Università degli Studi di Genova

Scuola Politecnica & Facoltà di Ingegneria

Laurea Magistrale in Ingegneria Strategica

International MSc in Engineering Technology for Strategy and Security of Genoa University

Optimizing Hydropower Scheduling in Competitive Electricity Markets through Deep Learning

Supervisor : Prof. Hossein Farahmand, NTNU Norway

Co-supervisor: Prof. Francesco Bellotti, Unige

Candidate : Rashid Miskeen (S5152104)

Academic Year: 2025/2026

Abstract

Global energy demand is growing rapidly. This leading to a greater focus on renewable energy sources. Hydropower is an integral part of the global renewable energy sector, which provides reliability, flexibility and storage capability in power generation. However, with growing market competition and the integration of variable renewable energy sources, optimizing hydropower scheduling has become a complex challenge. Traditional scheduling techniques suffer from computational inefficiencies, making them less responsive to real-time market conditions. In this research, we are using the application of machine learning models, including recurrent neural networks (RNN) and long-short-term memory (LSTM) networks, to better predict inflow, electricity prices, production and demand. The results indicate that the LSTM model is better at predicting for long-term trends and seasonal changes, which leads to more accurate forecasts and takes less time to calculate, based on the evaluation of various performance metrics such as NSE, MAE, MSE, RMSE and MAPE. These results support the integration of advanced machine learning methods for more economical and operationally efficient hydropower management in competitive electricity markets.

Preface

This master thesis is written for the course TET4900 Electric Power and Energy Systems, Master's Thesis at NTNU the spring semester of 2025.

I extend my gratitude to my supervisors, Prof. Hossein Farahmand and Prof. Francesco Bellotti, Unige, for their guidance and constructive feedback throughout my journey. Additionally, I would like to thank my co-supervisor, Ph.D. Candidate Rahmathulla Madathi Parambath for his opinion and contributions in providing the necessary links for finding online data for my thesis. His support has been very helpful in writing this master thesis.

Rashid Miskeen
Trondheim, June 2025

Contents

List of Figures

List of Tables

Abbreviations

List of all abbreviations in alphabetic order:

- **IEL** Department of Electric Energy
- **NTNU** Norwegian University of Science and Technology
- **AI** Artificial Intelligence
- **ML** Machine learning
- **ANN** Artificial Neural Networks
- **FNN** Forward neural networks
- **RNN** Recurrent Neural Network
- **LSTM** Long Short Term Memory
- **NSE** Nash–Sutcliffe Efficiency
- **MAE** Mean Absolute Error
- **MSE** Mean Squared Error
- **RMSE** Root Mean Squared Erro
- **MAPE** Mean Absolute Percentage Error
- **LTHS** Long Term Hydropower scheduling
- **LTHPS** Long-term hydroelectric power scheduling
- **SVR** Support vector regression
- **SDP** Stochastic Dynamic Programming

- **SDDP** Stochastic Dual Dynamic Programming
- **MILP** Mixed-Integer Linear Programming
- **No.3 Zone** Central Norway
- **EMPS** Elektrisitetsforsyningens Forskningsinistitutt's Multi-Area Power Market Simulator
- **NVE** Norwegian Water Resources and Energy Directorate
- **GPU** Graphics Processing Unit
- **TPU** Tensor Processing Unit
- **CPU** Central Processing Unit
- **GRU** Gated Recurrent Unit

Chapter 1

Introduction

We start the chapter with some background information about hydropower scheduling and our motivation for this project in Section 1.1. In Section 1.2 we define our research goal and research questions that can help achieve the goal. The main contributions are listed in Section 1.3 and the structure of the thesis is laid out in Section 1.4.

1.1 Background and Motivation

Hydropower is a method to generate electrical energy from flowing water and is the most widely used energy source in Norway with 136.49 terawatt-hours of yearly production out of a total of 156.1 terawatt-hours in a normal year (NVE,2023), which is about 88 % of total power production in Norway from hydropower [1]. Norway has a rich history in hydropower scheduling, and the Norwegian professional community is a leader in this field [2]. Currently, the EMPS model (Elektrisitetsforsyningens Forskningsinistitutt's Multi-area Power Market Simulator) is widely used as a decision support tool among players in the Nordic power market due to its suitability for hydropower systems [3].

Hydropower scheduling is uncertain due to various factors, with reservoir inflow being a primary contributor. Historical data serves as the foundation for predicting future inflows, and access to extensive reservoir data, as is the case in Norway, offers a significant advantage [4]. This large amount of reliable data allows for more accurate predictions. However, managing and processing large datasets alongside the inherent uncertainties in the system presents computational challenges.

Inflow volumes are strongly influenced by annual precipitation levels. As illustrated in Figure 1.1.1, historical precipitation data is presented for NO.3 zone Norway, with yearly values expressed as a percentage deviation from the average 2015-2023.

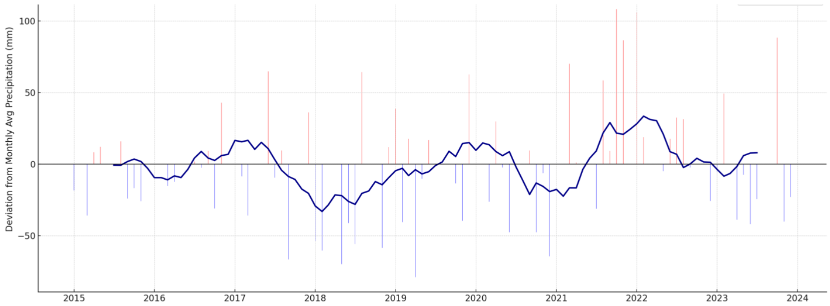

Figure 1.1.1: Comparison of precipitation levels in No.3 Zone as a percentage of the 2015-2023 average

As depicted in Figure 1.1.1, historical precipitation data for the No.3 zone reveal an upward trend compared to the long-term average. Furthermore, the figure highlights significant interannual variability, with pronounced deviations from the average years. For example, the year 2018 stands out for its dryness, while 2022 experienced a substantial surplus of rainfall. This observed variability underscores the inherent difficulty in predicting precipitation patterns and identifying a consistent annual trend. Consequently, hydropower scheduling becomes a more complicated task due to the multitude of potential inflow scenarios. This complexity emphasizes the importance of considering a diverse range of scenarios during the scheduling process.

According to Koestler [5], Norwegian hydropower plants can expect increased inflow volumes in the coming years. However, the projections also indicate that year-to-year inflow patterns will remain unpredictable. Figure 1.1.2 illustrates a comparison between a typical annual inflow profile and the corresponding power demand in Norway.

In Norway, typical seasonal variations influence reservoir inflow patterns. The highest inflow volumes typically occur during the spring flood season, coinciding with late spring and early summer. A secondary increase in inflow is also observed during autumn due to increased precipitation levels. Conversely, the demand exhibits a more predictable pattern, primarily driven by temperature fluctuations. Consequently, demand is generally higher during winter season for heating purposes and lower during summer season [6].

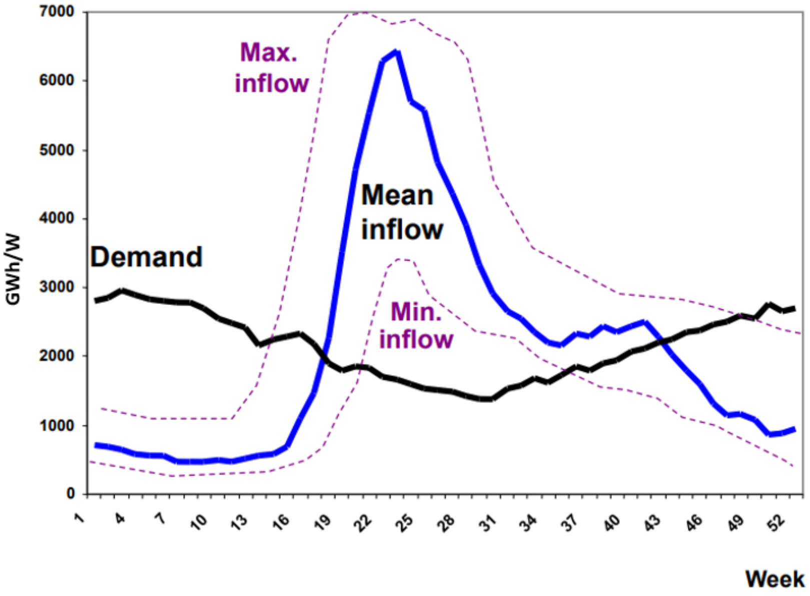

Figure 1.1.2: Inflow vs power demand for a typical year in Norway [Source: Course syllabus]

This dynamic underscore the critical role of hydropower scheduling problem involves deciding when to utilize water resources for power generation and when to save the water for future use. The goal of the problem is to maximize profits by finding the optimal generation schedule, which specifies how much water will be released for each time step in a planning horizon. To decide future schedules operators, rely on uncertain forecasts about future prices and how much water will flow into the reservoir. These anticipated changes are likely to introduce greater complexity and computational demands for achieving optimal hydropower scheduling strategies

Power prices are highly volatile and fluctuate hourly, offering opportunities for hydropower operators to achieve a higher realized price. As shown in Figure 1.1.3, power prices in the first half of September 2022 varied widely, ranging from below 30 EUR to over 500 EUR per unit. This volatility can significantly impact the revenue of the operator, potentially by more than a factor of 10. Consequently, choosing an optimal generation schedule can lead to large increases in profit for the owner of the hydropower plant. Because hydropower is such a large part of Norwegian power generation, increasing the profits by a tiny factor can give large increases in profits.

The Nordic countries have a liberialised common market, in which hy-

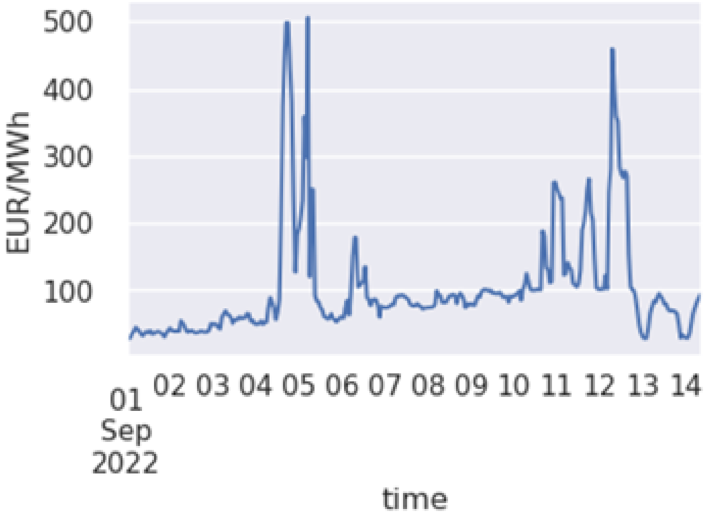

Figure 1.1.3: Power Price in September 2022 (Tronder Energi, 2022)

dropower constitutes about 50 % of total generation [7]. In Norway, hydropower accounts for 99 % of the total power generation. Due to the high fraction of hydro, market prices will depend on the hydrological situation, and taking the stochasticity in inflow into account is therefore essential. Assuming that the individual producers cannot influence the market price, it is also necessary to model stochasticity in price. In this thesis, stochastic models for price and inflow are developed in order to provide forecasts for the long-term scheduling. Parameters for these models are computed based on historical spot prices, forward and futures prices and historical inflow data collected from Professor and Ph.D Candidate.

1.2 Goals and Research Questions

The main purpose of this thesis is to find methods from the field of Artificial Intelligence that can be used to solve the hydropower scheduling problem. Formulating clear research questions is the first step in structuring the study, as these questions will guide the literature review, analysis, and overall direction of the work.

To support this goal, the research questions were designed to identify and review a wide range of relevant studies and data sources that highlight the use of ML techniques in hydropower optimization.

The following research questions have been developed for this thesis:

1. What are the state-of-the-art optimization and forecasting methods used in long-term hydropower scheduling?
2. Which techniques exhibit the best fit for addressing the hydropower scheduling problem?
3. What are the main expected benefits from applying machine learning techniques?

1.3 Contributions

In this thesis, the conducted work extends existing approaches by integrating a machine learning models into long term hydropower scheduling (LTHS). The main contributions are as follows:

- A forecasting framework was developed using RNN and LSTM models to improve prediction accuracy for inflow, electricity prices, production and demand in the NO3 (Central Norway) region.
- The proposed framework related to the long-term scheduling problem, allowing for data-driven decision support in reservoir operations and energy bidding.
- The performance of the machine learning models was evaluated through multiple case studies across different seasons and market conditions.
- The models were assessed using standard performance metrics and compared with baseline methods to analyze their robustness and suitability for practical application.

The entire model has been implemented and the results analyzed in Python, with custom modifications for data integration and model training.

1.4 Thesis Structure

This thesis is structured as follows:

Chapter 1: Introduces the background and motivation for the project, defines the problem, outlines the research objectives, and the thesis questions.

The conducted work is also described in this section.

Chapter 2: Describes the relevant and established theory to this thesis and discusses hydropower and its scheduling with a focus on Long-Term Hydropower Scheduling (LTHS) and also describes the theory associated with machine learning techniques and performance metrics. In this chapter also a literature review on the application of machine learning in hydropower scheduling and a description of the conducted work is done in this section.

Chapter 3: Details the methodological approach. This includes descriptions of the different machine learning methods implemented.

Chapter 4: Provides an overview of the dataset and case study are described in detail in this chapter. It also outlines the selected machine learning models and the evaluation metrics used in this study.

Chapter 5: Presents the results obtained from the different case studies, primarily on optimal Inflow, demand, scheduling accuracy, and computational time.

Chapter 6: This chapter discusses the outcomes of the experimental analysis presented.

Chapter 7: Finally, we summarize the main points and answer our research questions. We also discuss limitations and potential future work.

Chapter 2

Theory and Literature Review

This chapter outlines the relevant and established theory to this thesis and discusses hydropower and its scheduling with a focus on Long-Term Hydropower Scheduling (LTHS) and various scheduling approaches and describes the theory associated with machine learning techniques and performance metrics. The chapter also literature review on the application of machine learning in hydropower scheduling and summarizes the work conducted is done in this chapter.

2.1 Hydropower

In hydropower, the mechanical energy of flowing water is converted into electrical energy via the use of a turbine. Water has potential energy because it is pulled by gravity. The higher the elevation of the water and the more mass of water, and more potential energy is stored.
There are different types of hydropower plants. The most common in Norway is storage hydropower plants, which have a dam and a reservoir that can accumulate water. In Figure 2.1.1 shows some of the components of a hydropower dam. A reservoir can be released water into the penstock to drive a turbine. A generator converts the movement of the turbine into electrical energy, which is delivered to the power grid [8].

The Hydropower plants can be classified into two distinct categories: Reservoir-based and Run-of-river hydropower plants.

Reservoir-based hydropower plants feature huge storage facilities, enabling them to retain considerable volumes of water. The storage capacity allows the plant operators to schedule power generation at the most optimum pe-

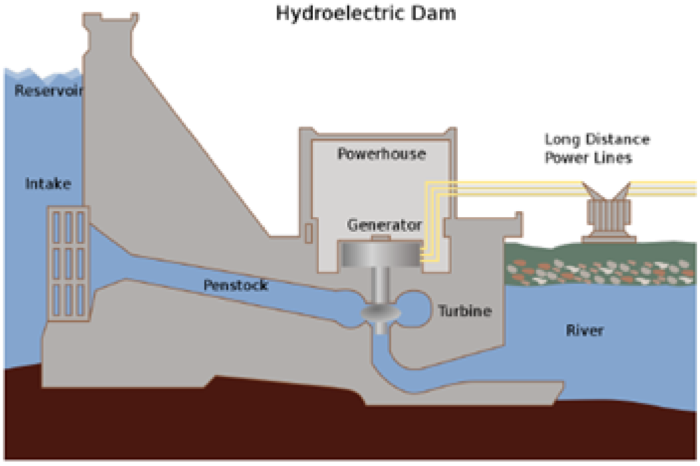

Figure 2.1.1: Hydropower plant
(Energy Information Administration, 2022)

riods. Reservoir-based hydropower plant creates electricity by capturing the potential energy of the water contained in the reservoir from the height difference between the reservoir and the generator. This may be stated by the following equation [9]:

$$P = \eta \times \rho \times g \times h \times Q \tag{2.1}$$

Where P is the power output, η is the efficiency of the plant, ρ is the density of water, g is the acceleration due to gravity, h is the height difference employed to create energy, and Q is the flow rate of water through the turbines.

On the other hand, run-of-river hydropower plants can't store water. They generally generate power from the river natural flow, as the name indicates. Only some of the river water is used to provide electricity, and then it goes back into the river. There is a notch in the weir that makes sure there is enough flow in the river all year round. This is because the flow rates change naturally. The turbines in these buildings transform the movement of the water into power. This shift in kinetic energy occurs according to equation [9]:

$$P = \frac{1}{2} \times \eta \times \rho \times A \times v^2 \tag{2.2}$$

Where P is the amount of electricity that comes out, η is the plant efficiency, ρ is the water density, A is the stream cross-sectional area, and v is the speed at which the water flows.

A typical reservoir-based hydropower plant consists of five major components: the reservoir, the tunnel, the turbine and generator, the discharge tunnel, and the discharge pool.

The reservoir holds and stores the water that naturally flows into a region [10]. The reservoir might take the shape of a natural pond, or an artificial body produced by a dam. The geographical characteristics of the location determine the choice of dam type and size. A regulation hatch is incorporated to release the water for maintenance or emergencies. To prevent structural damage to the dam, a spillway is commonly incorporated to discharge the water safely in the event of overflow or spillage. Any overflow or spillage from the reservoir results in the loss of potential energy and revenue for the plant.

The hydropower plant main parts are the turbine and the generator. The water potential energy is turned into electrical energy here [11]. The turbine turns because of the water pressure, which turns the generator to produce power. Depending on the height difference and the amount of water flow in the plant, several kinds of turbines are employed. There are many kinds of turbines for different height disparities. The efficiency of each turbine will depend on the water that is present [12]. By using them together between numerous pairs of turbines and generators, the difference in efficiency may be kept to a minimum.

The discharge tunnel takes the water to the discharge pool after it has gone through the turbine. The tunnel is lower than the reservoir so that a vacuum does not form in the water tunnel, which would make the plant less efficient. The discharge pool is the last part of the hydroelectric plant. It might be another reservoir, an uncontrolled river, or a direct connection to the sea

The total efficiency of a hydro power plant is the sum of all losses in the system, such as those that occur in the tunnel, the turbine, and any vacuum that is produced. This can be represented by the following equation.

$$\eta = 1 - \eta_{\text{tunnel}} - \eta_{\text{turbine}} - \eta_{\text{vacuum}} - \cdots \tag{2.3}$$

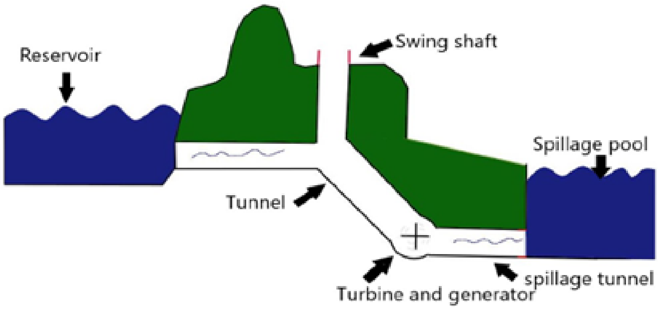

Figure 2.1.2: Cross-sectional view of a typical hydropower plant

2.2 Hydropower scheduling

As with wind, solar, and thermal power, hydropower contributes to the energy production within the power system. The determination of a marginal price for power generated from wind, solar, and thermal sources is a relatively straightforward process [13]. The decisions made regarding the immediate production of power from these sources may not necessarily influence the future capacity of power production at a later stage [14].

In contrast, hydropower have a storage capability, involves decisions made at one stage that impact reservoir levels in the future. Consequently, there is a necessity for strategies to determine the marginal value of accessible water. In addition to meeting immediate demand, future projections are significantly influence present decision-making. For instance, if forecasts indicate substantial inflow in the upcoming period, the typical course of action would be to generate power to prevent flooding. Conversely, if predictions indicate limited inflow, power production may be reduced to avoid future rationing [13].

Hydropower scheduling presents several challenges highlighted below:

- System size and topology: A hydropower system may consist of multiple modules with complicated topology, each regulated by distinct restrictions that increase to the total complexity. It also makes it much harder to connect these modules to other areas of the power system.
- Scheduling horizon and time step: The scheduling horizon varies from hourly intervals to several years, depending on the level of detail needed.

While it is preferable to use small time steps for scheduling, doing so over an extended horizon would be computationally infeasible.

- Uncertainty: Various uncertain parameters complicate the hydropower scheduling process, with the primary factors being inflow, demand and market prices.

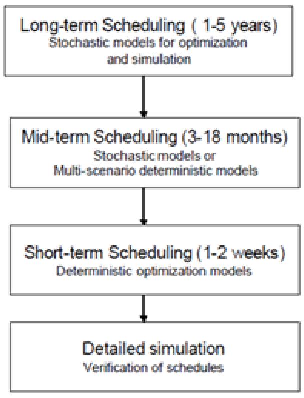

Figure 2.2.1: Hydropower scheduling hierarchy
Source: [15]

As illustrated in Figure 2.2.1, there is a connection between the scheduling horizons, and each of them plays a distinct role. Short-term scheduling is crucial because it is directly connected to decisions related to power production. Furthermore, it is influenced by the actual values of parameters such as inflow and consumption and interacts with the power market [15]. In contrast, LTHS serves as a reference point to guide seasonal and short-term scheduling, thereby facilitating the long-term management of resources. For hydropower companies, it is of great importance to have a comprehensive understanding of the future reservoir situation [15]. Seasonal scheduling serves as a bridge between LTHS and short-term scheduling. The three scheduling horizons in Figure 2.2.1, the focus of this thesis is only on LTHS.

2.2.1 Long-Term Hydropower Scheduling

The aim of LTHS is to achieve optimal hydropower production planning over a long planning horizon. For the largest reservoirs, it can take several years to fill up when they are empty [15]. Consequently, the decisions made at a time step can have an impact on the reservoir levels years ahead. In LTHS, forecasts of prices, demand, and inflow are of extremely importance. Due to the inherent uncertainty associated with these factors, LTHS is solved as stochastic dynamic optimization problems (SDP) [15]. The dynamic nature of the problem comes from the fact that a decision made at one time stage, affects subsequent decisions. However, SDPs struggle with "the curse of dimensionality" [15]. As the number of inflow scenarios for a given time step increases, the dimensionality of the optimization problem grows exponentially [16]. Furthermore, the computational burden increases with smaller time steps due to the increased complexity. In the context of LTHS extending over a year, the incorporation of numerous inflow scenarios results in an excessively high number of potential states in the long term. Figure 2.2.2 illustrates the reservoir level with stochastic inflow.

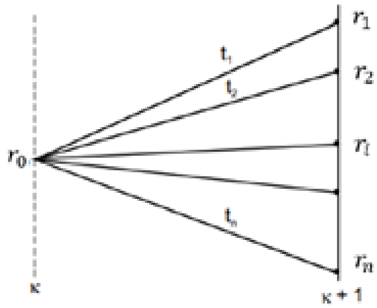

Figure 2.2.2: Example of reservoir levels with n stochastic inflow scenarios

Let r_0 be the reservoir level in time step κ. The inflows $t_1, t_2, \ldots, t_n$ yield n possible inflow scenarios, resulting in $r_1, r_2, \ldots, r_n$ as the possible reservoir levels in time step $\kappa + 1$. Assuming each reservoir level at time step $\kappa + 1$ leads to n new inflow scenarios, this process continues, resulting in new reservoir levels at time step $\kappa + 2$. Considering an excessive number of inflow scenarios per time step across all time steps in several years makes LTHS computationally infeasible due to the vast number of potential states to consider [15].

2.3 Model Framework

In LTHS with a disaggregated representation of the hydropower system, computational complexity remains a significant challenge. The model used in this thesis incorporates a machine learning-based framework designed to make scheduling more computationally feasible. The following sections describe the theory related to the structure and methods used in the framework.

2.3.1 Machine Learning and Neural Networks

Machine learning (ML) is set of techniques used to extract knowledge from data using computer force to learn relationships from datasets by iteration [17]. Instead of using hand coded rules of “if” and “else” decisions to process data, one feeds the data to algorithms that are constructed so that the program learns from experience autonomously and with minimal human interaction. By numerous iterations, the computer is tasked to find the most optimal route from input to output. Machine learning is a field that combines computer science and mathematics. It has a substantial portion of its theoretical core from probability theory and statistics as well as linear algebra and calculus that are nested together in the form of algorithms. Constructed to do different tasks such as finding the minimal error by gradient decent, these algorithms rely on heavy computer hardware to process and find patterns in large datasets. In geosciences, machine learning has been used for a long time in remote sensing applications such as automatic classification of landcover and geomorphology from satellite data, in weather forecasting and in hydrology.

The terms "machine learning" and "Artificial Intelligence" (AI) are often used interchangeably or in conjunction with one another. Today, artificial intelligence is seen as a more general term that covers computer vision, optimisation, data mining, robots, remote sensing, and other areas. Machine learning is a part or section of AI that focuses on using data to learn how to solve difficult problems. It is seen as one of the most important skills for the future of AI.

Remsan and Mathew’s 2015 book [18], "Data Driven Models in Hydrology," has many examples of how machine learning models can be used in hydrological forecasting. Because of this, support vector machines and different kinds of Artificial Neural Networks (ANN) are two of the most popular ways to predict time series forecasting [18]. Models vary in how complicated and adaptable they are. A lot of people use basic machine learning techniques, like Support vector regression or K-nearest neighbours are popular models in hydrological time series forecasting and they offer good predictability for a lower computational cost than the deeper networks.

The complexity level is crucial when choosing a model. More complex models are assumed to be better at simulating reality. This is certainly the case in the training phase (good fit with lower bias and higher variance). But less complex models may perform better when testing on unseen data. Less

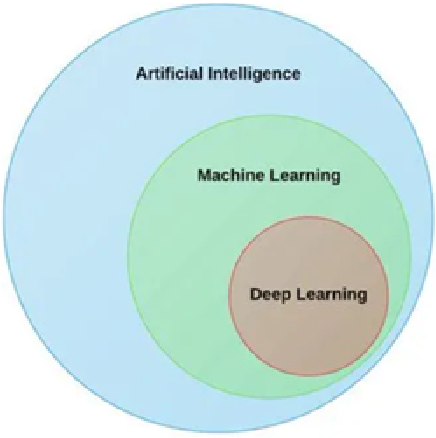

Figure 2.3.1: Relationship between Artificial Intelligence, Machine Learning, and Deep Learning

complex models can be more robust and generalize better, especially when working with shorter data series. Remsan and Mathew [18] argue that these issues are not adequately addressed in data driven hydrological modelling. Further they admit that it is difficult to decide on which model to use. Decisions are often based on familiarity, but complexity, length and type of data, computational resources and various other aspects should be considered when choosing a model.

In hydrological applications, machine learning has been widely applied to tasks such as inflow prediction, weather forecasting, and remote sensing analysis. Among the most commonly used models are Artificial Neural Networks (ANNs), including Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks. These models are computationally efficient and suitable for shorter time series, while deeper architectures offer greater flexibility but may risk overfitting.

Model selection in hydrology relies on factors like data size, complexity, and computational resources. For training data, complex models work best, but for smaller datasets, simpler models may be more effective for limited datasets. Forward neural networks (ANN) are the most common type, consisting of fully connected layers that flow data without loops. These networks are structured with input layers, dense hidden layers, and output layers. The number of layers and nodes varies depending on the application and data complexity. They are in high demand due to their enhanced forecasting capabilities and their capacity to integrate numerous data sources. This understanding sets the stage for more advanced architectures like Recurrent Neural Networks (RNN) and Long Short-Term Memory networks, which are discussed in the next sections.

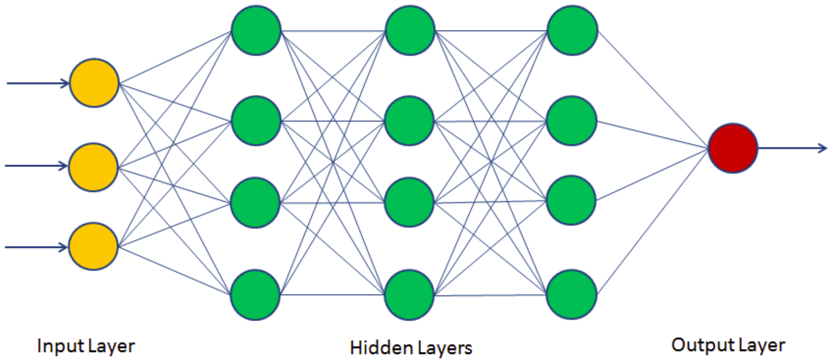

Figure 2.3.2: Basic overview of a neural network

2.3.2 Recurrent Neural Network

Recurrent Neural Networks (RNN), introduced by Rumelhart et al.(1986)[19], are a class of artificial neural networks designed to model sequential data with internal 'self-loops' that allow information from previous time steps to influence current predictions, making them suitable for time series analysis and forecasting. The application of RNNs in hydrology began in the late 1990s. The former conducted the research in a laboratory and demonstrated the ability to use RNN for event-based applications. Hsu et al. (1997) [20] compared the RNN against ANN and the models had about the same level of predictability. One of the most compelling advantages of RNN in hydrological modeling is their capacity for "internal memory," enabling the network to learn temporal dependencies. However, a major drawback of the early RNN was the phenomena with vanishing and exploding gradients in the backpropagation.

A simple RNN model is consists of a single cell that processes sequences of time series data. They contain self-loops that allow them to carry information from previous time steps, making them suitable for time series forecasting [20]. The cell is updated as each sequence is processed. In order to backpropagate we unroll the network as seen in the figure below:

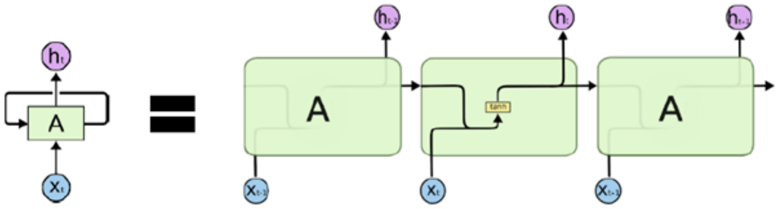

Figure 2.3.3: When unrolled, each cell can be viewed as one dense layer. Hence the network becomes analogous to a very deep artificial neural network Source:[21]

2.3.3 Long Short Term Memory Network

Long Short-Term Memory (LSTM) networks, proposed by Hochreiter and Schmidhuber (1997) [22]. Over the years it has become one of the most famous RNN and are still a popular choice for tasks such as handwriting recognition [23], music composition, traffic forecast [24] or other sequential problems. In a paper written by researchers from Google (Jozefowicz, et al., 2015) [25] the LSTM is forwarded as an extremely powerful and applicable framework for a broad variety of machine learning tasks. The LSTM are a refined version of RNN designed to overcome the limitations of standard RNN, particularly the vanishing gradient problem. Long short-term memories (LSTM) include memory cells and gating mechanisms to control the flow of data over time, allowing the networks to retain long-term dependencies more effectively.

LSTM are also briefly described in the book by Remesan and Mathews (2015) [18]. However, it is important to note that the majority of research articles that are used as references for the more advanced forms of RNN, such as LSTM or GRU in hydrology is published later. From 2015 the large tech giants revealed that LSTM based networks were used in their technology. For example, Google started using LSTM for speech recognition in 2015 on Google Voice, and in 2016 they started using it for their translation system reducing translation errors by 50% (Highfield, 2015) [26]. Similarly, in 2016 Apple and Amazon started using LSTM-based networks for translations, Quick type, and text-to-speech technology in their devices.

In the field of hydrology, LSTM are particularly well-suited for modeling the complex relationships between the variables that are input, such as precipitation, temperature, and other parameters, and the variables that are output, such as inflow etc. Their ability to capture both daily and seasonal patterns is especially valuable for inflow forecasting and hydropower scheduling.

The drawing below shows the LSTM architecture unrolled over time. The cell to the left is the LSTM cell at the previous time step while the one to the right is the cell at one step into the future. The current time step is in the middle. Three lines goes into the cell. In the bottom left corner, it receives the input X_t and the output from the previous time step (the output from the previous layer is in RNN called the hidden state abbreviated to h_{t-1}). The input X_t and the hidden state h_{t-1} are concatenated before it runs into the four gates marked as yellow boxes in the drawing. The third input the cell receives from the previous cell, runs as a straight arrow trough the upper part of the cells. This is the cell state and enables the LSTM to remember long term dependencies with a considerably smaller chance for the vanishing and exploding gradient problems seen in traditional RNN.

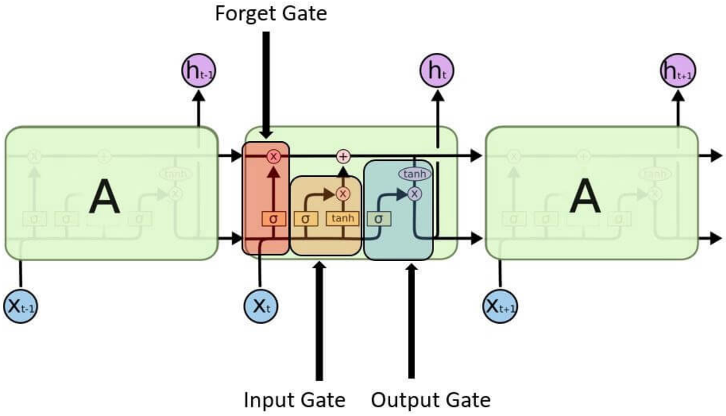

Figure 2.3.4: Overview of the LSTM network
Source:[21]

2.3.4 Performance Metrics

The overarching goal of this study is to assess the potential of machine learning algorithms in long-term hydropower scheduling, thus reducing the overall time required. To effectively evaluate the efficiency of machine learning techniques relative to traditional optimization methodologies, two principal factors must be considered: the computational time and the accuracy of the program.

Computational time, denoting the duration a computer algorithm requires to complete its task, provides an efficient and practical way of comparing distinct algorithms [27]. The hardware setup will impact the performance of computational time, it is, therefore, essential to run the algorithms on the same or equal computers when comparing.

To objectively measure the efficacy of the machine learning model, various performance metrics are employed, including the Nash–Sutcliffe Efficiency (NSE), Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and mean absolute percentage error (MAPE). This research will primarily utilize the first three metrics: MSE, RMSE, and MAE. Using diverse evaluation techniques allows for a more comprehensive analysis, considering each metric has unique strengths and limitations.

The **Nash–Sutcliffe Efficiency (NSE)** is a normalized statistic that determines the relative magnitude of the residual variance compared to the measured data variance. It is computed as follows[28]:

$$NSE = 1 - \frac{\sum_{i=1}^{N}(r_i - \hat{r}_i)^2}{\sum_{i=1}^{N}(r_i - \bar{r})^2} \tag{2.4}$$

where N is the total number of predicted values, $\hat{r}_i$ is the predicted values, r_i are the actual values. NSE values range from $[-\infty \, to \, 1]$, where a value of 1 corresponds to a perfect match between model predictions and observations. An NSE of 0 indicates that the model predictions are as accurate as the mean of the observed data, while values less than zero imply that the mean observed value is a better predictor than the model.

The **Mean Absolute Error (MAE)** is computed as follows [29]:

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|r_i - \hat{r}_i| \tag{2.5}$$

where N is the total number of predicted values, $\hat{r}_i$ is the predicted values, r_i are the actual values. MAE gives the mean absolute value between the predicted values and the measured data without considering the direction of the error.

The **Mean Squared Error (MSE)** is computed as follows [30]:

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(\hat{r}_i - r_i)^2 \tag{2.6}$$

where N is the total number of predicted values, $\hat{r}_i$ represents the predicted values, and r_i is the actual values. The MSE computes the average of the squares of the errors, emphasizing larger errors due to the squaring operation.

The **Root Mean Squared Error (RMSE)**, a derivative of MSE, is expressed as [31]:

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{r}_i - r_i)^2} \tag{2.7}$$

where N is the total number of predicted values, $\hat{r}_i$ signifies the predicted values, and r_i represents the actual values. The RMSE is the square root of the MSE, mitigating the heavy penalization of larger errors exhibited by MSE, thus providing a more balanced measure of model performance.

The **Mean Absolute Percentage Error (MAPE)** can be calculated through this equation:

$$MAPE = \frac{100}{N}\sum_{i=1}^{N}\left|\frac{r_i - \hat{r}_i}{r_i}\right| \tag{2.8}$$

where N is the total number of predicted values, $\hat{r}_i$ is the predicted values, r_i signifies the actual values. MAPE gives the mean absolute percentage value between the predicted values and the measured data without considering the direction of the error.

2.4 Literature review

This master thesis aims to investigate the performance of machine learning as a substitute for the optimization techniques used in long-term hydroelectric power scheduling (LTHPS). As hydropower plays a vital role in global energy systems, the need for robust and adaptive planning methods is growing due to increased system complexity, market volatility, and climate uncertainty.

The literature review for this thesis has three main goals:

1. To explore the state-of-the-art optimization techniques used in LTHPS.
2. To identify the machine learning methods currently applied to hydropower forecasting and scheduling.
3. The comparative performance of two different machine learning techniques against conventional optimization-based approaches.

To ensure the research is most relevant to the field, it is important to compare the results of machine learning to the state-of-the-art LTHPS optimization techniques while considering the main break throughs in LTHPS optimization techniques. Different machine learning techniques have been used in the field; therefore, it is important to ensure that the research in this thesis uses a new technique or approach to further advance the research in the field.

2.4.1 Hydro power scheduling

Hydropower scheduling is a crucial role in the efficient operation of power systems, particularly in areas with substantial hydro resources. The objective is to determine the optimal release of water over time to maximize revenue or meet demand, while also considering operational, environmental, and market constraints. In long-term planning, this issue becomes more complicated because of the unpredictability of reservoir inflows and fluctuating electricity market prices.

One of the foundational concepts in hydropower scheduling is the value of water, which represents the opportunity cost of using water for power generation now versus storing it for future use. This principle allows hydro operators to strategically allocate water resources across time. Early models focused on calculating the water value based on reservoir levels and inflow probabilities. Traditional hydropower scheduling methods, such as Stochastic Dynamic Programming (SDP), Stochastic Dual Dynamic Programming (SDDP), and Mixed-Integer Linear Programming (MILP), have long been employed for long-term energy planning. These techniques help in modeling uncertainties in inflow and market conditions and optimizing water usage across multiple time steps [32, 33]. Despite their effectiveness, these methods can become computationally expensive when applied to systems with several reservoirs or fine-grained time resolutions. Additionally, the issue becomes more complex due to the need for scenario-based modelling, which makes real-time implementation challenging [34]. These limitations have prompted the exploration of machine learning-based forecasting and optimization models, including the Simple RNN and LSTM-based approach investigated in this thesis.

2.4.2 Machine Learning in Hydro power scheduling

Machine learning (ML) techniques have been included into hydropower scheduling frameworks due to the rising complexity of power network, as well as the fluctuating electricity prices and unpredictable inflows. ML leverages historical data to discover patterns and makes predictions based on past data. In long-term hydropower scheduling (LTHPS), machine learning (ML) techniques are useful for forecasting reservoir inflows and market prices. Early applications used Artificial Neural Networks (ANN) for inflow prediction. While effective at capturing nonlinear relationships, ANNs and feedforward networks lack the ability to model temporal dependencies [35]. Recurrent Neural Networks (RNN) improved their performance by incorporating memory through internal states, which is suitable for time series forecasting. However, traditional Simple RNN suffer from vanishing gradient problems [36]. The development of Long Short-Term Memory (LSTM) networks was implemented in order to overcome these constraints. LSTM use memory cells and gating mechanisms to retain information over longer sequences, making them ideal for hydrological and price forecasting [37]. In hydropower scheduling, LSTM models forecast future inflows and prices, which are then used in simplified optimization models. These hybrid systems improve scalability and adaptability, especially in competitive electricity markets [38].

2.4.3 Conclusion of the Literature review

The literature highlights two machine learning models used in hydropower scheduling. Here can it be seen that Simple RNN and LSTM were two of the most used machine learning technique for hydropower scheduling. Simple RNN did not perform well in this field, but LSTM networks have demonstrated more promising results due to their ability to capture long-term dependencies. Therefore, the LSTM model is considered the most suitable choice for this master's thesis. To the best of the author's knowledge, LSTM networks have been applied only to a limited extent in hydropower scheduling, making their use in this thesis.

In the next chapter, the methodology and corresponding mathematical models will be presented, and further details will be given for the different parameters and variables of the real-world equivalent.

Chapter 3

Methods

This chapter outlines the methodology of the thesis. It starts by detailing the data used in the project and the preprocessing techniques applied. Then, it covers the feature engineering steps and the model development process for the recurrent neural network (RNN) and long-short-term memory (LSTM) architectures. This chapter also presents the mathematical formulation of both models and explains the criteria used for selecting the best model (including error metrics like NSE, MAE, RMSE etc). Finally, it outlines the approach used to evaluate the selected model and provides details on the software and hardware tools used for the implementation. Results from the preprocessing and training process are presented in this chapter, while the results directly related to the research questions are presented in Chapter 5.

3.1 Feature Engineering

Before training the models, the NVE dataset undergoes several pre-processing and feature engineering steps to make it suitable for machine learning. First, data cleaning was performed; any missing values or obvious outliers in the time series were addressed (by imputation using neighboring values or removing faulty entries) to ensure consistency. Next, all characteristic variables were standardized to have a zero mean and unit standard deviation; this scaling was applied to meteorological inputs (temperature, precipitation, etc.), operational inputs (such as past production, consumption, and price), and also to target outputs. Standardizing the inputs helps the neural network models train more effectively by ensuring that features are on a comparable scale and speeding up the convergence of gradient-based optimization. In addition to scaling, new features were derived to help the models capture important patterns.

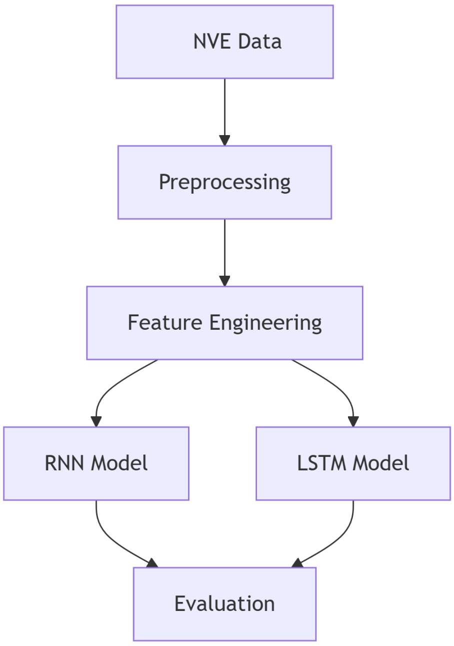

Figure 3.1.1: Flow Chart

Temporal features: We created time-related features like month or season indicators (e.g., specifying the month of the year). These can allow the model to consider seasonal effects that are not obviously present in the continuous input variables. Lag features: Because both RNN and LSTM learn from sequential data, we prepared the training data as sequences of past observations. For each day in the dataset, we constructed an input sequence consisting of the previous N days of characteristics (e.g. $N = 7$ days or $N = 30$ days, depending on what gave the best results during the experimentation). These lagged sequences allow the models to learn temporal dependencies. In practice, this meant that for each target prediction (say, forecasting the next day's inflow, price, or production), we included features like temperature, precipitation, etc., from the prior N days as part of the model input. We also considered aggregated features such as rolling averages or accumulated values (e.g., cumulative rainfall over the past week, average temperature over the past month). Such features can smooth out short-term noise and emphasize longer-term trends that could be relevant for long-term hydropower scheduling decisions. Feature selection was another aspect of our preparation. Given the number of potential input variables (multiple weather variables, lagged versions of each, past outputs like yesterday's price or production, etc.), not all features may be equally informative. We used a step-by-step process to pick the best features by training the initial models with different groups of features and checking how well they performed (mainly using metrices as a measurement, which we explain later in Section 3.2.2). This helped identify a set of the most predictive features without choosing the ones that add noise or redundancy. The final feature set included the main weather factors (temperature, precipitation, evaporation, and wind speed) and relevant operational factors (past inflow, past production/consumption, and past price) along with select engineered features such as lagged values and seasonal indicators. At the conclusion of this feature engineering process, we had the model.

3.2 Model Development and Training Process

In this thesis, two types of neural network models were developed for predicting hydropower-related outcomes: a Simple RNN and an LSTM network. All models were trained, validated [39], and tested on subsets of the data from 2015 to 2023. In this section, we introduce the models and explain the general training process, including how data was split and how hyperparameters were chosen. Figure 3.2.1 provides a high-level flowchart of the modeling process, from data preprocessing to model training and evaluation.

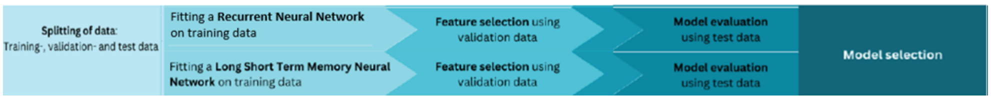

Figure 3.2.1: Flow chart showing the steps and data used during training, model evaluation and model selection

As illustrated in Figure 3.2.1, the workflow began with the raw data obtained from NVE and other sources, followed by preprocessing and feature engineering (as described in Section 3.1). After preparing the data, we split the dataset into training, validation, and test sets in chronological order. A typical split used in our experiments allocated the earliest portion of the timeline to training, an intermediate period to validation, and the most recent portion to testing. For example, one split scenario used data from 2015–2023 for training, 2021–2022 for validation (used during model tuning), and the final year 2023 as a hold-out test set for final evaluation. This chronological splitting strategy preserves the temporal order of events (ensuring that we always predict “future” data from the perspective of the model training) and avoids information leakage from the future into the past After splitting the data, we trained the two model architectures on the training set:

- Basic RNN model: a simple recurrent neural network with one or more recurrent layers (using tanh or ReLU activations) to serve as a baseline for sequence modeling.

- LSTM model: an advanced RNN variant incorporating LSTM layers, which include gating mechanisms to better capture long-term dependencies.

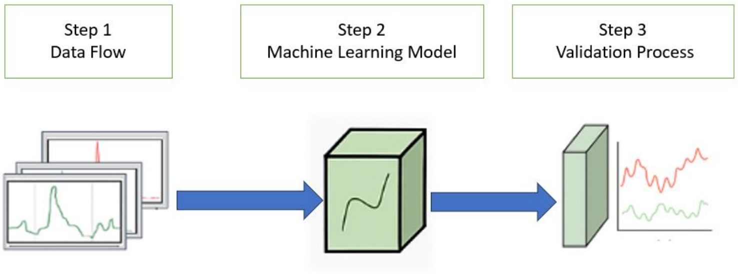

Figure 3.2.2: The framework of scenario reduction for long-term hydropower scheduling and validation process

During the model development, we construct the model with various setting for both type of architecture, such as the number of hidden layers, the number of units (neurons) in each layer, the sequence length (input time steps N), the learning rate for the optimizer, and the regularization techniques (like dropout rates) to prevent overfitting. We utilized the Adam optimizer for training these neural networks, as it generally provides fast convergence for time-series tasks. The loss function employed for training was the mean squared error (MSE), which is standard for regression problems – although we monitored mean absolute error (MAE) as well due to its interpretability in the context of our targets. Each model was trained for a fixed number of epochs, but we also implemented early stopping based on validation loss; if the validation lossess did not improve for a certain number of epochs, training was stopped to prevent overfitting to the training data.

Throughout training, the performance on the validation set was tracked. The basic RNN model training served as a baseline; once trained, we evaluated its validation error and noted its strengths and weaknesses (for example, we observed whether it could capture seasonal patterns or if it lagged during sudden changes in the data). Then, the LSTM model was trained under similar conditions. We expected the LSTM, with its memory cell and gates, to outperform the basic RNN especially on longer-term patterns, and we indeed found that the LSTM typically yielded lower validation error in our trials. The outcome of this stage was two trained model candidates (RNN and LSTM), each with tuned hyperparameters gleaned through the validation process. We next proceeded to a formal model selection phase to decide which architecture and configuration would be used for final testing.

3.2.1 Mathematical Modeling of RNN and LSTM

Recurrent neural networks are designed for sequential data processing and forecasting. In a Simple RNN, at each time step t, the model takes an input vector x_t, combines it with the previous hidden state h_{t-1}, and produces an updated hidden state h_t. The hidden state update can be described by the equation:

$$h_t = \tanh\big(W_{xh}\, x_t + W_{hh}\, h_{t-1} + b_h\big) \tag{3.1}$$

where W_{xh} is the weight matrix connecting input to the hidden state, W_{hh} is the recurrent weight matrix from the previous hidden state to the new hidden state, and b_h is a bias vector. The $\tanh(\cdot)$ is the activation function provides non-linearity and constraints on the value of the hidden state. The RNN produces an output y_t (for example, a prediction of the target variable at time t) based on the hidden state:

$$y_t = W_{hy}\, h_t + b_y \tag{3.2}$$

where W_{hy} maps the hidden state to the output and b_y is an output bias. In hydropower scheduling, y_t could represent a predicted value such as next day inflow, price, or production, and x_t would include features like recent observations of weather and other variables. Simple RNN, however, are known to struggle with learning long-term dependencies in sequences. Over long time periods, gradients tend to vanish or explode during backpropagation through time, making it difficult for the RNN to learn relationships that span more than a short number of time steps.

LSTM networks were introduced to address these limitations by providing a more complex memory structure. An LSTM unit at time t carries not only a hidden state h_t but also a cell state C_t that acts as a long-term memory. The LSTM uses gating mechanisms to control the flow of information into and out of this cell state. The equations governing an LSTM operation are as follows. First, the forget gate f_t decides how much of the previous cell state C_{t-1} to retain:

$$f_t = \sigma\big(W_f\, x_t + U_f\, h_{t-1} + b_f\big) \tag{3.3}$$

where $\sigma(\cdot)$ is the sigmoid activation function that squashes values between 0 and 1. Next, the input gate i_t determines how much new information to write to the cell, and the candidate cell state $\tilde{C}_t$ computes a new candidate value to potentially add to the cell state:

$$i_t = \sigma\big(W_i\, x_t + U_i\, h_{t-1} + b_i\big) \tag{3.4}$$

$$\tilde{C}t = \tanh\big(W_c\, x_t + U_c\, h + b_c\big) \tag{3.5}$$

where W and U are weight matrices for the input and recurrent connections of the respective gates (forget, input, candidate), and b_c are biases. The cell state is then updated by combining with the old cell content (scaled by f_t) with the new candidate content (scaled by i_t):

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \tag{3.6}$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t$$

where $\odot$ denotes element wise multiplication. This equation shows that when f_t is near 1 and i_t is near 0, the old cell state is mostly preserved (long-term memory maintained), whereas when f_t is small and i_t is large, the cell state is over written with new information. Finally, the output gate o_t decides how much of the cell state to expose to the hidden state, and the hidden state is updated as:

$$o_t = \sigma\big(W_o\, x_t + U_o\, h_{t-1} + b_o\big) \tag{3.7}$$

$$h_t = o_t \odot \tanh(C_t) \tag{3.8}$$

In short, the LSTM cell state C_t can carry information across many time steps, and the gating mechanisms (f_t, i_t, o_t) learn to protect or expose that information as needed. This architecture enables LSTM models to preserve long-term dependencies far better than simple RNNs. In our experiments, this means an LSTM is more capable of learning relationships such as "a dry winter (low precipitation and low inflow) can lead to high prices or low reservoir levels several months later" patterns that a basic RNN might fail to capture if the gap in time is large. The mathematical formulation provided above guided our implementation of the models using modern deep learning libraries, which have built in optimized routines for RNN and LSTM layers.

3.2.2 Model Selection

Once the Simple RNN and LSTM models were trained and tuned in the validation set, we carried out a model selection process to choose the final model for testing. Up to this point, the validation dataset had been used to iteratively assess performance during development. For the final selection, we considered different evaluation matrices to ensure that the chosen model was robust. During the feature selection and hyperparameter tuning phase, we focused on Nash-Sutcliffe Efficiency (NSE) and Root Mean Squared Error (RMSE) because they are important for evaluating both model accuracy and reliability in hydrological applications. Nash-Sutcliffe Efficiency (NSE) measures the predictive power of a model by comparing the predicted values to the mean of observed data. An NSE value close to 1, show that the model performance is good, but values below zero indicate poor model performance. RMSE, on the other hand, quantifies the standard deviation of prediction errors, providing insight into how much the model deviates from actual values, particularly penalizing larger errors. Although we also considered additional metrics like MAE, MSE, and MAPE for model evaluation, NSE and RMSE gave us the clearest view of overall accuracy and stability. In our experiments, the LSTM consistently achieved higher NSE values and lower MAE, MSE, and RMSE, than the Simple RNN, showing a better fit and less error variation. It also resulted in a lower MAPE, showing better percentage accuracy.

3.3 Software and hardware

The experimental part of this thesis was conducted using Python 3.11.13 within the Jupyter Notebook, an interactive development environment ideal for data exploration and iterative model prototyping. A Jupyter notebook service with access to free computing resources, including standard CPU, as well as accelerated GPU and TPU runtimes. It is connected to the Google Drive data storage service, which ensures convenient input and output files exchange. The most crucial Python libraries, from the thesis point of view, are listed in Table 3.3.1, including information about their versions used in the project.

Library	**Version**
TensorFlow	2.18.0
Keras	3.8.0
scikit-learn	1.6.1
pandas	2.2.2
NumPy	2.0.2
statsmodels	0.14.4
matplotlib	3.10.0
seaborn	0.13.2

Table 3.3.1: Python libraries used in the project

Chapter 4

Case Studies

This chapter outlines the data and case study used to develop and test the machine learning models for long-term hydropower scheduling. Section 4.1 describes the key datasets, including discharge, temperature, precipitation, wind speed, hydropower production, electricity consumption, and electricity price. Section 4.2 presents the Bratsberg hydropower plant as the case study for model application. The data and case study ensure practical relevance and support effective training, evaluation, and validation of the proposed models. This chapter ends with a summary.

4.1 Data

For this project, a combination of weather forecast data and observational data has been used as input in the machine learning models. These inputs include meteorological features such as temperature, precipitation, evaporation, and wind speed. In contrast, the model responses consist of inflow, electricity price, and hydropower production data. Daily data has been collected from 2015 to 2023 for hydropower plants in the Trondheim region. The data used in this research were sourced from both public and private repositories:

- Inflow and Weather Data: Daily inflow (discharge) volumes and meteorological parameters such as precipitation and temperature were obtained from the Norwegian Water Resources and Energy Directorate (NVE) with assistance from a PhD student.
- Production and Market Price Data: Daily records of electricity production and market prices specific to the NO3 region were provided by

the professor. This data includes real-time pricing from Nord Pool and associated reservoir operation schedules.

Norway is divided into five electricity bidding zones. The target study area for this thesis is the NO3 zone, located in central Norway. This region contains several major hydropower reservoirs and generation facilities, making it an ideal focus for evaluating long-term hydropower scheduling strategies. The geographic diversity and high reservoir density of NO3 offer a representative environment for applying machine learning-based forecasting and scheduling models. Figure shows the location of the NO3 bidding zone (highlighted in central Norway), which serves as the study area for this work (map courtesy of Statnett/NVE).

Figure 4.1.1: NO3 Bidding Zone and Target Area
(Source: Statnett/NVE)

4.1.1 Discharge

As shown in Figure 4.1.2, the discharge data analysis from 2015 to 2023 shows that how water flows, which is essential for the efficiency of hydropower operations in the Norway NO3 region. The time series shows significant seasonal fluctuations, with peak flow rates (60-90 m^3/s) occurring during spring snowmelt (April-June) and autumn precipitation (September-November), corresponding to increased turbine utilization periods, while lower values (50–150 m^3/s) coincide with winter freeze (December-February) and summer dry spells, requiring careful reservoir management. These patterns indicate the high responsiveness of hydropower production to climatic factors. Interannual trends also exhibit variability, as evident from wetter years such as 2021–2022, which sustained high flows, versus the abrupt declines observed in 2020, which can be characteristic of either drought or operational changes. These trends need to be understood in order to enhance machine learning models applied in inflow forecasting and reservoir scheduling and hence facilitate efficient energy production and grid stability. The data underscores the value of considering both seasonal cycles and extreme weather events in long-term hydropower planning.

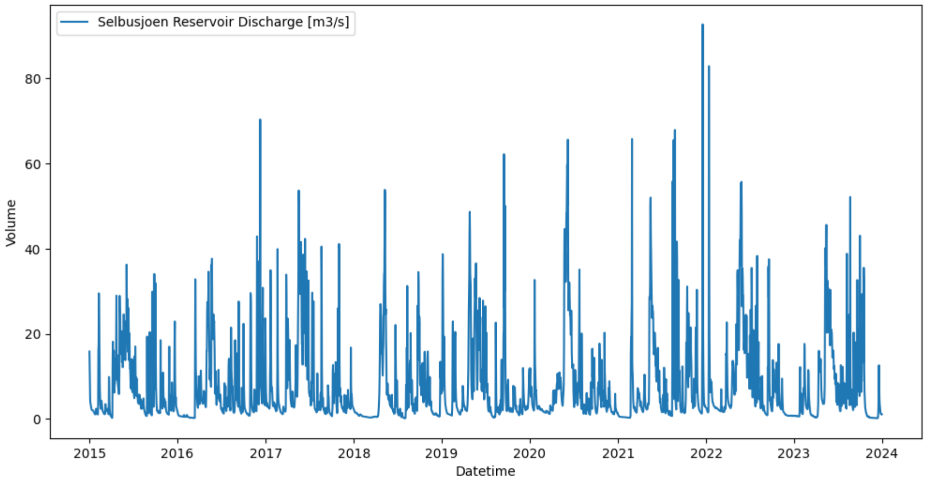

Figure 4.1.2: Discharge data (m3/s)

4.1.2 Temperature

As shown in the figure 4.1.3, the temperature time series reveals critical thermal patterns influencing NO3 hydropower operations, showing distinct seasonal fluctuations between winter lows (often below 0°C) and summer peaks (exceeding 15°C). These variations directly affect hydrological processes; winter freeze periods reduce inflow through ice formation, while summer warmth accelerates snowmelt, generating discharge peaks. These fluctuations directly influence hydrological processes; winter freeze periods diminish input due to ice formation, while summer warmth expedites snowmelt, resulting in discharge peaks. The data demonstrates significant interannual variability, with warmer years (e.g., 2020) indicating heightened evaporation losses and modified melt timing, in contrast to cold years requiring the conservation of winter reservoirs. The 2022-2023 timeframe exhibits an abnormal temperature increase that presumably affects glacier retreat rates and seasonal water supply. Hydropower scheduling requires adaptive solutions, such as temperature-correlated inflow forecasting and climate-resilient reservoir management, to sustain generating efficiency in the context of Norway rising climate. The evident thermal cyclicity further substantiates the efficacy of integrating temperature indicators as essential predictors in machine learning models for optimizing energy

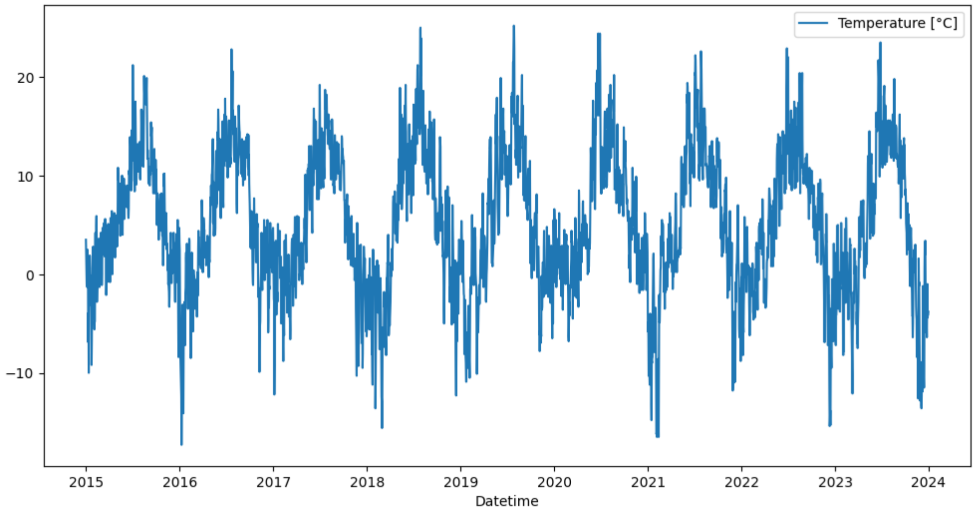

Figure 4.1.3: Temperature (°C)

4.1.3 Precipitation

The precipitation data from 2015-2023 indicate significant fluctuation in precipitation amounts and having a direct impact on the water availability for hydropower generation in the NO3 zone. From Figure 4.1.4, the monthly precipitation totals (mm) time series indicate the occurrence of high and low precipitation in alternate succession. There are greatly wet years (2018) where there was above-normal rainfall, potentially causing problems such as reservoir overflow, and dry years (2020) with water shortage problems for power generation. There is a clear seasonal pattern; rainfall rises in autumnal months, combined with springtime snowmelt, creating a two-peaked inflow regime necessary for reservoir release scheduling. According to the dataset, extreme events and interannual variability have increased recently. For instance, late 2022 had high rainfalls, and this suggests the necessity for efficient forecasting given that weather climate change has the potential to be rendering rainfall behavior high unpredictable. Such precipitation trends, especially when integrated with temperature and discharge records, provide critical information on optimizing reservoir management. Periods of over- and under-rainfall directly relate to generating capacity for power, so our simulations must become capable of anticipating such based on historical trends. Figure 4.1.4 presents the fluctuations in rain, further establishing that there is a demand for advanced forecast models that can handle such fluctuations in hydrology input data.

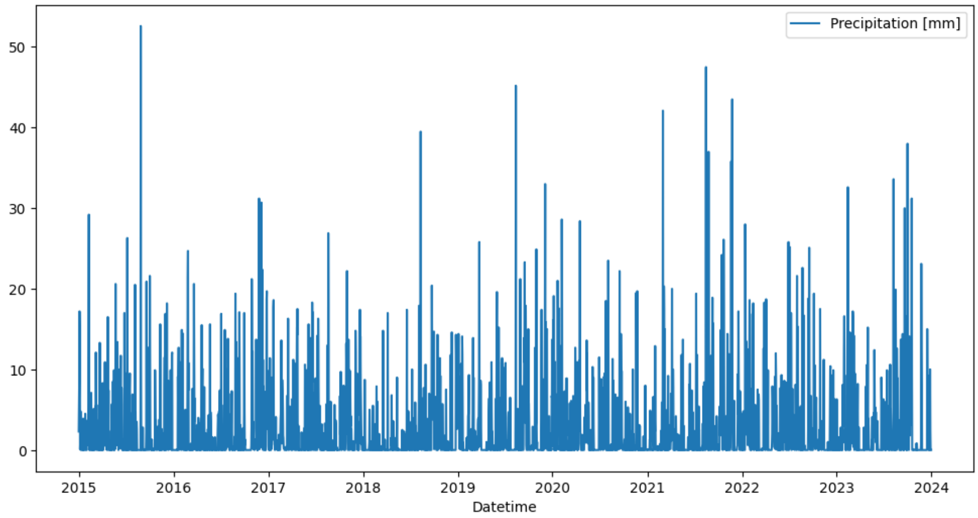

Figure 4.1.4: Precipitation (mm)

4.1.4 Wind Speed

Also under analysis were wind speed observations in the NO3 region, as indicated by Figure 4.1.5. Wind speed is an indirect driver of hydropower; it is relevant in the broader energy system context and can indirectly affect market conditions. Wind time series exhibits strong seasonality with the higher winds usually in winter months (frequently greater than 8–10 m/s during storms) and lower winds during summer (occasionally less than 4 m/s). Clear interannual tendencies are apparent: 2018 and 2022 stand out as particularly windy years, with more highwind days, which would have boosted wind power production in the region, whereas 2020 had long spells of low wind. The NO3 area has a winter-dominated wind regime that is complementary to hydropower – in times of low winter inflows (frozen months), there is high wind power available as an alternative power source. The records also reflect greater wind variability in recent years, which can be related to evolving weather patterns. The NO3 area has a winter-dominated wind regime that supplements the hydropower in the context that during low winter inflow (frozen), the wind power would be at its maximum, and it can be utilized as a back-up power source. The records also show heightened wind variability in the recent past, which may be a result of changing weather patterns. For our purposes, wind speed was included as a feature to capture these external influences on the power system. Although our model focuses on hydropower scheduling, considering wind availability can be important for anticipating market price fluctuations (e.g., abundant wind can reduce power prices). Figure 4.1.5 shows the wind speed trends and variability over the years, highlighting why integrated renewable resource data can be valuable in a comprehensive energy scheduling model.

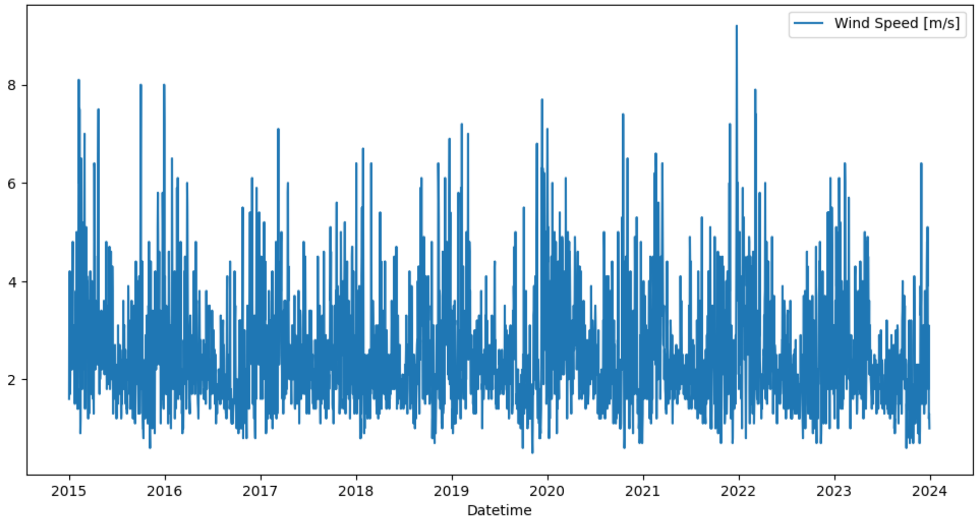

Figure 4.1.5: Wind Speed (m/s)

4.1.5 Hydropower Production

Hydropower generation in NO3 region from 2015 to 2023 shows clear annual variability, ranging from about 2,000 to 4,500 MW. Production peaks usually occur in spring and autumn, driven by snowmelt and rainfall, while winter shows lower output due to frozen inflows and reservoir conservation. Production instability in hydropower production increased after 2020 due to shifting weather patterns and market turbulence. The system reliance on water availability led to high production in 2018 and low production in 2020. The clear correlation between annual production levels and hydrological conditions validates the need for integrated forecasting systems that combine weather predictions, market signals, and reservoir analytics to optimize NO3 hydropower.

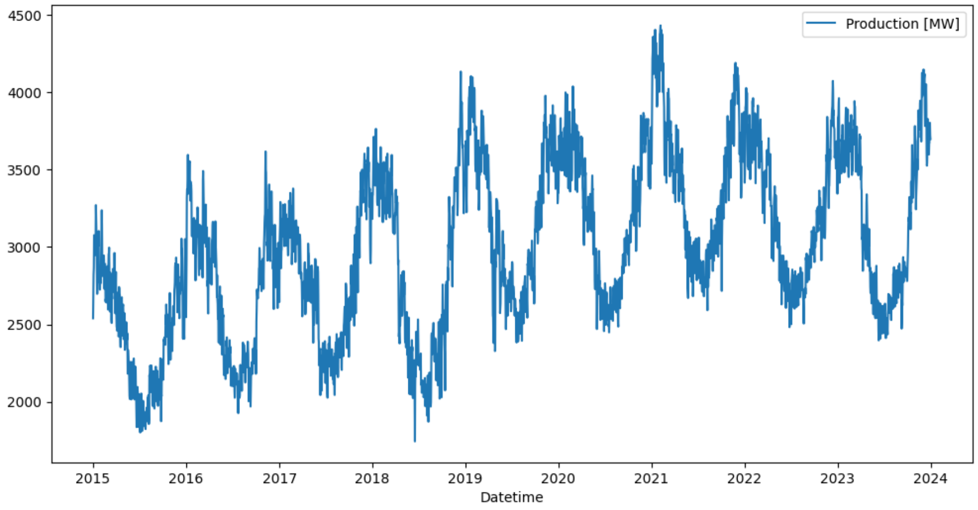

Figure 4.1.6: Production (MW)

4.1.6 Electricity Consumption

The electricity consumption in NO3 region between 2015 and 2023 shows the seasonal and interannual variability, with demand ranging between 2,000 and 4,000 MW. Peaks in winter often exceed 3,500 MW because of the increased demand for heating, while summer consumption ranges around an average of 2,500 MW. Some interesting trends include a decline in demand in 2020 because of the COVID-19 pandemic, followed by a rise and signs of rising demand after 2021 because more homes will have electricity. Changes in consumption patterns and extreme weather conditions after 2020 contrast with the relatively stable consumption patterns of 2015-2019. The load variability demands changes to the hydropower dispatch plan, and such changes necessitate the need for proper reservoir management in winter peak-demand seasons. The variability in load underlines the importance of demand forecasting in hydropower scheduling and emphasises the demand side.

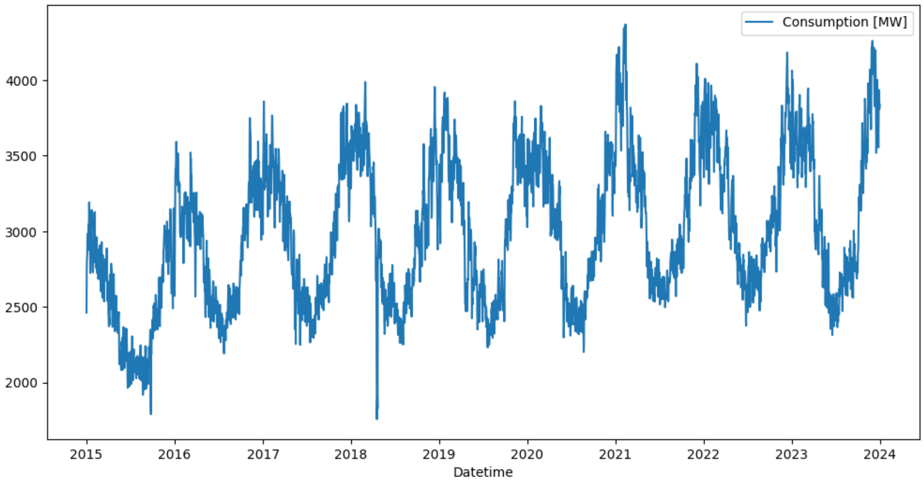

Figure 4.1.7: Consumption (MW)

4.1.7 Electricity Price

Electricity prices in the NO3 zone from 2015 to 2023 display hydrology and energy-driven volatility. Figure 4.1.8 illustrates low prices and sharp spikes. Prices rise in winter from higher demand and lower hydroelectric intakes, and fall in summer from surplus supply and lower demand. The statistics highlight two different periods. Between 2015 and 2019, there were stable prices as a result of normal hydrology and a balanced Nordic market. From 2020 and especially 2021 and 2022, the prices were extremely volatile with record highs from low reservoir levels, increasing European electricity prices, and increasing integration with the European market. Late 2022 recorded price peaks, highlighting NO3 increasing vulnerability to external shocks despite its hydropower dependence. These changes impact generation scheduling. Higher prices bring more production, and lower costs conserve water. Price spikes and drought highlight the value of merging market predictions and hydrological forecasts. Price is used as a forecast target in the model and includes past prices for making assumptions regarding the future. Figure 4.1.8 illustrates the impact of price changes on hydropower scheduling.

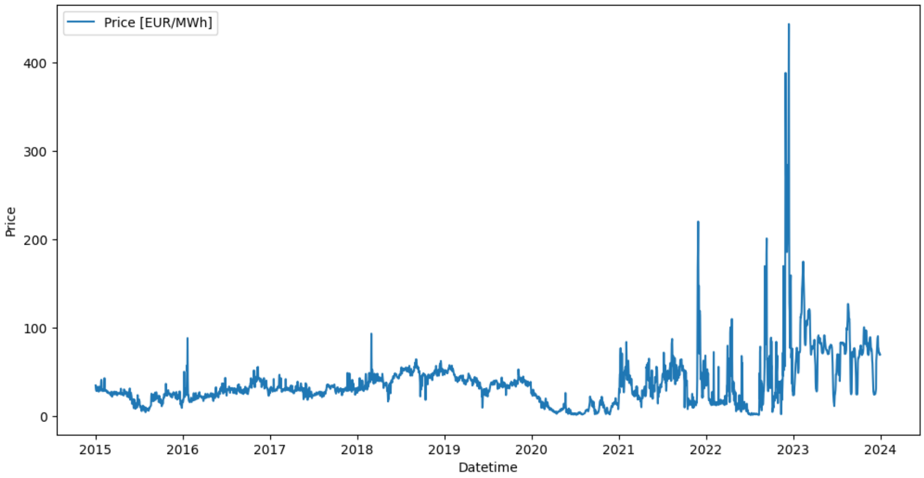

Figure 4.1.8: Price (EUR/MWh)

4.2 Case study: Bratsberg hydropower plant

To be able to test our model, we have found data online from the NVE website for the Bratsberg hydropower plant, a facility owned and operated by the integrated electric utility company Statkraft. Apart from sharing the relevant characteristics of the plant, NVE has also provided us with historical time series for inflow and production. The Bratsberg Hydropower Plant is a medium-sized and is located in the NO3 area in Central Norway. The plant is supplied by a single reservoir, Selbusjøen, which allows for controlled discharge operations with two Francis turbines and maximum discharge of each turbine is 60 m3/s. The main inflow to the reservoir comes from the Nea River (Neaelva), which originates from the Nea Valley and the Tydal mountains, receiving runoff from snowmelt, rainfall, and lakes. The outflow from Selbusjøen continues through the Nidelva River toward Trondheim with an average hydraulic head of 147 meters. The plant's operation is subject to physical reservoir constraints and local summertime regulatory restrictions, making it a representative case for studying long-term hydropower scheduling with real-world limitations. In Table 4.2.1, we have listed the physical boundaries of the reservoirs. The hydropower plant is by-passing a rather long section of the original river, and an environmental flow in this bypass reach is maintained at 30.0 m3/s, adding to the outlet flow from Bratsberg providing the flow in the lower parts in the river. So, when the Bratsberg power plant is shut down, there is still 30.0 m3/s in the river according to the new technology article [40].

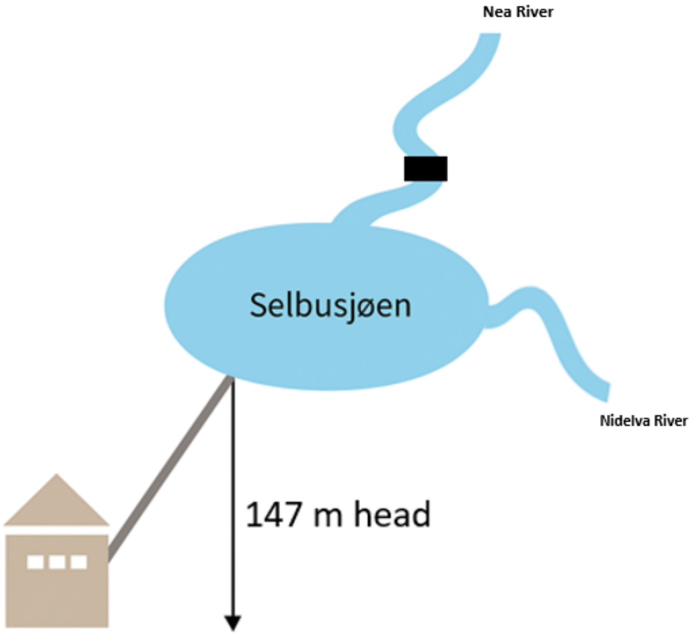

Figure 4.2.1: Bratsberg Hydropower plant

Characteristic	**Value**
Installed Capacity	124 MW
Annual Production	650 GWh
Water Source	Selbusjøen Reservoir
Gross Head	147 m
Number of Turbines	2 × Francis turbines
Tunnel Length	12 km
Max Turbine Discharge	60 m^3/s per turbine
Reservoir Water Height	155.0 m

Table 4.2.1: Characteristics of Bratsberg Hydropower plant

4.3 Machine learning models

After going through the relevant literature review, a Long short-term memory (LSTM) neural network was chosen as the best fit to be tested for implementation into the LTHPS. This ML technique was, as far as the author of this thesis knows, LSTM has not yet been applied specifically to hydropower scheduling. However, other neural network models, such as Recurrent Neural Networks (RNN), have shown promising results, but LSTM might give good results for this case study. The performance of the machine learning model is evaluated using NSE, MSE, MAE, RMSE and MAPE metrics.

4.4 Conclusion of Case study

In this chapter, the case study is presented. In addition, the method for generating data has been discussed to increase the dataset. The chapter ended with a discussion on which machine learning models are going to be used in this thesis and how it is are going to be evaluated. In the next chapter, the results will be presented.

Chapter 5

Results

In this chapter, we start by showing our experimental plan to answer our third research question on the effectiveness of machine learning techniques for long-term hydropower scheduling. We focus on four key forecasting tasks: inflow, production, consumption, and electricity price. For each task, we compare the performance of a simple RNN and an LSTM model highlighting their accuracy, efficiency, and comparative strengths. Model performance is evaluated using metrics such as NSE, MAE, MSE, RMSE, and MAPE, providing a comprehensive view of accuracy and consistency. The results are analyzed through tables and visualizations, highlighting the strengths and weaknesses of each model.

5.1 Inflow Forecasting

Inflow forecasting is essential for effective hydropower scheduling because reservoir levels and turbine operations directly depend on anticipated water availability. To evaluate the performance of the machine learning model for this task, both a Simple Recurrent Neural Network (RNN) and a Long Short-Term Memory (LSTM) network were trained to predict inflow volumes over the planning horizon. To performance of these models on the test dataset is summarized in Table 5.1.1, using key evaluation metrics.

Model	NSE	MAE	MSE	RMSE	MAPE
Simple RNN	0.5696	0.5201	0.6684	0.8175	90.97%
LSTM	**0.6988**	**0.4389**	**0.4677**	**0.6839**	**88.11%**

Table 5.1.1: Evaluation of RNN vs. LSTM for Inflow Forecasting

As shown in Table 5.1.1, the LSTM model performed better than the RNN in all the evaluation metrics. The LSTM achieved a higher Nash–Sutcliffe Efficiency (NSE = 0.6988), indicating that it explains a greater proportion of the variance in the inflow data relative to the mean inflow baseline. This implies a stronger alignment between the LSTM predictions and the actual observed inflows.

In terms of error metrics, LSTM also recorded a lower Mean Absolute Error (MAE = 0.4389) and Root Mean Squared Error (RMSE = 0.6839), suggesting that the deviations between its predictions and true values were both smaller and less variable. The Mean Absolute Percentage Error (MAPE) was also lower for the LSTM is 88.11% than the RNN is 90.97%, indicating predictions closer to the actual inflow values.

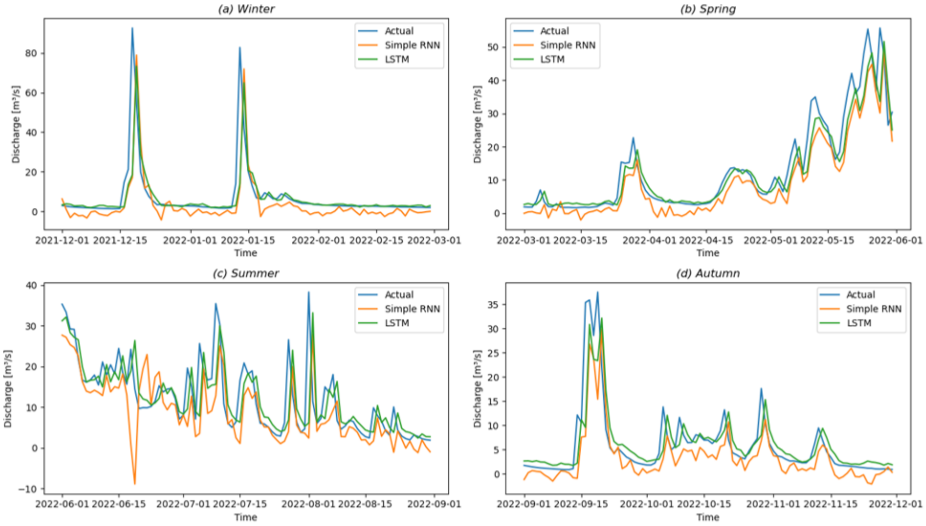

Figure 5.1.1: illustrates the predicted vs. actual inflows across different seasons.

The enhanced performance of the LSTM is a function of its design, which incorporates memory gates that are specifically designed to preserve information over longer sequences. Hydrological inflows typically exhibit strong seasonality patterns; such as snowmelt in spring or high autumn rainfall, which require the model to capture long-term dependencies. While RNNs may struggle with vanishing gradients over long sequences, the LSTM gated mechanisms allow it to preserve and utilize information from earlier time steps more effectively.

5.2 Production Forecasting

In table 5.2.1 the evaluation results for predicting the production in the No3 region are presented. This shows the Nash–Sutcliffe Efficiency (NSE), the mean square error (MSE), the mean average error (MAE), the root mean square error (RMSE), and the mean average percentage error (MAPE) for the two different machine learning techniques, RNN and LSTM in predicting hydropower generation across seasonal timeframes. MAPE is expressed as an percentage, while the other metrics are shown as decimal values.

Model	NSE	MAE	MSE	RMSE	MAPE
Simple RNN	0.8882	0.0450	0.0033	0.0576	31.62%
LSTM	**0.8949**	**0.0444**	**0.0031**	**0.0558**	**31.58%**

Table 5.2.1: Evaluation of RNN vs. LSTM for Production Forecasting

As shown in table 5.2.1, the LSTM model is better than the RNN across all evaluation metrics. The Nash–Sutcliffe Efficiency (NSE) for LSTM was 0.8949, slightly higher than the RNN value is 0.8882, indicating that LSTM explained more variance in the actual production data. Both models demonstrated strong predictive ability, but the LSTM provided slightly tighter agreement with observed generation values.

In terms of error metrics result, The LSTM model achieved a marginally lower MAE (0.0444) and RMSE (0.0558) compared to the RNN, suggesting that its predictions were both more accurate on average and exhibited fewer large deviations. The MAPE for both models was around 31.6%, but the LSTM still showed a modest edge in minimizing percentage errors.

The seasonal performance figure 5.2.1 illustrate the alignment between actual and predicted production values for both models in winter, spring, summer, and autumn. Both models follow production trends well for winter and autumn but notice slightly greater deviations in summer that can result from variable inflow conditions along with greater variability in demand. Despite these seasonal fluctuations, both models maintain high predictive accuracy, with LSTM consistently delivering smoother and more accurate forecasts. These results confirm the performance and suitability of LSTM for real-world hydropower generation forecasting scenarios.

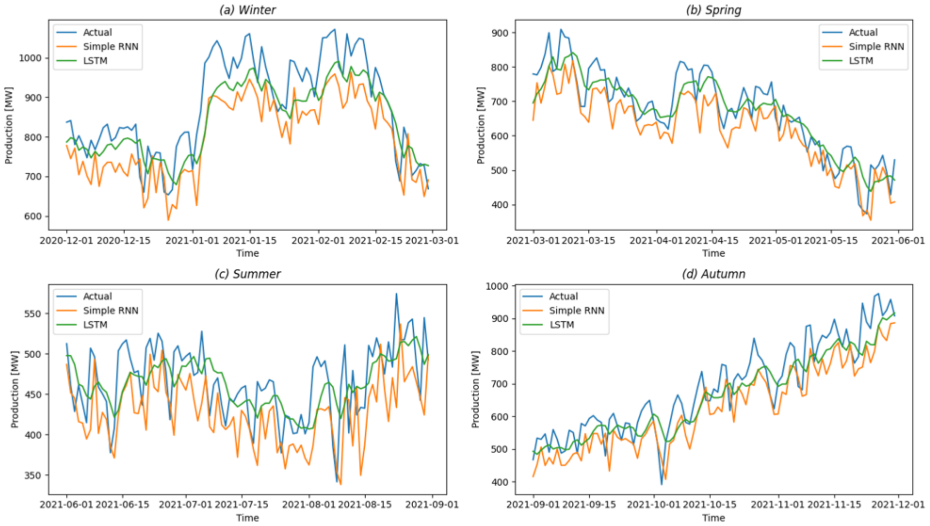

Figure 5.2.1: illustrates the predicted vs. actual production across different seasons.

5.3 Consumption Forecasting

In table 5.3.1 can the evaluation results for predicting the demand in the No3 region are presented. This shows the Nash–Sutcliffe Efficiency (NSE), the mean square error (MSE), the mean average error (MAE), the root mean square error (RMSE) and the mean average percentage error (MAPE) for the two different machine learning techniques, RNN and LSTM in predicting electricity consumption across seasonal timeframes. MAPE is expressed as an percentage, while the other metrics are shown as decimal values.

Model	NSE	MAE	MSE	RMSE	MAPE
Simple RNN	0.9049	0.0468	0.0031	0.0554	41.14%
LSTM	**0.9107**	**0.0413**	**0.0029**	**0.0537**	**38.65%**

Table 5.3.1: Evaluation of RNN vs. LSTM for Consumption Forecasting

As shown in table 5.3.1, the LSTM model is better than the RNN model across all evaluation metrics. The Nash–Sutcliffe Efficiency (NSE) for LSTM was 0.9107, slightly higher than the RNN value is 0.9049 and a lower Mean Absolute Percentage Error (MAPE) of 38.65%, compared to the Simple RNN value of NSE is 0.9049 and MAPE of 41.14%. These results highlight LSTM enhanced ability to model complex temporal patterns in consumption data with greater precision. The MAPE dropped from 41.14% with Simple RNN to 38.65% with LSTM, highlighting a notable improvement in percentage-based accuracy.

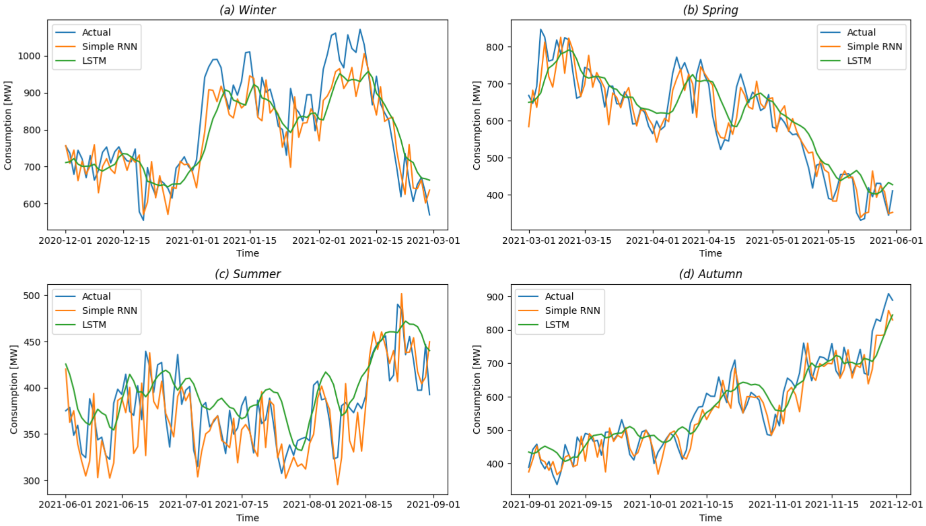

Figure 5.3.1: illustrates the predicted vs. actual consumption across different seasons

Figure 5.3.1 illustrate the seasonal consumption forecasts produced by both models compared with actual values. Both models follow the actual trends closely across seasons, though LSTM predictions tend to align more closely with observed peaks and troughs. During winter and autumn, the models capture strong seasonal demand increases, while summer and spring show more variability. LSTM consistently provides smoother and more stable forecasts, especially during periods of high fluctuation. These results confirm the effectiveness of LSTM in long-term electricity consumption forecasting, making it a valuable tool for hydropower system planning and decision-making.

5.4 Price Forecasting

In table 5.4.1 can the evaluation results for predicting electricity prices in the NO3 region are presented. This shows the Nash–Sutcliffe Efficiency (NSE), the mean square error (MSE), the mean average error (MAE), the root mean square error (RMSE) and the mean average percentage error (MAPE) for the two different machine learning techniques, RNN and LSTM in predict electricity prices across different seasons using historical price data.MAPE is expressed as an percentage, while the other metrics are shown as decimal values.

Model	NSE	MAE	MSE	RMSE	MAPE
Simple RNN	0.5315	0.0467	0.0062	0.0785	87.76%
LSTM	**0.6553**	**0.0454**	**0.0045**	**0.0665**	**55.58%**

Table 5.4.1: Evaluation of RNN vs. LSTM for Price Forecasting

As shown in table 5.4.1, the LSTM model is better than the RNN model across all evaluation metrics, with a higher NSE of 0.6553 compared to the NSE value of RNN is 0.5315. The LSTM also recorded reduced MAE (0.0454 vs. 0.0467) and RMSE (0.0665 vs. 0.0785). The MAPE value for LSTM are 55.58% is represents a substantial improvement over RNN is 87.76%, which reflecting a notable gain in relative forecasting accuracy. This improvement is particularly relevant in the energy markets, where even minor price mispredictions can lead to significant financial consequences.

Figure 5.4.1: illustrates the predicted vs. actual price across different seasons.

Figure 5.4.1 illustrates the actual and predicted electricity prices across all four seasons of 2021. The plots show that electricity prices show significant seasonal variability, especially in winter and autumn, often driven by demand surges and market dynamics. While both models generally follow seasonal price trends,the LSTM model showed a more precise tracking of price spikes and dips, particularly in colder months when price volatility was higher. RNN model tended to lag slightly in reacting to sudden changes. These findings underscore the pros of LSTM memory-based structure in dealing with the unpredictable and seasonal nature of electricity price data. Thus, LSTM is the more reliable model for this forecasting task, supporting more informed and profitable scheduling strategies.

5.5 Summary of Results

The experimental evaluation indicates that for long-term hydropower scheduling purposes, LSTM-based models offer superior forecasting accuracy for inflows, prices, and power production/ consumption. They capture temporal dependencies more effectively and thereby provide more reliable predictions, which is critical for optimizing operations. The RNN models, while faster and simpler, consistently underperformed in accuracy, making them less suitable as primary forecasting tools in this context. Going forward, the slight reduction in efficiency with LSTM is a reasonable trade-off for the gains in forecast precision. These findings support the recommendation to use LSTM (or similarly advanced recurrent architectures) in practical hydropower scheduling systems to improve decision making and overall system performance.

In the next chapter, the results will be discussed the optimization of scheduling decisions informed by these improved forecasts and potential ways to further enhance model efficiency without sacrificing accuracy.

Chapter 6

Discussion

This chapter discusses the results of the experimental analysis presented in chapter 5, providing a critical discussion of the forecasting inflow, electricity price, hydropower production, and consumption. The discussion relates these findings to the research questions and objectives, and specifies how the machine learning algorithms; Simple RNN and LSTM, have been able to perform in the context of long-term hydropower planning. It also evaluates the broader implications of the results for the energy sector and outlines the strengths and limitations of the model in practical applications.

6.1 Interpretation of Key Findings

The evaluation of forecasting performance shows that LSTM consistently performs better than the RNN model across all four forecasting tasks. In inflow forecasting, LSTM provided greater accuracy and better recognition of seasonal patterns, making it more suitable for hydrological time series with long-term dependencies. For energy price forecasts, where volatility and sharp changes are common, LSTM has captured trends and fluctuations more effectively, thereby reducing the risk of misjudging market behavior. Both in production forecasts and consumption forecasts, LSTM has again provided superior performance, especially in adapting to predicted and observed peak events.

These results confirm the assumption that models with memory capabilities (like LSTM) are more suitable for complicated and time-dependent datasets in hydroelectric systems. The consistent improvements in NSE, MAPE, and RMSE optimizations in all the scenarios indicate that the capability of LSTMs to learn long-term dependency is converted into practical benefits in forecast accuracy. This supports its integration into hydroelectric energy planning operation procedures.

6.2 Evaluation of the Model's Strengths

The main strength of the LSTM model lies in its long-term memory mechanism, which effectively handles for long-term dependencies in time series data. This was particularly advantageous for influx and price predictions, where trends extend over weeks or months. The model also showed resilience to sudden changes and seasonal variations, which led to a reduction in percentage errors across all activities.

From a practical point of view, the accuracy of LSTM improves decision-making in watershed management and energy market operations. It provides a better understanding for turbine planning, water discharges, and bidding strategies. Moreover, both models; RNN and LSTM were able to learn with relatively modest data sets and data processing resources, which suggests their adaptability to other regions or hydroelectric plants with similar data limitations.

6.3 Limitations of the Results

Despite the promising performance of the models, some limitations can be mentioned. First, the models were trained on historical data for a Norwegian region (NO3), and generalization to other geographic locations or market conditions may require to be further testing. Secondly, although the daily resolution of the data was sufficient for long-term planning, an increased temporal resolution (e.g., daily) could reveal further insights, but also increase the model complexity and computational costs.

A limitation is that exogenous externalities like policy changes, random power outages, or severe weather conditions have not been modeled explicitly but affect the forecast reliability. Also, while the LSTM model was more precise than the RNN model, it was slower to train and required more parameters to tune, which may be a shortcoming for real-time applications.

6.4 Unanticipated Findings

While the primary aim was to evaluate the forecasting accuracy of RNN and LSTM models, there were some surprising discoveries through experimentation. One of these was the Simple RNN relatively good performance in predicting consumption, particularly at time intervals with steady seasonal patterns. Although overall it was worse than LSTM, it was unexpectedly competitive under low-variance conditions. This implies that for some less volatile stable time series, simpler models can still yield satisfactory accuracy and conserve computational time.

The other surprise result was the relative stability of the training time of the LSTM model across different forecasting applications despite the various volatility and complexity of the data. This implies that once calibrated, LSTM models possess stable performance across different domains and hence can be useful tools even beyond hydropower scheduling.

Finally, the error analysis showed that both models occasionally struggled with predicting extreme events (e.g., sudden inflow spikes due to storms or sudden price jumps). This points to the value of exploring hybrid or ensemble models that will better capture rare but significant events.

6.5 Broader Impact on Energy Policy and Sustainability

The results of this thesis hold significant implications for sustainability and energy policy. Better long-term prediction of hydropower inflows, electricity prices, production, and demand allows hydropower resources to be more suitable and optimally scheduled. By improving the precision of these forecasts, utility operators can reduce on fossil fuel backup systems, contributing to a more sustainable and decarbonized energy mix.

Moreover, the findings support integrating machine learning into the operating systems of grid management and market planning. This is all within the broader trend of smart energy systems, where data-driven decision-making enhances the efficiency, robustness, and environmental sustainability of power systems.

From a policy perspective, improved forecasting supports better management of water resources, especially in Regions No.3, where hydropower is a major contributor to electricity supply. It also enhances transparency in market pricing, potentially informing regulatory strategies and helping prevent price manipulation.

Chapter 7

Conclusions

This master's thesis presented an overview of hydropower scheduling and machine learning through an extensive theory chapter, a literature review, an in-depth model presentation, and a case study containing a hydropower plant with one reservoirs. The literature review gilded the state-of-the-art long-term hydropower scheduling and a possible machine learning to test the optimization model. The case study used real life data provided by professor for a hydropower plant with 1 reservoirs. Using real-world data from the Norwegian NO3 bidding zone (2015–2023) from NVE, the study demonstrated the feasibility and benefits of applying deep learning models to support decision-making in renewable energy systems. The case study was conducted using two machine learning techniques: Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) models to the forecasting tasks critical for long-term hydropower scheduling. This included inflow prediction, electricity price forecasting, and the estimation of both hydropower production and consumption.

Among the models evaluated, LSTM consistently outstanding performance than the simple RNN architecture across all forecasting tasks. LSTM models achieved higher predictive accuracy (as evidenced by better NSE, lower MAE, MSE, RMSE, and MAPE), particularly in datasets characterized by long-term dependencies or high volatility—such as river inflows and electricity market prices. While RNN were faster to train and more lightweight, their performance suffered from limitations in memory retention and difficulty learning from long sequences.

In conclusion, this master thesis confirms that machine learning models, in this case, LSTM networks can be applied as effective optimization techniques in long-term hydropower scheduling. Their ability to deliver more reliable forecasts can lead to substantial benefits in the operation of hydropower system. Such benefits are enhanced reservoir management, improved electricity generation planning, and more profitable market participation. The results validate the integration of LSTM-based models into hydropower scheduling systems and prove the pivotal role of data-driven tools in the management of renewable energy.

7.1 Future work

This chapter will discuss the areas of this master's thesis, several directions for further research and enhancement are recommended:

- **Model Sensitivity Analysis**: Future work should investigate the sensitivity of forecasting models to key input variables such as precipitation, temperature, and previous inflow. Understanding which features most strongly influence the accuracy of the model can guide improved data collection strategies and model refinement.

- **Direct Integration into Hydropower Scheduling**: Although this study focused on forecasting tasks, an important next step is to directly integrate these forecasts into real-time hydropower scheduling frameworks. This would involve combining predictive models with reservoir optimization systems to inform decisions on water, storage management, and electricity generation decisions.

- **Forecast Accuracy Improvements**: Accuracy is also improved by trying ensemble learning methods (e.g., combining multiple LSTM or hybrid models together), season model tuning (e.g., creating separate models for winter and summer), or including external variables such as energy market signals and climate indices.

- **Explainability and Interpretability**: Improving the transparency of ML models through explainable AI techniques can support trust and adoption among hydropower operators and policymakers.

- **Application to Multiple Reservoirs**: The current framework can be extended to more complex hydropower systems involving multiple interconnected reservoirs. This would allow for more realistic simulations of cascading water systems and upstream-downstream dependencies.

- **Multiple Power Stations and Geographic Transferability**: To assess generalizability, the models should be tested in different power plants located in various regions with different hydrological and climatic patterns. This would validate the robustness of the approach under diverse operating conditions and inform a wider deployment.

Final Remarks: This thesis demonstrated the strong potential of machine learning to improve the forecasting and planning of hydro power system. With careful selection of model, data handling, and integration into operational workflows, these tools can significantly contribute to smarter, more sustainable energy systems. As energy systems become more data-driven and reliant on renewable sources, tools such as LSTM will play a central role in bridging the gap between data and optimized operations.

Bibliography

[1] Energifakta Norge, "Kraftproduksjon," energifaktanorge, 2024. [Online]. Available: https://energifaktanorge.no/en/norsk-energiforsyning/kraftproduksjon/

[2] S. Jaehnert, "Hvorfor vannverdi er viktig for strømpris og forsyningssikkerhet," Sintef, 2023. [Online]. Available: https://blogg.sintef.no/energi/vannverdi-er-viktig-for-strompris-og-forsyningssikkerhet/

[3] I. Graabak, S. Jaehnert, M. Korpås, and B. Mo, "Norway as a battery for the future european power system–comparison of two different methodological approaches," in *The International Workshop on Hydro Scheduling in Competitive Markets.* Springer, 2018, pp. 76–83.

[4] I. Graabak, S. Jaehnert, M. Korp, and B. Mo, "Norway as a battery for the future european power system—impacts on the hydropower system," *Energies*, vol. 10, no. 12, p. 2054, 2017.

[5] V. Koestler, A. Østenby, C. Birkeland, F. Arnesen, and I. Haddeland, "Vannkraftverkene i norge får mer tilsig," *Norges vassdrags-og energidirektorat*, pp. 1–32, 2019.

[6] G. L. Doorman, "Course elk: Hydro power scheduling," 2018, lecture material.

[7] A. Gjelsvik, B. Mo, and A. Haugstad, "Long-and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming," *Handbook of power systems I*, pp. 33–55, 2010.

[8] "Hydropower," Energy Information Administration, 2022, april 20, 2025. [Online]. Available: https://www.eia.gov/energyexplained/hydropower/

[9] Omnicalculator, “Hydroelectric power calculator,” 2023, april 21, 2025. [Online]. Available: https://www.omnicalculator.com/ecology/hydroelectric-power

[10] S. Vaca-Jiménez, P. Gerbens-Leenes, and S. Nonhebel, “The monthly dynamics of blue water footprints and electricity generation of four types of hydropower plants in ecuador,” *Science of The Total Environment*, vol. 713, p. 136579, 2020.

[11] B. A. Nasir, “Design considerations of micro-hydro-electric power plant,” *Energy procedia*, vol. 50, pp. 19–29, 2014.

[12] K. Kumar and R. Saini, “A review on operation and maintenance of hydropower plants,” *Sustainable Energy Technologies and Assessments*, vol. 49, p. 101704, 2022.

[13] S. Jaehnert, “Hvorfor vannverdi er viktig for strømpris og forsyningssikkerhet,” *SINTEF Energy Blog*, 2023, april 22, 2025. [Online]. Available: https://blogg.sintef.no/sintefenergy-nb/vannverdi-er-viktig-for-stromprisog-forsyningssikkerhet/

[14] A. Olabi, K. Elsaid, K. Obaideen, M. A. Abdelkareem, H. Rezk, T. Wilberforce, H. M. Maghrabie, and E. T. Sayed, “Renewable energy systems: Comparisons, challenges and barriers, sustainability indicators, and the contribution to un sustainable development goals,” *International Journal of Thermofluids*, vol. 20, p. 100498, 2023.

[15] G. L. Doorman, *Course Lect-11 Hydro Power Scheduling.* NTNU, 2018, course material.

[16] X. Wu, S. Yin, C. Cheng, Z. Chen, and H. Su, “Ssdp model with inflow clustering for hydropower system operation,” *Water Resources Management*, vol. 37, no. 3, pp. 1109–1123, 2023.

[17] J. Alzubi, A. Nayyar, and A. Kumar, “Machine learning from theory to algorithms: an overview,” in *Journal of physics: conference series*, vol. 1142. IOP Publishing, 2018, p. 012012.

[18] R. Remesan and J. Mathew, “Hydrological data driven modelling,” *Earth System Data and Models*, vol. 1, 2015.

[19] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *nature*, vol. 323, no. 6088, pp. 533–536, 1986.

[20] V. Øyri, “Long short-term memory (lstm) recurrent neural networks for urban hydrological modelling,” Master’s thesis, 2020.

[21] C. Olah. (2015, 08) Understanding LSTM networks, blog post. April 22, 2025. [Online]. Available: https://colah.github.io/posts/2015-08-Understanding-LSTMs/

[22] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[23] A. Graves and J. Schmidhuber, “Offline handwriting recognition with multidimensional recurrent neural networks,” *Advances in neural information processing systems*, vol. 21, 2008.

[24] Z. Zhao, W. Chen, X. Wu, P. C. Chen, and J. Liu, “Lstm network: a deep learning approach for short-term traffic forecast,” *IET intelligent transport systems*, vol. 11, no. 2, pp. 68–75, 2017.

[25] R. Jozefowicz, W. Zaremba, and I. Sutskever, “An empirical exploration of recurrent network architectures,” in *International conference on machine learning.* PMLR, 2015, pp. 2342–2350.

[26] V. Highfield. (2015) Google voice improves speech-to-text using the power of ’thinking computers’. Alphr. April 26, 2025. [Online]. Available: https://www.alphr.com/technology/1001258/google-voice-improves-speech-to-text-using-the-power-of-thinking-computers/

[27] D. Harris-Birtill and R. Harris-Birtill, “Understanding computation time: a critical discussion of time as a computational performance metric,” in *Time in Variance.* Brill, 2021, pp. 220–248.

[28] B. Schaefli and H. V. Gupta, “Do nash values have value?” *Hydrological processes*, vol. 21, pp. 2075–2080, 2007.

[29] H. in a Machine World. (2018) Mae and rmse: Which metric is better? Medium. Online article. [Online]. Available: https://medium.com/human-in-a-machine-world/mae-and-rmse-which-metric-is-better-e60ac3bde13d

[30] A. H. Murphy, “Skill scores based on the mean square error and their relationships to the correlation coefficient,” *Monthly weather review*, vol. 116, no. 12, pp. 2417–2424, 1988.

[31] W. Wang and Y. Lu, “Analysis of the mean absolute error (mae) and the root mean square error (rmse) in assessing rounding model,” in *IOP conference series: materials science and engineering*, vol. 324. IOP Publishing, 2018, p. 012049.

[32] M. V. Pereira and L. M. Pinto, “Multi-stage stochastic optimization applied to energy planning,” *Mathematical programming*, vol. 52, pp. 359–375, 1991.

[33] A. Turgeon, “A decomposition method for the long-term scheduling of reservoirs in series,” *Water Resources Research*, vol. 17, no. 6, pp. 1565–1570, 1981.

[34] A. Helseth and B. Mo, “Hydropower aggregation by spatial decomposition—an sddp approach,” *IEEE Transactions on Sustainable Energy*, vol. 14, no. 1, pp. 381–392, 2022.

[35] R. Remesan and J. Mathew, “Hydrological data driven modelling,” *Earth System Data and Models*, vol. 1, 2015.

[36] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[37] V. Mishra and A. D. Tiwari, “Sub-seasonal prediction of drought and streamflow anomalies for water management in india,” *Journal of Geophysical Research: Atmospheres*, vol. 127, no. 3, p. e2021JD035737, 2022.

[38] Z. Liu, J. Zhou, X. Yang, Z. Zhao, and Y. Lv, “Research on water resource modeling based on machine learning technologies,” *Water*, vol. 16, no. 3, p. 472, 2024.

[39] J. B. Slimane, “Deep hybrid neural networks for predicting missing segments in semg time series data,” *International Journal on Information Technologies & Security*, vol. 16, no. 3, 2024.

[40] O. Saberi, P. Storli, and K. Alfredsen, “New technology to increase hydropower plant operational flexibility,” *Int. J. Hydraul. Eng*, vol. 10, pp. 1–7, 2021.

Appendices

Table A: Descriptive Statistics of Input Features

Feature	**Count**	**Mean**	**Std**	**Min**	**25%**	**50%**	**75%**	**Max**
Discharge	3287.0	7.6556	9.8076	0.1609	1.7484	3.6827	9.4701	92.7194
Evaporation	3287.0	41.8873	1616.2535	0.2000	0.5000	1.1000	2.5000	65535.0000
Precipitation	3287.0	2.7987	5.1685	0.0000	0.0000	0.3000	3.5000	52.6000
Temperature	3287.0	5.1762	7.1165	-17.3000	0.2000	4.7000	10.7000	25.2000
Wind speed	3287.0	2.5151	1.1485	0.5000	1.7000	2.2000	3.0000	9.2000
Price	3287.0	36.4498	30.0427	1.0900	21.5200	30.5500	44.1500	443.7400
Production	3287.0	3003.8748	543.8285	1743.5400	2612.9800	2968.4600	3448.8150	4431.5800
Consumption	3287.0	2998.3915	482.1379	1756.7500	2608.8750	2954.8300	3384.6250	4368.6700