially be mitigated by averaging over a few runs of the simulation when evaluating the performance during the training.

The evaluation of the EMODPS solutions displayed substantial performance differences when tested on different climate scenarios. This highlights the importance of using different climate scenarios, especially those that represent climate change, in the training and evaluation of our models. The unpredictable effects of climate change lead to a wide variety of future scenarios, making it essential to consider multiple different climate scenarios. Interestingly, models which were trained on baseline data seemed to perform better on unseen varying climate conditions than those trained on it. This is possibly due to bias introduced by the traces. Since the output results only take the results of the last iteration, it could happen that the last iteration had a more favourable scenario, making the results better. This would suggest that a solution that possibly was less optimal had higher performance metrics.

Overall, this study has made milestones in improving the water management simulation for the Nile Basin, future work should address the above-mentioned limitations when developing more accurate and reliable models.

6 Conclusion

This research had two main objectives: refactoring the simulation to fit into MONES or any other Markov decision-making process-based reinforcement learning algorithm and testing EMODPS on the simulation with two types of data sets. Both aims of the research were successfully achieved, which provided insights on the application of RL in water management problems.

The simulation has been successfully refactored and integrated into the MONES framework. This proves that advanced RL algorithms can be applied to optimize complex water management issues. While MONES did not perform well as compared to EMODPS, further research and experimentation are needed to make definitive conclusions about MONES' efficiency.

When the EMODPS algorithm was applied to two different data sets, the solutions of the varying climate conditions faced significant challenges in finding solutions that were as optimal as those trained on the better climate conditions. This underscores the need to improve the training process for such scenarios, as they represent the future climate conditions better. This is important as once these climate conditions are encountered, the policy chosen by decision-makers might fail to efficiently utilize the resources. It is through the incorporation of diverse climate scenarios that the optimization solutions will be robust and efficient for real-world applications.

Further research in the field has to enhance the accuracy and reliability of water management simulations. This can be done by including a bigger context in the simulation. More meteorological factors such as river evaporation rates should be considered. A better representation of objectives is necessary, as it could capture the complexities of real-world decision-making processes. Utilizing more accurate and recent data to reflect and project the climate conditions will improve algorithms' relevance and accuracy.

This project is aimed at improving water management by exploring different machine-learning algorithms. The goal of this project is to assess whether these algorithms can be efficiently used for complex water management problems, rather than be used as a direct controlling policy in real-life situations. Incorporating varying climate conditions into our simulations is not intended to alarm the public about the future, but rather an encouragement and spotlight on the ongoing problems. By highlighting potential climate change scenarios it is aimed to draw attention to the urgent need for solutions and higher research efforts to tackle ongoing and possible future water management challenges.

In conclusion, this study highlights the potential of improving the machine algorithms used in water management, especially in the face of climate change. It spotlights the necessity for robust and flexible approaches to modelling and resolving water management challenges, to ensure a more effective and sustainable water resource management in the future.

A Nile Basin simulation overview

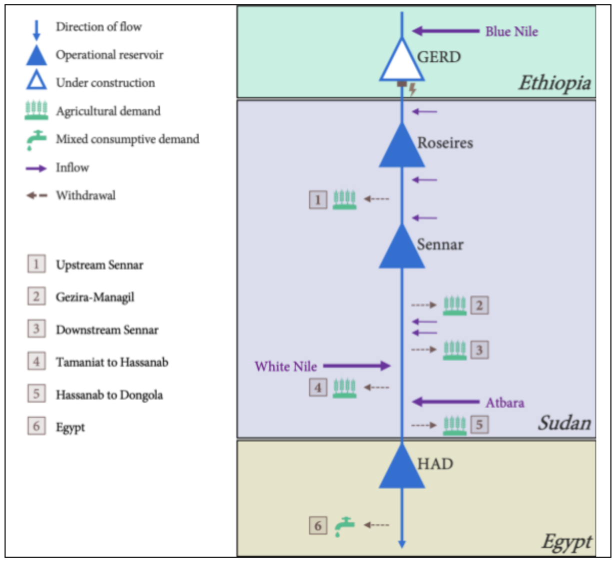

Figure 12: Nile Basin simulation overview. Source: Adapted from [12]

.

B EMODPS Solutions' Evaluation results

Table 2: EMODPS solutions' cross evaluation results.

Solution	Egypt Irr. Deficit BCM/year	Egypt Low HAD, %	Sudan Irr. Deficit BCM/year	Ethiopia Hydropower, TWh
BonB Best Ethiopia Power	2.23	0.14	1.27	15.21
BonB Best Percentile Threshold	2.62	0.27	0.44	15.16
BonB Best Absolute Threshold	1.74	0.19	0.32	15.15
BonH Best Ethiopia Power	2.56	0.18	1.27	15.11
BonH Best Percentile Threshold	3.09	0.3	0.44	15.06
BonB Best Absolute Threshold	2.13	0.23	0.34	15.05
HonB Best Ethiopia Power	3.28	0.32	0.6	15.18
HonB Best Percentile Threshold	1.75	0.19	0.32	15.15
HonB Best Absolute Threshold	2.43	0.24	1.23	15.2
HonH Best Ethiopia Power	3.75	0.34	0.6	15.07
HonH Best Percentile Threshold	2.13	0.23	0.34	15.05
HonH Best Absolute Threshold	2.85	0.27	1.22	15.01

BonB - Solution trained on Baseline data, tested on Baseline scenarios
BonH - Solution trained on Baseline data, tested on Hurst scenarios
HonB - Solution trained on Hurst data, tested on Baseline scenarios
HonH - Solution trained on Hurst data, tested on Hurst scenarios

References

[1] Gymnasium documentation. `https://gymnasium.farama.org/`, 2023.

[2] Mo-gymnasium documentation. `https://mo-gymnasium.farama.org/`, 2023.

[3] M. Ahmed, R. Abdelrehim, M. Elshalkany, and M. Abdrabou. Impacts of the grand ethiopian renaissance dam on the nile river's downstream reservoirs. *Journal of Hydrology*, 633:130952, 2024.

[4] E. Bartholomew and J. H. Kwakkel. On considering robustness in the search phase of robust decision making: A comparison of many-objective robust decision making, multi-scenario many-objective robust decision making, and many objective robust optimization. *Environmental Modelling Software*, 127:104699, 2020.

[5] T. A. Gado and D. E. El-Agha. Climate change impacts on water balance in egypt and opportunities for adaptations. In M. Abu-hashim, F. Khebour Allouche, and A. Negm, editors, *Agro-Environmental Sustainability in MENA Regions*, Springer Water. Springer, Cham, 2021.

[6] M. Giuliani, A. Castelletti, F. Pianosi, E. Mason, and P. M. Reed. Curses, tradeoffs, and scalable management: Advancing evolutionary multi-objective direct policy search to improve water reservoir operations. *Journal of Water Resources Planning and Management*, 142(2):04015050, 2016.

[7] C.F. Hayes, R. Rădulescu, E. Bargiacchi, and et al. A practical guide to multi-objective reinforcement learning and planning, 2022.

[8] M. E. McClain. Balancing water resources development and environmental sustainability in africa: A review of recent research findings and applications. *AMBIO*, 42(5):pp. 549–565, 2012.

[9] S. Melaku. Effect of climate change on water resources. *Journal of Water Resources and Ocean Science*, 5:14, 02 2016.

[10] C.A. Nikiel and E.A.B. Eltahir. Past and future trends of egypt's water consumption and its sources. *Nature Communications*, 12, 2021.

[11] N. Riquelme, C. V. Lücken, and B. Baran. Performance metrics in multi-objective optimization. In *2015 Latin American Computing Conference (CLEI)*, pages pp. 1–11, 2015.

[12] Y. Sari. Exploring trade-offs in reservoir operations through many objective optimisation: Case of nile river basin. Master's thesis, Delft University of Technology, 2022.

[13] E. Seif, H. Shokr, A. M. Youssef, M. Sharaf Eddeen, Y. Diab, and Y. E. Imam. Climate change effect on water temperature and evaporation in lake qarun. *Water Science*, 37(1):1–10, 2023.

[14] B. Shavazipour, J. H. Kwakkel, and K. Miettinen. Multi-scenario multi-objective robust optimization under deep uncertainty: A posteriori approach. *Environmental Modelling Software*, 144:105134, 2021.

[15] P. Stott. How climate change affects extreme weather events. *Science*, 352:pp. 1517–1518, June 2016.

[16] K.M. Strzepek, D.N. Yates, and D.E. El-Quosy. Vulnerability assessment of water resources in egypt to climatic change in the nile basin. *Climate Research*, 6:89–95, 1996.

[17] L. Votruba. *"6 Simulation Models of Water Resource Systems"*, volume 32 of *"Developments in Water Science"*. Elsevier, 1988.

[18] K. G. Wheeler, M. Basheer, Z. T. Mekonnen, S. O. Eltoum, A. Mersha, G. M. Abdo, and S. J. Dadson. Cooperative filling approaches for the grand ethiopian renaissance dam. *Water International*, 41(4):pp. 611–634, 2016.

[19] K. G. Wheeler, J. W. Hall, G. M. Abdo, S. J. Dadson, J. R. Kasprzyk, R. Smith, and E. A. Zagona. Exploring cooperative transboundary river management strategies for eastern nile basin. *Water Resources Research*, 54:pp. 9224–9254, 2018.

[20] D. Yan, F. Ludwig, H. Q. Huang, and S. E. Werners. Many-objective robust decision making for water allocation under climate change. *Science of The Total Environment*, 607-608:294–303, 2017.

[21] J. Z. Salazar, P. M. Reed, J. D. Herman, M. Giuliani, and A. Castelletti. A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances in Water Resources*, 92:172–185, 2016.