

Short-term hourly hydropower prediction: Evaluating LSTM and MILP-based methods

Y. Villeneuve, S. Séguin, A. Chehri, K. Demeester

G-2025-34

April 2025

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : Y. Villeneuve, S. Séguin, A. Chehri, K. Demeester (avril 2025). Short-term hourly hydropower prediction: Evaluating LSTM and MILP-based methods, Rapport technique, Les Cahiers du GERAD G- 2025-34, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2025-34>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: Y. Villeneuve, S. Séguin, A. Chehri, K. Demeester (April 2025). Short-term hourly hydropower prediction: Evaluating LSTM and MILP-based methods, Technical report, Les Cahiers du GERAD G-2025-34, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2025-34>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2025
– Bibliothèque et Archives Canada, 2025

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2025
– Library and Archives Canada, 2025

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Short-term hourly hydropower prediction: Evaluating LSTM and MILP-based methods

Yoan Villeneuve ^{a, b}

Sara Séguin ^{a, b}

Abdellah Chehri ^c

Kenjy Demeester ^d

^a Université du Québec à Chicoutimi (Qc), Canada, G7H 2B1

^b GERAD, Montréal (Qc), Canada, H3T 1J4

^c Royal Military College of Canada, Kingston (On), Canada, K7K 7B4

^d Rio Tinto Complexe Jonquière, Saguenay (Qc), Canada, G7S 4L2

yoan.villeneuve1@uqac.ca

s1seguin@uqac.ca

abdellah.chehri@rmc-cmr.ca

Kenjy.Demeester@riotinto.com

April 2025
Les Cahiers du GERAD
G–2025–34

Copyright © 2025 Villeneuve, Séguin, Chehri, Demeester

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : Hydropower generation plays a crucial role in the global energy landscape, offering a renewable and sustainable source of electricity. Accurate forecasting of hydropower output is essential for efficient energy management and maintaining grid stability. This paper presents an autoregressive Long Short-Term Memory (LSTM) model designed to predict short-term hydropower production, specifically targeting the hourly water output decisions of two interconnected hydropower plants located on the Péribonka River in Québec, Canada. Given the critical role of efficient scheduling in hydropower operations, especially within the Short-Term Hydropower Scheduling (STHS) problem, our model aims to offer a viable machine learning-based solution to complement traditional optimization approaches. We evaluated the LSTM model by comparing its predictive performance with historical operational data and results derived from a deterministic Mixed-Integer Linear Programming (MILP) model. Our analysis covers multiple validation instances, showcasing the capabilities of the model and highlighting its strengths and limitations. The results demonstrate that the autoregressive LSTM approach successfully captures the underlying patterns in water discharge decisions, providing predictions that are generally aligned with operational realities and optimized benchmarks. However, the study also underscores challenges such as maintaining reservoir volume constraints, particularly in periods of high inflow variability. Despite these challenges, the LSTM model presents promising predictive performance, laying the foundation for further improvements in integrating machine learning into short-term hydropower management. To our knowledge, this is the first study to apply an autoregressive supervised LSTM model to predict hourly water flow decisions in hydropower systems, thus significantly contributing to the advancement of machine learning applications in hydropower scheduling.

Keywords: Hydropower generation, long short-term memory, machine learning, optimization, Mixed-Integer Linear Programming, operational research

Résumé : La production hydroélectrique joue un rôle crucial dans le paysage énergétique mondial en offrant une source d'électricité renouvelable et durable. La prévision de la production hydroélectrique est essentielle pour une gestion efficace de l'énergie et le maintien de la stabilité du réseau électrique. Cet article présente un modèle autorégressif à mémoire court et long terme (LSTM) conçu pour prédire la production hydroélectrique à court terme, en ciblant spécifiquement les décisions horaires de débit d'eau de deux centrales hydroélectriques interconnectées situées sur la rivière Péribonka au Québec, Canada. Compte tenu du rôle critique d'une planification efficace dans les opérations hydroélectriques, particulièrement dans le cadre du problème de planification hydroélectrique à court terme (STHS), notre modèle propose une solution viable basée sur l'apprentissage automatique pour compléter les approches d'optimisation traditionnelles. Nous avons évalué le modèle LSTM en comparant ses performances prédictives avec des données opérationnelles historiques et des résultats issus d'un modèle déterministe de programmation linéaire mixte en nombres entiers (MILP). Notre analyse couvre plusieurs instances de validation, mettant en évidence les capacités du modèle ainsi que ses forces et limites. Les résultats démontrent que l'approche autorégressive LSTM parvient à capturer efficacement les motifs sous-jacents dans les décisions de débit d'eau, offrant des prédictions généralement alignées avec les réalités opérationnelles et les solutions optimisées. Toutefois, l'étude souligne également des défis, tels que le respect des contraintes de volume des réservoirs, en particulier lors des périodes de forte variabilité des apports en eau. Malgré ces défis, le modèle LSTM présente des performances prédictives prometteuses, posant ainsi les bases pour des améliorations futures dans l'intégration de l'apprentissage automatique à la gestion hydroélectrique à court terme. À notre connaissance, cette étude est la première à appliquer un modèle LSTM autorégressif supervisé pour prédire les décisions horaires de débit d'eau dans les systèmes hydroélectriques, contribuant ainsi significativement à l'avancement de l'apprentissage automatique dans la planification hydroélectrique.

Mots clés : Génération hydroélectrique, réseau récurrent à mémoire court et long terme, apprentissage automatique, optimisation, programmation linéaire en nombres entiers mixtes, recherche opérationnelle

1 Introduction

Hydropower generation plays an essential role in the global energy landscape, providing a renewable and sustainable source of electricity. In the province of Québec, Canada, hydropower accounts for 95% of the province's electricity needs and represents more than half of the annual electricity generated in the country [9, 21]. The effective scheduling of hydropower operations is crucial to ensure optimal use of water resources while meeting energy demand, especially given the current context of climate change [9]. Due to the scale and complexity of the hydropower scheduling problem, it is divided into three main optimization time frames, which are short-, medium-, and long-term scheduling problems. Short-term optimization models focus on optimizing day-to-day production to use the plant's components and available resources in the most efficient way possible [4, 8]. Medium-term optimization usually revolves around optimizing the reservoir volume of one or many hydropower plants for the foreseeable future. These problems involve a great deal of uncertainty caused by the water inflows, but also the prices when the hydropower system is in a deregulated market setting [17]. Long-term planning is used over a horizon of several years in order to quantify the impact of a major modification to the network. Long-term models are used to plan the expansion of a system or to measure the impact of maintenance on a turbine generator unit, to name a few. The models in this term are not used on a regular operational basis. The literature also discusses a fourth term, real-time scheduling, which deals with optimizing plant production in real time [7, 17]. Various methods have been developed to optimize hydropower operations, ranging from traditional mathematical optimization techniques [4], to more recently machine learning algorithms, as reviewed in [60, 63]. The Short-Term Hydropower Scheduling (STHS) problem has not seen as much advancement in the realms of machine learning as the medium- and long-term problems [8].

Papers were published in the last few years to review the state-of-the-art of the STHS problem. Some papers [4, 38] explain the process of formulating the hydropower scheduling problem and discuss some of the many articles on the optimization of the STHS problem. Most of these review articles for STHS encompass the unit commitment problem (UC) and unit load distribution problem (ULD), collectively known as the Hydro Unit Commitment (HUC) problem [38]. It is noted that there is potential to develop larger models merging different time horizons, given a feasible and practical model, which could lead to an improvement in the quality of the solution. Some reviews [7, 8, 60, 63] are related to machine learning in the field of hydropower production. Machine learning is still in an exploratory stage in the field of hydropower, as the few papers found in the literature demonstrate. In addition, most of the articles published focus on the medium-term inflow forecasting problems. This could be explained by the fact that the STHS is already a mathematically difficult problem, being nonconvex and nonlinear, coupled with the short delay for decision-making and the need of a large dataset to represent this problem over many years. In this review [60], the authors mention how artificial intelligence algorithms outperform classical optimization techniques in solving multi-objective optimization problems, while the contrary is also true for single-objective problems, with the added benefits of stability. Another issue mentioned by [7] is the few models that make use of reinforcement learning and the lack of models developed for run-of-river plants. Run-of-river hydropower plants are optimized by the construction design of the plant [62] and the production function is linear and depends on the flow of the river with little or no water storage. In this sense, optimizing the hydropower production of run-of-river plants is difficult since there is no water storage and that the water discharge is mainly controlled by the flow of the river, but a few articles [1, 62] shows that there is an interest in developing models for these types of hydropower plants.

Dynamic Programming (DP) is widely employed due to its ability to handle complex decision-making processes. This approach involves breaking down the scheduling problem into smaller and easier sub-problems and then iteratively solving each sub-problem to find the overall optimal solution [6]. Dynamic programming [2] is used for medium- and long-term hydropower problems, but progress in this field is slow, mainly due to the lack of scalability when faced with complex problems, where it usually falls short compared to metaheuristic models. In [14], a stochastic sampling DP

algorithm is presented for medium-term water management in multiple reservoirs. Using production function approximations and a linear objective function, the model efficiently predicts midterm water inflows, optimizing energy production for a system with four reservoirs. Stochastic Dual Dynamic Programming (SDDP) is implemented in [28] for medium-term hydropower scheduling problems to optimize hydropower production by effectively addressing the nonconvexity of the power and discharge function with the stochastic elements of market value and water inflows. Integer cuts are used to help with convergence, which resulted in an enhanced decision-making process, especially when omitting the energy's market price. The SDDP algorithm is also presented in [26] to achieve the highest possible economic outcome while considering environmental and operational constraints. Constraint relaxation and time-dependent auxiliaries on lower reservoir volume bounds are combined to address state-dependent maximum discharge constraints. Although the proposed approach increases expected annual profits, it also leads to overly conservative reservoir management strategies. A DP model that uses successive approximation and relaxation strategies is proposed by [25] to address the long-term joint power generation scheduling problem for large systems, achieving better results in terms of total power generation and calculation time compared to other methods. In [49], an innovative model is developed for the estimation of water with a precise representation of inflows and volume-dependent environmental constraints, leading to the integration of nonconvexity in the problem formulation. Development with DP for the STHS problem is often overlooked because of the high computation time constraint and computational difficulty for large hydropower systems. Recent advances in STHS with DP include the ULD model [33] for the hydropower plants with multiple units and serving multiple power grids. With a multicore parallel DP method, it is shown that it is possible to optimize a large ULD problem, such as the 18 units of Xiluodu stations in China, in a reasonable amount of time. While DP can yield accurate results, a major drawback is that computational complexity grows exponentially with the size of the problem, which is referred to as the infamous "curse of dimensionality". This limiting factor poses a significant challenge in scaling up large-scale hydropower solutions with DP.

Mixed-Integer Linear Programming (MILP) models are widely used for hydropower scheduling [42, 51, 63]. MILP formulations permit the inclusion of various constraints, allowing for a better representation of the STHS problem, such as reservoir storage, water release policies, and energy generation targets, to name a few [4, 18, 39]. In [57], a two-phase STHS optimization model is developed to first obtain the water discharge, the volume of the reservoir and the number of units working in each period, then determine which combination of turbines to use. Splitting the optimization in a two-stage model makes the computation time much shorter than with a larger model. A study from [15] has demonstrated the effectiveness of MILP models to optimize the short-term hydropower unit commitment problem. The problem is solved using efficiency points to represent the water discharge and the power produced at the maximum reservoir capacity for every combination of turbines. The solution is then adjusted to the current state of the reservoir, achieving improvements in operational efficiency and energy production. The paper by [52] introduces an MILP model for daily operations that convert the nonconvex and nonlinear Hydropower Production Function (HPF) to a convex and linear approximation by using three heuristics in the model. Each heuristic offers either dynamically adjusts the unit input-output curve or to have a stable unit input-output curve. A comparison between a nonlinear stochastic MILP model and a nonlinear heuristic model in a deregulated market is developed in [30]. These models take into consideration the European Nordic electricity markets, with efficient bidding in the day-ahead market for price-taking hydropower producers. Although the nonlinear MILP approach yields better results, the heuristic method can provide a viable solution within a shorter time frame. Similarly to dynamic programming, MILP models for the STHS problem suffer from a scalability problem related to the computational load, which does not fit well with the time frame available to make a decision in a short-term setting.

In addition to traditional MILP techniques, machine learning algorithms have emerged as promising tools for hydropower scheduling, with a lot of recent papers published on their uses in hydropower [8, 63]. Problems related to the hydropower scheduling problem with machine learning in

medium-term [16, 36, 46], and long-term horizons [19, 41], have been studied in the recent literature, but very few advancements have been made in regards to machine learning in the STHS problem.

Machine learning is used to address the limitations of optimization models, mainly to decrease the complexity of the hydropower scheduling problem when scaling the model [60]. In [32], the Approximate Dynamic Programming (ADP) method introduced in the study improves the optimization of hydropower reservoir by leveraging a Long Short-Term Memory (LSTM) model as a response surface model. The method addresses the challenges of traditional DP methods, offering a more efficient and accurate approach to reservoir operation optimization by reducing computation time and improving performance through accurate power output estimations. The research from [31] also proposes an ADP algorithm for the HUC problem. This approach preserves the nonconvex and nonlinear nature of the STHS problem, opting to map the original model's value function with machine learning. This adaptation enables hourly predictions and achieves a favourable balance between efficiency and solution optimality. In [66], a short-term model is developed, using various machine learning and optimization techniques, to predict the daily energy production of Mahabad Dam, located in the province of West Azerbaijan in Iran. The inputs of this model are transformed into frequency, then, an LSTM is used to capture the temporal dependencies and patterns in the data, followed by a random forest to make the prediction on the final output. This model marks a notable progress by applying various algorithms and techniques effectively, which results in improved predictive accuracy. Using 38 years of daily hydropower generation data, the study presented in [48] evaluated the performance of three distinct machine learning algorithms : Autoregressive Integrated Moving Average (ARIMA), Artificial Neural Network (ANN), and Support Vector Machine (SVM). The algorithms were then tested with three scenarios: daily, monthly and seasonal power generation patterns. Each model uses a window of the plant's previous power output as input, adjusting the window size and time frame according to the observed scenario. The case study consists of forecasting power generation at the Three Gorges Dam in Hubei Province. In general, SVM and ANN exhibit strong predictive performance in forecasting hydropower when using wider windows as input. For the daily scenario, both ANN and SVM demonstrate strong predictive capabilities, which highlights the effectiveness of machine learning algorithms in optimizing hydropower generation predictions for daily operations. Similarly, [47] forecasted changes in reservoir water level in Malaysia with two different time frames (daily and weekly) with various machine learning algorithms : Boosted Decision Tree Regression (BDTR), Decision Forest Regression (DFR), Bayesian Linear Regression (BLR) and Neural Network Regression (NNR). Using input data on water levels, rainfall, and water sent out of the Kenyir Dam, an analysis of the results shows that BDTR gives the best prediction for both daily and weekly forecasts.

This paper aims to contribute to the field of STHS by developing and evaluating an LSTM model to predict future hourly water output decisions of a system composed of two hydropower plants. The performance of the LSTM model is compared with the results of a deterministic MILP model [15] and real-life decisions, all of which are made on Rio Tinto's Chute-du-Diable and Chute-à-la-Savane hydropower plants in the province of Québec, Canada. The key criterion to evaluate the performance of the LSTM model is the total water flow to be processed by the plant, the volume of water of the plant's reservoirs and the energy produced for each hour of the planning horizon. This paper illustrates the potential of machine learning algorithms as a viable type of algorithm for hydropower scheduling in the future. By demonstrating the ability of an LSTM model to predict hydropower scheduling decisions, this research contributes to the advancement of efficient and sustainable hydropower operations. To our knowledge, this is the first paper that develops a supervised autoregressive LSTM model using the previous plant states to predict the future water flow of a hydropower system. Moreover, this paper offers a clear benchmarking between MILP optimization and recursive neural network prediction, which is not seen in the current hydropower literature. Overall, this paper addresses a significant gap in the literature and lays the foundation for further research in the application of machine learning algorithms for the STHS problem.

2 Case study

The hydropower system presented in this paper is privately owned and operated by the Rio Tinto company, an aluminum producer in the Saguenay-Lac-Saint-Jean region. Their system, shown in Figure 1, produces around 90% of their energy needs. The energy market value is not considered because the price is fixed by their contracts with Hydro Quebec, a government-owned corporation that is in charge of the production, distribution and transmission of the energy in the Province. Therefore, the electricity price is not considered in this paper. The system consists of six hydropower plants, each with a reservoir.




Figure 1: Rio Tinto's power plants designated by "R" pins in the Saguenay Lac-Saint-Jean region [61].

This hydropower system is part of a watershed that covers a surface area of about $73,800 \text{ km}^2$. Since the Saguenay Lac-Saint-Jean region is subject to a lot of snowfall during the winter, this watershed is subject to floods between April and June. These conditions require intricate reservoir management within the region to ensure both year-round water management and the safety of its residents. High inflows in spring cause a lot of erosion, hence the importance. In summer, there are also many boaters on the lake and public beaches, so managing the level in the summer is important to allow activities on the lac, which led the region to develop a strong tourism-based economy.

This research studies two interconnected hydropower plant located on the Péribonka River. As shown in Figure 2, the Chute-du-Diable (CdD) plant is the first of Rio Tinto's system to receive upstream water flow.




Figure 2: Location of CdD and CaS hydropower plants, located on the Péribonka River [20].

2.1 Chute-du-Diable and Chute-Savane hydropower plants

The CdD plant and the CalS plant are separated by around 20 km of river length. The water flow output of CalS is then sent down the river and into the lake Saint-Jean. Table 1 presents the characteristics of the power plants used in the case study for the paper.

Table 1: Characteristics of the CdD and CalS powerplant [61][35].

Characteristics	CdD	CalS
Number of units	5	5
Dam height	37.8 <i>m</i>	39.62 <i>m</i>
Gross water head	33.1 <i>m</i>	34.5 <i>m</i>
Reservoir surface area	47 <i>km</i> ²	18.5 <i>km</i> ²
Holding capacity	1200 <i>hm</i> ³	625 <i>hm</i> ³
Installed capacity	224 <i>MWh</i>	245 <i>MW</i>
Maximum Water output	850 <i>m</i> ³ / <i>s</i>	810 <i>m</i> ³ / <i>s</i>

Both plants have 5 turbine units available for power production and have similar values in relation to their size and capacity. One key difference is the size of the reservoir, with CdD having almost twice the holding capacity of CalS. Together with CalS downstream of CdD, this means that the outflow of CdD has a significant influence on the water management of CalS. As the formulation of the STHS problem implies, the main objective for the scheduling of both power plants involves determining the quantity of water to discharge and spill every hour.

2.2 Data

Rio Tinto collects data at intervals of 2 minutes from their power plants. For this project, a dataset of hourly data is provided starting from December 2010 to December 2022. The content of the dataset relates to historical data from different aspects of each power plant. The characteristics registered for each period are detailed in the table 2.

Table 2: Characteristics contained in the 12 years dataset of the CdD-plant production.

Feature	Unit
Time	hour
Natural Inflow	<i>m</i> ³ / <i>s</i>
Elevation	<i>m</i>
Volume	<i>hm</i> ³
Water Discharged	<i>m</i> ³ / <i>s</i>
Water Spilled	<i>m</i> ³ / <i>s</i>
Energy	<i>MW</i>
unit State (x5)	[0,1]

For the sake of confidentiality in regards to Rio Tinto, the real data is classified. Instead, a conversion to the percentage is used in this paper by using Eq.(1). This equation is taken from [15] and shown here:

$$percentage(\%) = \frac{current\ value \times 100}{\max(historical\ values)}. \quad (1)$$

2.2.1 Input selection

This project is based on two datasets that comprise 12 years of hourly data describing the state of the CdD and CalS hydropower plants from December 2010 to December 2022. The raw data was used from the original dataset, except for the natural inflows and water discharge. The equipment used to register the natural inflows in a reservoir is prone to errors caused by natural phenomena

that agitate the water and generate waves, which can disturb the sensors in the area. These errors cause either an abnormal spike in the series or an inflow value so low that it becomes negative. These values are normalized to the average of the previous and next values of the affected periods. Another modification of the feature is the water discharge. Hydro-Québec, a crown corporation, negotiates fixed price contracts with Rio Tinto that require the company to feed them with the energy needs for a short period of time, usually 1 to 3 hours. This generated spikes in energy production has been normalized in order to stay consistent with the energy needs of Rio Tinto.

For the dataset, volume and elevation are related by an analytical equation, where the volume is inferred using the net water head value. Because both features are redundant, the volume is chosen to represent the amount of water available in the reservoir. This is also consistent with the optimization model used to compare the results from this paper's model.

The raw natural inflow cannot be used in its current state as an input for prediction, since the observed inflow at any given period is not computed immediately. To be usable, this series is shifted forward so that the inflow value represents the inflows from the previous period. The state of each unit is simplified into a single feature compared to a binary feature for each, representing the number of active units for each period.

As the machine learning model is compared to an MILP model that uses deterministic natural inflows known in advance, a similar feature must represent these values in some way during training. It is achieved by using a moving average to represent the average inflow of each future interval of 24 hours. The Moving Average (MA) equation is given by :

$$MA_{d,t} = \frac{1}{24} \sum_{i=1}^{24} \text{natural inflows}_{t+i,d} \quad \forall t \in T, \forall d \in D, \quad (2)$$

where t represents the hours in each day d added to the MA for the whole dataset of T hours and D days. Figure 3 shows the moving average when $d = 1$ for CdD.




Figure 3: Moving average of the first day ahead of the natural inflows at CdD from late 2010 to late 2017.

The use of this method helps to better represent the historical inflows by toning down the high fluctuations, which is also useful in the context of machine learning by keeping the natural inflow more consistent between instances.

2.2.2 Water outflow and power production

To make a decision on the scheduling of the power plant, more precisely the quantity of water to discharge, a new feature is to be created in the dataset, as shown in Eq. (3) that represents the total

water value output of the hydro plant at each period. The outflow of a hydropower plant is broken down into the discharge flow and water spilled. The discharge flow relates to the body of water used to produce energy, and the spilled water relates to the unused water that is released directly from the reservoir to the downstream. Figure 4 shows these features in the dataset.

$$water\ outflow_t(m^3/s) = discharge\ flow_t + water\ spilled_t \quad (3)$$




Figure 4: Percentage of the water outflow, including both discharged and spilled, at CdD from late 2010 to late 2017.

The water inflows from CalS are different from those of CdD because of the influence of the outflow from CdD that is added to CalS. Since the two plants are close together, the delay in water travel is ignored. Therefore, the total inflow from Eq. (4) replaces the feature “natural inflow” for the CalS dataset.

$$CalS(total\ inflow) = CalS(natural\ inflow) + CdD(water\ spilled + water\ discharged) \quad (4)$$

The two plants are now linked in the CalS dataset through this new feature, which helps to correlate their operations for future modelling.

2.2.3 STL decomposition of the inflows

Spillage does not often occur in the dataset. As shown in Figure 3, snowmelt occurs between April and July, where spikes in the inflows can be spotted at yearly intervals. Inflow values fluctuate a lot in this set, mainly due to the sensors used to register the inflow in a reservoir. These sensors are sensitive to disturbance when values are recorded, such as a gust of wind that causes the water level to rise or fall. Therefore, a pattern can be seen in the data caused by the recurring seasons. To help analyze the data, a decomposition method called Seasonal-Trend Decomposition using Locally Estimated Scatterplot Smoothing (LOESS) or STL Decomposition [10] is applied to the inflow time series to separate and smooth the series. STL Decomposition allows to decompose a time series into three subseries referred as the trend, seasonal and residual component. Using LOESS [11] to transform the series into multiple regressions in a local area of the data, the method separates the data into three configurable components. As shown in Eq. (5), this decomposition method does not lose its original information, because each component is additive and can be summed back to its original values:

$$natural\ inflow = trend + seasonal + residual \quad (5)$$

This manipulation helps to visualize the data by eliminating the noise without a loss of information. Figure 5 is the decomposition of the natural inflows of CdD, with the parameter of the decomposition algorithm set to monthly for the seasonal component and yearly for the trend component.




Figure 5: STL decomposition of the inflows features at CdD from late 2010 to late 2017.

The season component now clearly highlights the flood period while also highlighting an ascending trend in the yearly number of inflows. The mean value of the residual component represents about 2.03% of the data range, which is acceptable considering the ideal mean value of zero for this component. This means that there might still be patterns not captured by the method or that the series might be multi-seasonal.

3 Methodology

In this Section, the methods used to achieve a predictive model for the STHS problem on a two-hydropower plant system are described. In this paper, a modified version of recurring neural networks (RNN) is utilized, which allows for the prediction of water discharges for each hydropower plant in the system. The dependency between plants and the prediction method for the system is also addressed. Lastly, a MILP model developed for this system prior to this project is briefly explained, as it is used in the next Section of this paper for results analysis. A short comparison between RNN and MILP is also provided to address potential biases when comparing both methods. This Section highlights the key component of this paper that defines the effectiveness and applicability of the proposed approach.

3.1 Recurrent neural network

Machine learning is a term used to describe a type of algorithm that uses data to build a model capable of predicting an output based on given input data. In contrast to optimization models that search for an optimal solution, machine learning is used to predict the solution to a problem given as input data. These types of algorithms are increasingly used in various fields as the various methods evolve and data availability becomes more accessible than ever before. There exists a wide variety of machine learning methods and models, each of which can be used based on the problem's context. For this paper, a supervised learning method is used to predict the water discharge for each hydropower plant available in the hydropower system. Supervised learning means that, as the model is trained on historical data, the model is informed of the quality of its predictions based on the real values present in the dataset.

A neural network is a type of machine learning algorithm. Similarly to the human brain, neural networks consist of interconnected neurons that process data in layers, allowing the model to learn

complex patterns [24]. Among the various types of neural networks, RNNs stand out for their ability to handle sequential datasets. Unlike traditional neural networks, RNNs have connections that form directed cycles. This class of neural network generally has an added mechanism that allows one to retain information about the previous output of the model, allowing one to maintain a hidden state that captures information from previous time steps. Figure 6 shows a simple architecture of an RNN with a single recurrent layer.




Figure 6: Depiction of a neural network with a recurrent layer receiving information on the previous output of the model.

The general architecture of an RNN is typical of any neural network model, including an input layer, one or more hidden layers, and an output layer. The specificity of the RNN is that the hidden layers are augmented with recurrent connections, which ties the previous output of the model to the hidden layer. This added step makes this type of network well suited for problems with temporal dependencies, where understanding the sequence and order of data points is crucial [27]. Because of this feature, the RNN architecture fits the STHS problem, where the context of the previous hours significantly influences the current predicted output.

3.1.1 Autoregressive LSTM model

The added connections of the RNNs mean that they are susceptible to the vanishing and exploding gradient problem because the gradients either decrease to near zero or grow exponentially during backpropagation [67]. This problem often hinders the training process and prevents the network from training effectively [50]. In this paper, a variant of the RNN is used, namely the LSTM model. These types of models use a gating mechanism that makes them more effective in capturing long-term dependencies and mitigate the problem of vanishing / exploding gradients [29, 50]. The LSTM is implemented as a hidden layer that contains LSTM cells as in Figure 7.

An LSTM cell is composed of a forget gate, an input gate and an output gate. The forget gate receives information x_t from the previous layer, long-term memory c_{t-1} and short-term memory h_{t-1} . This gate determines what percentage of the long-term value should be remembered. Similarly, the input gate determines what new information should be kept in long-term memory c_t . The output gate regulates what information to output based on both short-term and long-term values. The output of this gate, which is the output of the cell, also becomes the new short-term memory h_t . The activation functions of tanh and sigmoid (depicted as σ) transform the values into specific ranges: $[0,1]$ for sigmoid and $[-1,1]$ for tanh. In an LSTM, the sigmoid function specifically acts as 'gates' that control the information flow., while tanh functions create and regulate the actual content being processed.

The output of an RNN model usually only gives information on the next predicted value based on the given input. In order to turn a single output prediction of the model into a forecast, the model




Figure 7: An LSTM cell depicting the forget, input and output gate.

uses its own prediction for the next period as input to predict the subsequent periods. This method is known as an autoregression. By iteratively feeding the previous prediction back into the model, called the feedback loop, the RNN can generate a sequence of predicted values over several hours. With the help of TensorFlow's documentation [59], the LSTM model is turned into an autoregressive LSTM model by adding a feedback loop to the base model. In Figure 8, the input of the model is shown to be a window of the last 48 hours available in the dataset before the desired prediction.




Figure 8: Representation of the process of the autoregressive model from the input data to the prediction sequence output.

In this figure, 24 periods are used as input to predict the next 24 periods. Each period in the input undergoes a warm-up phase to set up the internal state of the LSTM units using the inputs. This step is crucial to prepare the internal states of the model, as it ensures the ability to make accurate predictions in subsequent steps by leveraging relevant information from the input. After the warm-up, the model outputs a prediction along with the state, which is then used as input to predict the following time step. This process establishes a feedback loop in which each output becomes an input to the model until the specified prediction sequence length is reached. A key aspect of making the model autoregressive is that each prediction is the same size as a single instance. Although the model outputs a complete prediction for each feature, only the total water discharge values are compared with the corresponding label values to train the model. This methodology for the prediction of a single hydropower plant is illustrated in Figure 9.

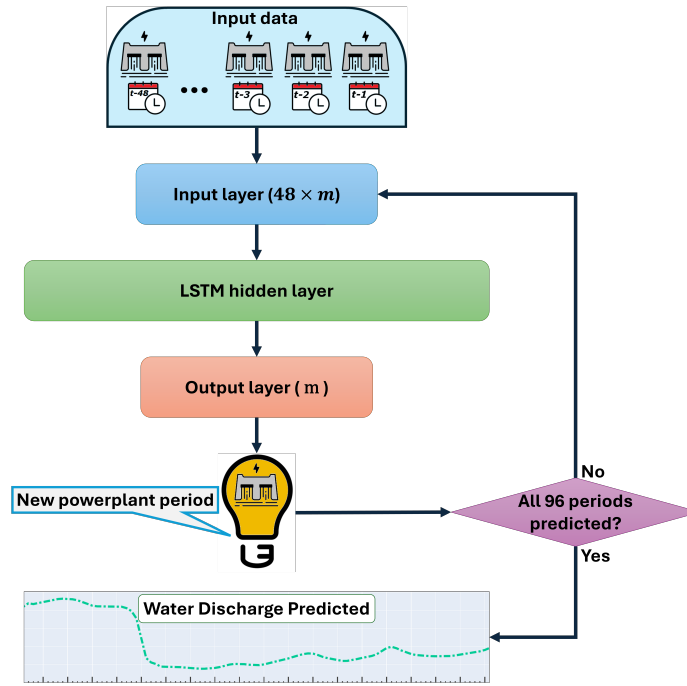


Figure 9: Architecture of the autoregressive LSTM model for the prediction of the total water discharge for a single hydropower plant.

The input is a sequence of the last 48 periods before the current prediction times m , the number of features at each period. Based on the different configurations tested, 48 periods as input gave the best results. Adding more increased computational load and time to train, with little to no significant increase in performance. The input is sent to the input layer to be processed in the feedback loop until all future periods are predicted. The number of periods to predict is set to 96 periods, which means that the water discharge for 96-hour periods will be predicted in the end. This number is chosen in order to have an even quantity of output as the MILP model, presented in Section 3.3. The size of the LSTM hidden layer was tested incrementally until it reached 256 cells, a further increase did not yield better results.

Many combination of hyperparameters was tested to configure the training step. A batch size of 32 gave good results while keeping the training time lower. The optimizer used is NADAM, as it is a slight upgrade to the more commonly used ADAM. The loss function is applied on the Mean Squared Error (MSE), which favour overall stability in the model. The learning rate is halved from default to 0.0005, because a slower learning rate resulted in a better end-score performance, at the cost of an increase in training time. The dataset is split 80% for model training and 20% for the test, the split occurring during the month of July 2020 in the dataset. The performance of the model is evaluated with the Mean Absolute Error (MAE), MSE and the Root Mean Squared Error (RMSE). These performance metrics are not the only deciding factor of the model, because a result could give better results than the label. For this reason, the quality of a model is also judged by evaluating the feasibility of the predictions and the amount of energy produced. The feasibility of the forecast is mainly based on the bounds of the water level in the reservoir, which is subject to constraints related to the accepted maximum and minimum level.

3.2 Training and system prediction

The objective of this project is to predict a forecast of the total water discharge for every power plant in the system. Before training both models, since there is a model trained for each power plant, the

dataset is scaled, which standardizes every feature by removing the mean and scaling to unit variance. For each feature, Eq. (6) is applied to standardize the data in the following manner :

$$z = \frac{(x - \bar{m})}{std}, \quad (6)$$

where z is the result, x the original value, $\bar{m}$ the mean of the feature and std represents the standard deviation. As mentioned above, the autoregressive LSTM model takes as input the last 48 periods in the historical data. The with the same base features in the dataset, detailed in Table 2, to describe each period in the dataset. Table 3 displays the features present for each period in the 48-hour sequence used as input in this paper model for CdD and CalS.

Table 3: Input features used in the autoregressive LSTM model for the Chute-du-Diable and Chute-à-la-Savane hydropower plant.

Chute-du-Diable	Chute-à-la-Savane	
Reservoir Volume	Reservoir Volume	Future Avg. Tot. Inflow (73-96)
Number of Active Units	Number of Active Units	Avg. Upstream Prediction (1-24)
Previous Natural Inflow	Previous Natural Inflow	Avg. Upstream Prediction (25-48)
Total Water Discharge	Previous Total Inflow	Avg. Upstream Prediction (49-72)
Future Avg. Inflow (1-24)	Total Water Discharge	Avg. Upstream Prediction (73-96)
Future Avg. Inflow (25-48)	Future Avg. Tot. Inflow (1-24)	
Future Avg. Inflow (49-72)	Future Avg. Tot. Inflow (25-48)	
Future Avg. Inflow (73-96)	Future Avg. Tot. Inflow (49-72)	

For each hourly period, the inputs used are the volume of the reservoir (hm^3) in the reservoir, the number of active units, the water inflows (m^3/s) from the previous period, as well as the total water discharge (m^3/s) and the average inflows (m^3/s) for the next four days. Concerning the latter, these averages are computed using the Eq. (3) at first and then replaced by the deterministic inflow scenario.

The CalS prediction model has five more features in the inputs, added to consider the dependency of the upstream hydropower plant operation. One of the features is the total inflow (m^3/s), which is the sum of the natural flow with the total CdD inflow for a given period. No time delay between the power plants is considered given their geographical proximity. The four other features are the average predicted upstream discharge (m^3/s) for the next four days, including the CdD outflow with the natural inflow. After training the model, this feature is computed with the predicted CdD water discharge. The trained CdD model is applied to the historical dataset of the plant to obtain this feature. Hypothetically, the same operation could also be performed on an additional hydropower plant in the system, allowing the model to be scaled to the desired size.

Once the models are trained, a pipeline is utilized to go from the input data to making a prediction for each plant. To better visualize, Figure 10 demonstrates the principal steps to achieve a prediction.

There are two arrows that represent the path for each plant. At first, both input are improved by computing the average future inflow. For CdD, this is the only modification made on the initial input values taken from the dataset. This input is sent to the LSTM model and the result is stored in memory. In regards to CalS, the stored result of CdD is needed to compute the average upstream prediction. The input is then ready to be sent to the trained model. From the validation instances used in this paper, the true inflow values are taken at each plant to calculate the hourly reservoir values over the predicted 96 hours.

3.3 MILP model

MILP models are widely used to formulate the STSH problem [4, 12, 42]. It consists of an objective function that is maximized or minimized, based on the problem formulation. In hydropower production the profits are usually maximized or the costs are minimized. A classical MILP model is composed of




Figure 10: Steps leading to a full prediction of the system's next 96 periods of water release.

parameters, known values that define the problem, variables, values to optimize the objective function, and constraints, which are restrictions on the problem, and bounds, which must be respected by the variables values in order to obtain a feasible solution [39]. The following formulation represents a generic form of a MILP minimization :

$$\min \quad c^T x \quad (7)$$

$$\text{subject to} \quad Ax \leq b, \quad (8)$$

$$x \in \Omega, \quad (9)$$

$$x_i \in \mathbb{Z} \quad \forall i \in I. \quad (10)$$

where Eq. (7) is the objective function, Eq. (8) represents a set of linear constraints, Eq. (9) defines the feasible domain Ω , which includes both continuous and integer variables, and Eq. (10) explicitly states that the decision variables in the set I must be integers.

3.3.1 Maximize power generation in STHS

To compare the results of this project's RNN, a MILP model developed prior to this project is used. The paper from [15] presents a MILP model for the STHS unit commitment problem for the CdD and CalS hydropower plants. The objective of the model is to maximize energy generation by optimizing the quantity of water committed to units of a plant during the next hourly periods. The MILP model introduces predetermined values, called efficiency points, to select the unit combination that maximizes the energy output. Efficiency points are optimal combinations of power produced and water discharges in a reservoir at full capacity. These are values at which a power plant usually operates, as deviating from these points will generate less efficient production performances. Figures 11 show the water discharge values corresponding to the efficiency points for each combination of three or more active units, because a lower number of units used is fairly uncommon.




Figure 11: Percentage of Water discharge for each efficiency points at CdD, grouped by unit combination.

In Figure 11, the relationship between the quantity of water released and the combination of units used shows the optimal values with a full capacity reservoir. Each combination has between 20 and 10 efficiency points, for a total of 195 points for CdD and 200 points for CalS. A novelty of this MILP model is that the variables are chosen to be efficiency points over a period of 96 hours. The following MILP model presented in this Section is the same model as presented in [15] without any modification added. The Eq. (11) represents the objective function in three terms, where the first computes the power output $P_{k,t}^b$ at each period for each combination at the maximum volume of the reservoir V_{max} , the second corrects for the difference of power produced at the current volume to account for energy losses θ^c when not at V_{max} . The last term penalizes unit start-ups. The sum of the three terms is multiplied by Δt to convert power to energy.

$$\max_{y,v,z} \Delta t \times \left[\sum_{c \in C} \sum_{t \in T} \sum_{b \in B} \sum_{k \in K_b^c} P_{k,t}^c \times y_{k,t}^c - \sum_{c \in C} \sum_{t \in T} \theta^c \times (V_{max}^c - v_t^c) - \sum_{c \in C} \sum_{t \in T} \sum_{j \in J_t} \epsilon^c \times z_{j,t}^c \right] \quad (11)$$

The MILP model is subject to constraints such as limited unit start-ups, water balance, and reservoir initial and final volumes. It incorporates energy losses from unit start-ups and uses next-period inflow data to provide optimal solutions for water commitment of each turbine unit. The model determines optimal values for water discharges, energy generation, reservoir volume, and unit commitment over the next 96 periods.

$$\begin{aligned} v_{t+1}^c &= v_t^c - \Delta t \\ &\times \left[\sum_{b \in B} \sum_{k \in K_b^c} (q_{t,k}^c \times y_{k,t}^c \times \beta) + (\delta_t^c \times \beta) - (d_t^c \times \beta) \right. \\ &\left. + \sum_{l \in U^c} \sum_{b \in B} \sum_{k \in K_b^l} (q_{t,k}^l \times y_{k,t}^l \times \beta) + (d_t^l \times \beta) \right], \quad \forall c \in C, \forall t \in T \end{aligned} \quad (12)$$

$$\sum_{b \in B} \sum_{k \in K_b^c} y_{k,t}^c = 1, \quad \forall c \in C, \forall t \in T \quad (13)$$

$$\sum_{b \in B} \sum_{k \in K_b^c} y_{k,t+1}^c \times A_{t+1,k,j}^c - \sum_{b \in B} \sum_{k \in K_b^c} y_{k,t}^c \times A_{t,k,j}^c \leq z_{j,t}^c, \quad \forall c \in C, \forall t \in T, \forall j \in J_t \quad (14)$$

$$\sum_{t \in T} \sum_{j \in J_t} z_{j,t}^c \leq N_{max}^c, \quad \forall c \in C \quad (15)$$

$$v_{min}^c \leq v_t^c \leq V_{max}^c, \quad \forall c \in C, \forall t \in T \quad (16)$$

$$v_1^c = v_{ini}^c, \quad \forall c \in C \quad (17)$$

$$v_T^c = v_{final}^c, \quad \forall c \in C \quad (18)$$

$$y_{k,t}^c, v_t^c, z_{j,t}^c \in \mathbb{B}, \quad \forall c \in C, \forall t \in T, \forall b \in B^c, \forall k \in K_b^c, \forall l \in U^c \quad (19)$$

$$d_t^c, d_t^l, v_t^c \in \mathbb{R}^+, \quad \forall c \in C, \forall t \in T, \forall l \in U^c \quad (20)$$

Constraints (12) ensure water balance in powerhouses, (13) enforce a single operating point per period per powerhouse, (14) link start-up variables to selected combinations, (15) impose a maximum number of unit start-ups, (19) defines binary variables and (20) the real variables. A constraint of interest is (16) that bounds the level of the reservoir water to prevent the model from over flowing, and the constraints (17) and (18) specify the initial and final volumes, which are set to the initial and final true value of the observed data. These last constraints are not present in the LSTM developed in this paper, and historical data shows that the volume bounds are often not respected in the historical dataset at any time of year.

3.3.2 MILP and machine learning comparison

In hydropower, the modelling of an MILP for the STHS problem is based on transforming the problem into mathematical functions, parameters and variables, using proper algorithms to find the optimal value, given the scope of the problem [39]. Compared to the MILP formulation, a machine learning model is built with historical values related to the problem. This method relies on the patterns and correlations in the data to make predictions [24]. Unlike MILP, which provides explicit constraints and objective functions for optimization, machine learning models learn from past data to generalize and predict future outcomes by trying to capture complex nonlinear relationships. Handling non-convex Mixed Integer Non-Linear Problems (MINLP) is significantly more difficult than MILP due to the existence of multiple local minima, which make them less desirable than their linear counterpart [37].

As these two fields of research evolve over time, researchers often question which technique should be used for a given project [53]. If the purpose of a program is to give the optimal solution on a given problem, an optimization model can be built by transforming the problem into objectives and constraints. This type of model requires the intervention of experts to improve performance and ensure that the model is always parameterized in a way that represents the context of the problem. The problem itself must also be scaled in a way that is computationally feasible, which can significantly restrict the problem formulation. Machine learning requires less interaction from experts, as it uses data on the problem to find a solution and can be kept up-to-date with the latest input. Although ML models may not guarantee optimality, they offer significant advantages in terms of computational speed and lower maintenance, appealing to problems with large datasets, real-time processing requirements, or the need for adaptive learning. In the context of the STHS problem, optimization models are needed for a theoretically optimal solution and the resulting strategy to be compliant with operational constraints. For machine learning models, they excel in scenarios where the system must adapt to changing conditions quickly and provide solutions rapidly, making them a good fit in the short-term and real-time decision-making related to STHS [63].

Due to their shared objective of finding solutions to given problems, machine learning and optimization methods are often compared, with results varying based on factors such as problem formulation and the techniques used. For example, [5] tested two new optimization models on a binary classification task against classic, well-established machine learning methods, finding that the optimization models favour both optimization models for their accuracy with precision and recall. It is important to remember that optimization is a key component in the training of a machine learning model, where the objective function typically measures the error or loss between the predicted and actual outputs, helping to find the best parameters to achieve this goal [55]. Hyperparameter tuning also relies on optimization techniques, such as grid search, random search, and Bayesian optimization, to improve model performance. As highlighted in [3], numerous papers explore hybrid models that leverage the strengths of both optimization and machine learning, with the aim of combining these methodologies for greater efficiency. Although this paper compares the two techniques, it is essential to recognize that each has its own strengths and weaknesses.

4 Results and discussion

In this Section, the performances of the autoregressive LSTM model are analyzed. First, the results of the training and testing step are presented. The validation instances provided in this project are then used to predict the sequences of water discharges for each hydropower plant. These predictions are shown in contrast to the results from an MILP model, taken from [15], and the real operational decision at each plant. The analysis of the water discharges is coupled with the evolution of power and volume resulting from the predictions for each plant. This is followed by an analysis of individual instances to better understand the LSTM results. This Section highlights the difficulties of representing the STHS problem as a data-driven model, the effect of making prediction under high uncertainty and the consequence of prediction inaccuracies on the evolution of external values.

4.1 Training set results

As mentioned in Section 3.1.1, the LSTM model is trained with the Nadam optimizer on the loss value, which is the MSE. A lower MSE value helps to monitor and minimize the chances of outliers in the predictions, where the average of the difference between the original and predicted values are squared, which results in higher values when the difference is significant. The MSE is given by Eq. (21) :

$$MSE = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y})^2, \quad (21)$$

where y_i and $\hat{y}$ represent the historical discharge value and the prediction of the models, respectively, and N represent the amount of data points. The MSE is used to ensure that the model's predictions remain within a realistic range. In this project, it is crucial to prioritize consistency and prevent highly infeasible results. This concern is due to a lack of the ability to constrain the model when making predictions of water discharge values. Ideally, predicted water discharge values would be based not only on pattern recognition in the dataset, but also on the constraints of the hydropower plant in an uncertain environment. This could be fixed through a post-prediction step, but would not be fit for a fair comparison with the MILP model. A post-prediction step for the predicted results could be explored in future work.

Because some input features, such as expected inflows, are modified to fit the context of each validation instance, it is not necessary to have a perfect accuracy score. Again, consistency in the results is much more preferable. To help analyze the results, the MAE and RMSE are also computed, see Eq. (22) and Eq. (23). Respectively, these two performance metrics provide insight into the model's average accuracy and sensitivity to strong errors.

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}|, \quad (22)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y})^2}. \quad (23)$$

As presented in Table 4, the dataset is divided 80% into the training set and the remaining 20% into the test set for each hydropower plant to build a model and analyze the performance on unseen data. For each hydropower plant, the training set is made up of 84,000 instances from December 2010 to July 2020, while the test set is made up of 20,968 instances from July 2020 to November 2022. Details of the features are shown in Table 3. The validation is performed on 10 instances from 5 months chosen within the test set covered period. For both hydropower plants, the training was carried out on dedicated servers from Calcul Québec [45], a regional partner of the Digital Research Alliance of Canada [40]. The cluster used consists of 125 central processing units and 111 V100-16G

graphics processing units. Each model took about 10 hours to train on 4 graphics processing units and 47000M memory made available.

Table 4: Summary of the dataset partitioning into training, test and validation sets.

Dataset	% of Total	# Instances	Time Period Covered
Training	80%	84,000	December 2010 – July 2020
Test	20%	20,938	July 2020 – November 2022
Validation	–	10	2020-09, 2020-12, ,2021-07, 2021-09, 2021-10
Total	100%	104,938	December 2010 – November 2022

Performance scores are relevant to assert the performance of the model in the training phase. The analyze performed on the set of validation instances uses modified inputs to swap the knowledge of true historical inflows in the dataset to the expected values included in the validation instances. The performance scores for the training of both hydropower plant are shown in Table 5. The values are standardized as per Eq. (6), meaning that the predicted water discharge is scaled where the mean is zero and the standard deviation at one. Therefore, the closer the scores are to zero, the better.

Table 5: Performance values of each model based on the test set from training.

	Chute-du-Diable	Chute-à-la-Savane
MSE	0.3332	0.1576
MAE	0.5772	0.397
RMSE	0.3843	0.2412

These values represent the overall difference between the real and predicted values. Considering both hydropower plants, CalS has much better performance values than CdD. This is expected, because CalS have added features related to the water output of CdD. The performance results are not representative of the model performance, because the water inflows for future days have not yet been modified to fit with expected inflows. The following Figures 12 display three water discharge sequences prediction from random periods in the test set for each hydropower plant, with the label representing the true discharge values in the dataset.




Figure 12: Six set of the water discharges prediction chosen at random in the test data for Chute-du-Diable (a) and Chute-à-la-Savane (b), with the input discharge values and the historical discharge has the label.

The first 48 values in these figures represent the previous value of the discharge of water from the inputs. After the 48-hour mark, the input plot split into two, where the full line represents the label, which is the actual historical water discharge decisions, and the predicted discharge is represented by circle dots. Water discharge values are kept scaled with the StandardScaler function from scikit-learn [43], to keep the data anonymized. Based on the given input data, the model is able to find patterns in the data and can predict a plausible sequence. The general trend of historical water flows is followed by the predictions with relative accuracy. Most of the error occurs when there is a sudden change in the quantity of water discharges. At some point in many predictions, historical values show a clear change in the quantity of water discharge from one hour to another. The prediction values change gradually for three to four hours before reaching a new stable level. This is problematic, because it means that without post-processing to correct the gradual change in discharge values, the prediction sequence cannot be accurately utilized, although adding an extra step after prediction would be trivial.

4.2 Validation instances analysis

Several validation instances were provided to experiment with the model in this paper on true historical instances. These validation instances contain the necessary information to compute the results with the MILP model from [15]. In order to have a meaningful comparison between the MILP model and the LSTM model in this paper, the input for machine learning is expected to mimic those of the MILP model's parameters detailed in Section 3.3.1. From the given instances, 10 validation instances were chosen based on the duration of each instance and the feasibility with respect to the MILP model. These instances cover the months of September, October, November and December, for both the year 2020 and 2021, which is outside the training set. For each hydropower plant, the validation instances contain the following data over a period of 14 days:

- The available turbine units;
- The natural inflows observed;
- The observed volume of the reservoir;
- The power produced by the plant;
- The simulated scenarios of natural inflows

Each of these features is tailored to the MILP model, as the validation instances were built for the purpose of the MILP in [15]. Because the data in the instances are daily values, a linear interpolation is performed to transform each day into hours. The volume of the reservoir is used to obtain the initial and final reservoir values. The simulated inflow contains plausible scenarios of daily inflow for an instance. According to the original article, only the median scenario is based on the total sum of water inflows. Further reading on the specifics of scenario trees modelling can be found at [56]. As the scenarios used in this paper are deterministic, the inflow used as input might differ from historical inflows. Future work on this project will consider stochastic inflows.

When preparing the input of the LSTM model, the future average inflow features are adjusted to fit the instance scenario. After the prediction, the reservoir volumes and the energy produced are computed for both hydropower plant. Because the power production of a hydropower plant is dependent on both the water discharge (q) and the water volume (v), the theoretical power (p) generated in kW by a turbine unit can be calculated with Eq. (24) :

$$p(q, h_n) = G \times \eta \times h_n(Q, v) \times q \times \rho, \quad (24)$$

where G is the gravitational acceleration constant of the Earth $9.8m/s^2$, η the efficiency of the units, ρ the density in kg/m^3 and p the power output in W . The net water head function h_n in m is dependent on the quantity of water discharged plus spilled Q in m^3/s and the volume of stored water in hm^3 :


$$h_n(Q, v) = h_f(v) - h_t(Q) - h_p(Q, q), \quad (25)$$

where h_f , h_t and h_p represent the elevation forebay, tailrace and friction loss in the penstock, respectively [57].


4.2.1 Comparison of all validation instances

This Section presents the analysis of the predictions of water discharges on the validation instances and a comparison with the results from a MILP model and historical values. Ten validation instances based on a real case of hydropower production are used. The instances range between the months of July and December 2020 and 2021. The year 2020 had significantly more natural inflows than the year 2021, which can also be inferred in the plots of this Section. The MILP model was developed in [15] to solve the STHS problem specifically for this time of year. The LSTM model is trained on every period of the year. Although not in the scope of this paper, the LSTM model could make predictions at other periods of the year.

Using the same process shown in Figure 10, each validation instance data is used to calculate the results presented in this section. This is also true for the MILP model presented in Section 3.3. The 48 previous hours are gathered from the dataset for each instance and the deterministic inflow calculated from the instances data is applied on both the LSTM and the MILP models. For confidentiality reasons, the results were transformed using Eq. (1). For both models, the same deterministic scenario is used. The results of the MILP model and historical decisions are used for comparison to illustrate and understand the performance of the LSTM. For example, results that perform similarly to reality or tend towards the optimized decision can be regarded as promising. Figure 13 shows a box plot where each instance includes three sequences of water discharges from the real decision values, the predicted values of the LSTM model and the MILP results from the model presented in Section 3.3.



(a) validation instances results of water discharge values (%) for the Chute-du-Diable hydropower plant.



(b) validation instances results of water discharge values (%) for the Chute-à-la-Savane hydropower plant.

Figure 13: A box plot comparison of each 96-hour instance showing real operational values (dark), predicted values (grey), and optimized values (light grey) in regard to water discharge variation. Instances are labelled by their month and year, with a number in parentheses to distinguish multiple instances within the same period.

In this figure, each component consists of five statistics: the minimum, first quartile, median, third quartile, and maximum. Some outlier values may appear in the form of dots under or above the component. For the purpose of this analysis, a box plot provides a comprehensive view of variation across all instances. Some component that represents the results of the MILP shows no variation. This appends simply because the chosen efficiency point stays consistent with no variation over the sequence. Compared to LSTM, the discharge results for 2020 are mostly lower than the MILP and the real decision values. For the year 2021, the median of the LSTM follow in most cases the historical data, with the MILP presenting often lower quantity of discharge. The fact that the LSTM fall mostly in between of the MILP and real discharge in 2021 may imply a sort of compromise between the two during low inflows. These results show that the prediction result stays mostly within reasonable values. It is a good thing when prediction results also show signs of stability. It is also expected that real decision is prone to a lot of variation, because reality is influenced by unforeseen factors, like sudden change in the electricity demand, and natural interferences, both not accounted in the MILP and LSTM model. The comparison of results between the CdD Figure 13a and CalS Figure 13b shows generally similar results, which is to be expected, since the two hydropower plant has similar features and are separated by only 20 km of river distance. With CalS being downstream and with less than half the reservoir area of CdD, this implies a strong influence on its production. One difference that can be seen is the number of outliers for CalS compared to that for CdD. This might indicate that CalS acts more as a “run-of-river” type of hydropower plant.

The energy generated for each instance plays an important role in the viability of the solutions. Even if the quantity of water discharged differs from one method to another, the resulting amount of energy produced over the horizon is most important for the operational profitability of a hydropower plant. Figure 14 represents the level of energy produced based on the discharge results and volume of each period.




Figure 14: A box plot comparison of each 96-hour instance showing real operational values (dark), predicted values (grey), and optimized values (light grey) in regard to the energy variation. Instances are labelled by their month and year, with a number in parentheses to distinguish multiple instances within the same period.

Again, there seems to be a discrepancy between the strategy used for the years 2020 and 2021. Looking at the MILP for 2021, the results are often similar to the real decision, but the MILP has a conservative approach to water release. This translates into lower amounts of energy produced when compared. Because optimization results must comply with a set of constraints, it lacks decision-making freedom and foresight in regard to the evolving context of production, where operators have more freedom to adjust production values in real time. For the year 2021, the situation is reversed, where the MILP results have higher energy levels than the real decisions, which were more conservative. For the prediction model, the same conclusions can be extracted as for Figure 13. The results stay mostly in between optimization and reality, reinforcing the idea of a compromise between the two methods.

The balance of reservoirs and the conservation of water is a key aspect of hydropower production. Based on Eq.(24) to generate power in a hydropower plant, a good model should maintain a high volume of water when inflows are expected to be low and low volume when inflows are expected to be high. In reality, the volume of the reservoir is subject to many stochastic factors such as expected inflows [13, 22, 23]. It is also important to comply with the bounds of the reservoir water level for the safety of the infrastructure and to avoid environmental hazards. Although the LSTM model lacks many of these contextual features, having the real inflows gives enough insight to compute the reservoir volume and infer some information when compared with real production data.

The volume of the reservoir gives information on the viability of the autoregressive LSTM model. For the LSTM model, respecting the reservoir water bounds is a complex task. Supervised machine learning models make prediction based on the historical pattern. To control the results of the predicted values, the training data must be engineered in a way that promotes a desired output. The model can also be parameterized to better fit the desired prediction behaviour. However, beyond these approaches, there is no mechanism in our LSTM model to constrain the prediction to respect the reservoir bounds. In comparison, MILP has bounds on water discharges and reservoir volumes (Eq. (16)) and a constraint in regard to the final volumes (Eq. (18)) on the last periods. Figure 15 illustrates a box plot of the results for reservoir volumes, calculated based on the initial historical volume, the results of water discharges and the observed inflows for each validation instance.

These two box plots present many results where the volume exceeds the maximum bound (Eq. (16)) of the reservoir for the LSTM, as represented by the horizontal dotted line at the 100% mark. This is especially true for CalS predictions (Figure 15b), where some instances go well above the bounds compared to CdD (Figure 15a). Looking at the real value for CalS, it is clear that whatever the instance date, the level barely changes around 96-97%, which is noticeably different from the values of CdD. This is caused by the difference in reservoirs size between the two hydropower plants, where CdD has a holding capacity of $1,200hm^3$, while CalS has almost half the reservoir capacity of $625hm^3$. Therefore, CalS production is heavily dependent on the upstream plant, making the margin of error much smaller in this reservoir. As a result, any significant alteration in the CdD water discharge strongly impacts the reservoir balance, requiring special attention and careful consideration at CalS. The added features related to the upstream reservoir here does not suffice to maintain water balance.

For the CdD plant, there are only three instances where the upper quartile of the components goes over the reservoir maximum bound constraint, all in instances of 2020. As stated in this Section, year 2020 had significantly more inflows than 2021. Compared to the CalS instances, those three are the same instances with the highest maximum value for this plant prediction. Therefore, these three instances can be considered to be more difficult to predict. Looking at the quantity of water discharged in Figure 13, the LSTM model predicts lower quantities of water discharges compared to the real and optimized decision. This implies that, in comparison, the prediction model did not have enough meaningful information to produce a feasible solution. For the rest of the CdD instances, the volume bounds are respected and are comparable to those of the real and MILP decisions. Instances 12-2020 (1) and 10-2021 (1) are a bit different from the rest. Instance 12-2020 (1) shows the values of volumes which are under the compared sequences. Looking at Figure 13a, the volume seems stable,

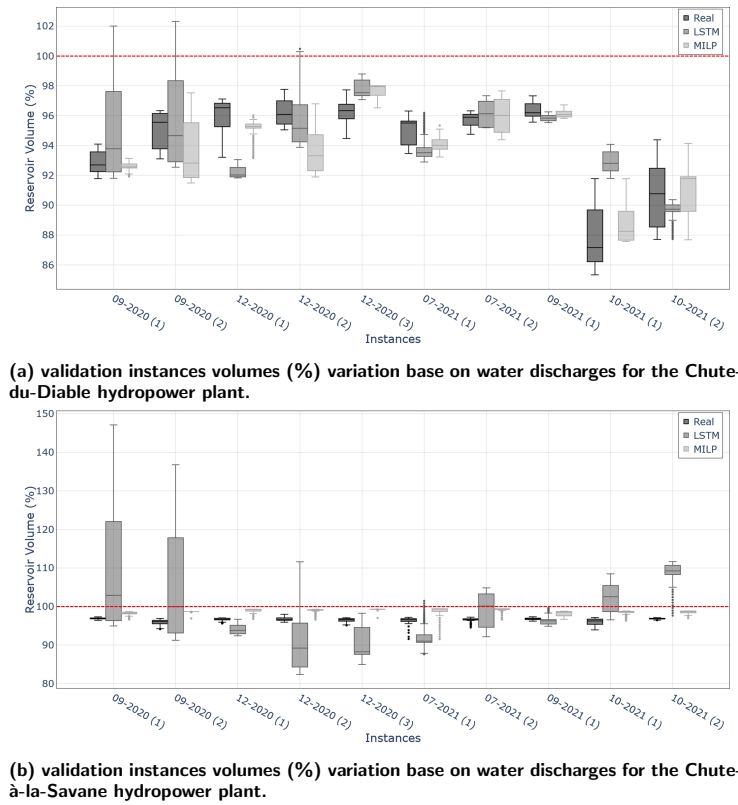


Figure 15: A box plot comparison of the reservoir water volume(%) for real operational values (dark), predicted values (grey), and optimized values (light grey). Instances are labelled by their month and year, with a number in parentheses to distinguish multiple instances within the same period.

compared to a higher discharge for real decisions and only the upper quartile component of the MILP model results going above the prediction box. With only a quartile higher, the volume results for the MILP model are still well above those of the LSTM model. In contrast, the predictions in instances 10-2021 (1) are above the maximum reservoir capacity for both hydropower plants.

The results demonstrate that the LSTM model can provide water discharge predictions that are often within a reasonable range of both the real operational data and the optimized MILP model. For 2020, the model tends to predict lower discharge values compared to reality and optimization, while for 2021, it finds a balance between the two. This suggests a capability to adjust to different inflow conditions, but also highlights the model difficulties to extract meaningful patterns in the given input when inflows are high. The LSTM model also maintains reservoir volumes of CdD close to the MILP in most cases, although there are some instances where it exceeds reservoir limits. In particular for the CalS plant, the prediction model has a lot more difficulties with regard to the water balance of the CalS hydropower plant. This discrepancy highlights the challenges posed by smaller reservoirs coupled to the influence of upstream operations for this type of model. More experiments are needed to better model the latter. The predicted power production aligns well with real and optimized values, strengthening the ability of the model to produce acceptable results in terms of energy generation. A major drawback of this method is the lack of explicit constraints in the LSTM model, which can lead to deviations, particularly in complex instances with high inflows. The LSTM model shows potential to approximate real operational decisions and offers a compromise between optimization and real-time flexibility. Improvements could be made to better respect reservoir limits and improve prediction accuracy in challenging conditions. In the next Section, a closer look into individual instances results is done to better comprehend the observation in this Section.

4.2.2 Comparison of individual instances

This Section focuses on specific validation instances that were highlighted in the previous Section. The goal is to observe the behaviour of the LSTM model prediction. Performance, stability and viability are the main appeal for the comparison between optimize and true decision. At first, this Section presents the predicted results of validation instances that were highlighted as feasible in the previous Section. Second, validation instances that resulted in unfeasible predictions are observed to understand and improve future models. In this Section, instances analysis is split again into CdD (a) and CalS (b), with two figures of the sequence of water discharges and volumes. Each figure shows three plots, the sequence of the real (filled), predicted (lines and dots) or MILP (dotted) decisions. All plots start 47 hours before the results of the model to present the prior decisions, which is also the production period used as input for the LSTM model.

Figure 16 shows the results sequence for the instance 12-2020 (1).




Figure 16: Close-up on the sequences of 96 hours of water discharge (left plot) and volume (right plot) of the 12-2020 (1) instance for real (dark), predicted (grey) and optimized (light grey) decisions.

The quantity of water discharged by the LSTM model is different. The CdD values are on a fix discharge and the CalS results values are on an upward trend. Compared to both real and MILP decisions, real decisions increase the quantity of water release for the third day ahead. MILP did the opposite by discharging less. In the end, both end their prediction with less water volume than the true decision by about 4% and 3% less water in the reservoir for CdD and CalS.

Another feasible solution is presented in Figure 17, with instance 09-2021.

For this instance, the predicted results follow the trend of both the real and predicted solution. Here, the result exposes a weakness in the way the LSTM model makes prediction sequences, because changes append progressively for water discharge values. Changes in the MILP model and reality always append suddenly, where the MILP follows the efficient points. However, the LSTM model could predict the upward trend with relative precision and with a similar final volume for CdD. Based on the observation made in the last Section in regard to volume stability, it is understandable that the final volume of CalS differs from both real and optimization, but up until the 80-hour mark, the prediction for this hydropower plant were still plausible. With a small adjustment to the quantity of water discharges, this could be fixed.




Figure 17: Close-up on the sequences of 96 hours of water discharge (left plot) and volume (right plot) of the 09-2021 (1) instance for real (dark), predicted (grey) and optimized (light grey) decisions.

Let us take a look at instances where the volume exceeded a reservoir capacity. One thing to note is that the three instances in which both hydropower plant results cause excess water in the reservoir are in the year 2020. It is known that 2020 received a lot more natural inflow than 2021, which gives a hint to understand the model behaviour. Table 6 presents the difference between the deterministic scenarios used in the input compared to the historical value that was replaced, for instances 09-2020 (1), 09-2020 (2) and 2020-10 (1).

Table 6: Comparison of the percentage between the mean deterministic scenario of natural inflows for each instance of CdD and CalS compared to the real historical inflow value received.

Plant	09-2020 (1)		09-2020 (2)		10-2021 (1)	
	Scenario	Historical	Scenario	Historical	Scenario	Historical
CdD	91.6	94.41	130.86	120.88	72.41	48.33
CalS	8.76	13.86	15.97	21.64	5	7.05

This Table shows that in regard to the CalS hydropower plant, the deterministic inflow scenarios are lower than the true historical quantity. For instance 10-2021 (1) of CdD, the difference is sizeable, with 24.08% more water expected than has been received. It should be noted that this data represent the natural inflow, but the CalS hydropower plant also receives additional inflows provided by the water output from the CdD hydropower plant. Figures 18 illustrate the results of instance 09-2020 (1).

In this instance, the MILP model decreased the quantity of water discharges for the first day and a half, which is the same result observed by the LSTM model. Up until the 40th hour, the LSTM acts similarly to the MILP and the reservoir volume stays within the expected amount for this first period. After that, the LSTM seems to have expected a lower quantity of inflows than in reality. Table 6 shows that the quantity of water expected in the deterministic scenario is always lower than the true values, with a wider gap for days 2 and 3. This resulted in a lower quantity of water discharges at this moment, causing the volume to rise above the 100% mark. Even if both hydropower plants have similar results, a smaller reservoir capacity for CalS shows that this irregularity has a much greater impact on the volume overtime.

Figure 19 displays a similar case in which the water discharge starts with good results, but does not keep up with the uncertainty of the inflow.




Figure 18: Close-up on the sequences of 96 hours of water discharge (left plot) and volume (right plot) of the 09-2020 (1) instance for real (dark), predicted (grey) and optimized (light grey) decisions.




Figure 19: Close-up on the sequences of 96 hours of water discharge (left plot) and volume (right plot) of the 09-2020 (2) instance for real (dark), predicted (grey) and optimized (light grey) decisions.

Here again, the first 40 hours of water discharge values seems inline with both the MILP and the real decision values. Similarly to Figure 18, the LSTM does not appear to recognize the increase in the quantity of natural flows received during the second half of the predicted values. This can be inferred from the sudden upward trend in the volume at the 30-hour mark. Even if Table 6 shows higher historical inflows than the deterministic scenario of the CalS hydropower plant and lower quantity of inflow scenario at CdD. This implies that there is another factor that causes a deviation in the CdD results.

The ninth instance differs from the last two instances in some aspects. Only the volume from the CalS reservoir exceeds its holding capacity, but the volume of CdD compared to the MILP and the real decision show that it is still higher than expected, but within the reservoir bounds. An in-depth

look at this instance is presented in Figure 20, with the results of CdD presented in Figure 20a and CalS presented in Figure 20b.




Figure 20: Close-up on the sequences of 96 hours of water discharge (left plot) and volume (right plot) of the 2020-10 (1) instance for real (dark), predicted (grey) and optimized (light grey) decisions.

For both hydropower plants, the evolution of the reservoir volume trend upwards at the beginning and goes downward after half of the sequence. In regard to CalS, it goes beyond the 100% threshold at first, but surprisingly ends at a final volume similar to the real decision values. This differs from CdD, with a final value almost 5% above the real decision values. The water discharge predicted is at first much lower than it should have been based on the MILP values for the first half of the decision sequence, but the prediction does seem to adjust on the second half, be it in a very noisy way. The inflows in Table 6 provide an explanation, where the inflows for CdD for days 3 and 4 differ greatly from the historical quantities. The difference between the first and second half of the CdD results sequence explains the difference between the LSTM and MILP results, which had a repercussion on the volume variation of the CalS reservoir.

4.3 Discussion

In this paper, an LSTM model is built and trained to predict the water discharge of the next 96 hours for two hydropower plants. The analyses made in this Section show that predictions are not always viable solutions for about half of the instances tested. This is mainly due to the reservoir volume when prediction and inflows are applied. Some instances demonstrate a higher amount than is allowed in regards to the hydropower plants water level bounds (Eq. (16)). The unfeasible results obtained open new challenges that must be addressed in future papers.

First, the difference in the size of the reservoir and the location of the hydropower plant makes for a very sensitive environment. The larger holding capacity for CdD upstream compared to the downstream CalS imply that the downstream plant must be operated together. Given that these two hydropower plants are close to each other, a large portion of the inflow comes from the CdD water output. In this paper, both hydropower plants share similar features, as presented in Table 1. This means that CalS is constrained to act like a run-to-the-river hydropower plant, despite its holding capacity. Therefore, a different approach should be applied in future research. The predictions made should be less independent to each other, as each hydropower plant has its own neural network. A merge of the output for each plant in a single multi-output neural network could help to obtain better

results about their dependency. However, the autoregressive LSTM model demonstrates its ability to capture the temporal patterns inherent in hydrological data. This is a crucial strength because water discharge in hydropower systems is a highly time-dependent process. The model ability to make prediction based on previous hours indicates that the underlying architecture is well suited for short-term hydropower data.

Second, the use of the true historical inflows during training to substitute for a deterministic scenario of the validation instances seems to have a positive effect when the expected natural inflow does not differ much from reality. Most unfeasible solutions originate in instances where inflows are highly uncertain. However, this is a well-documented and established challenge in STHS [22, 36, 56, 58]. Given the LSTM results, future work should implement natural inflow scenarios in the model that represent the stochastic nature of inflows. Adding a noise factor to the inflow feature might lead to more generalized predictions. Another prospect could be to generate new inflow scenarios based on historical values in the dataset, which would lean more on the medium-term horizon. Furthermore, recent studies show that current climate change should be considered, as historical data may not accurately represent natural inflows in the years to come, [34, 54, 64, 65].

Third, reducing the prediction range of the water discharge sequence would greatly improve the prediction accuracy. The choice of 96 hours predicted comes from the MILP model that [15] uses to compare the results. As seen in the previous Section 4.2.2, instances that are deemed unfeasible because of the volume evolution mostly append between the 40-hour and 50-hour points. Therefore, reducing the prediction length so that, say, 48 hours would improve the accuracy of the model. A lower number of periods to predict would simplify the complexity of the problem and allow faster model training. This would also allow for better scalability of the model with the addition of more features or hydropower plants. Overall, the LSTM model offers a strong research opportunity. The fact that the baseline performance is already reasonably aligned with expected water dynamics provides a solid foundation, proving that the performance of the model could improve further.

Fourth, the subject of this paper focuses on the STHS for the months of July to December, but the LSTM model is trained on a dataset with all-year-round production data. Due to the objective of comparing with a MILP model built to optimize the production of this period of the year, this aspect of the LSTM is not explored in this paper. The STHS problem in winter requires a different configuration of constraints and objectives, as it must deal with a low water resource due to snow accumulation, [54]. Furthermore, spring is a specific case characterized by a large quantity of natural inflows as a result of snowmelt. A random forest model was developed to forecast streamflow in snowmelt-dominated watersheds in [44] with great success. The fact that these seasons were used to train the LSTM model means that data that do not specifically match the validation instances that were used. Even if there is a temporal aspect to the recurrent nature of the LSTM, this means that the data outside the range of the validation instances periods may act as noise. The decision to keep these in the dataset came down to diminished size of the dataset, cutting the dataset in half. However, this also means that the LSTM model developed in this paper is trained and capable of making predictions at any time of the year, which makes it a versatile tool compared to the MILP model. These types of test are not conducted as this is outside the scope of this paper, but are to be explored in future works.

5 Conclusion

This paper demonstrates the efficacy of an autoregressive Long Short-Term Memory (LSTM) model to predict short-term hydropower production. By capturing the temporal patterns and sequential dependencies inherent in the STHS problem, the model shows promising performance in predicting water discharge over a 96-hour horizon. The integration of machine learning techniques, particularly LSTM and autoregression, into the prediction process marks a significant advance in the management of hydropower resources. The integration of autoregression with LSTM represents a significant advancement in applying machine learning to energy resource management. By effectively capturing the

temporal dependencies and patterns inherent in historical hydropower generation data, the proposed model provides a responsive tool for operational decision-making.

Comparative analysis with traditional computation methods highlights the capacity of the LSTM model to adapt to the complexities and nonlinearity related to the STSH problem. Although the model exhibits limitations with high natural inflows and uncertainty, these cases reveal opportunities for improvement. Irregular water discharge sequences in some scenarios, particularly those with high natural inflow conditions, lead to a deviation in regard to the reservoir capacity. Addressing this challenge in future work is expected to greatly improve the prediction capacity of the model. This research addresses a critical gap in the literature by providing a clear benchmark between machine learning approaches and conventional optimization techniques, such as Mixed-Integer Linear Programming (MILP). This work also brings forward valuable considerations for future development. These include shortening the prediction horizon to improve accuracy, modelling inflow uncertainty more effectively using stochastic methods or scenario generations, and adopting a multi-output neural network architecture to better represent dependencies between closely linked hydropower plants.

Furthermore, the findings underscore the importance of incorporating advanced predictive models in the context of increasing reliance on renewable energy sources. As the energy landscape evolves, the use of machine learning in the STHS will be essential to accurately schedule hydropower production to ensure grid stability and optimize energy management strategies.

Overall, this study sets the groundwork for future research in the application of machine learning algorithms in STHS and opens avenues for further exploration of hybrid models that combine the strengths of both machine learning and traditional optimization methods. The potential for improved efficiency and sustainability in hydropower operations is significant, paving the way for more resilient energy systems in the face of climate change and fluctuating energy demands.

References

- [1] Mehmet Bilgili Et Al. One-day ahead forecasting of energy production from run-of-river hydroelectric power plants with a deep learning approach. *SCIENTIA IRANICA*, pages 1838–1852, 2022.
- [2] Abdus Samad Azad, Md Shokor A. Rahaman, Junzo Watada, Pandian Vasant, and Jose Antonio Gamez Vintaned. Optimization of the hydropower energy generation using meta-heuristic approaches: A review. *Energy Reports*, 6:2230–2248, 2020.
- [3] Beatriz Flãmia Azevedo, Ana Maria A. C. Rocha, and Ana I. Pereira. Hybrid approaches to optimization and machine learning methods: a systematic literature review. *Machine Learning*, 113(7):4055–4097, Jul 2024.
- [4] Imène Bankalai and Sara Séguin. Hydropower optimization. *Les Cahiers du GERAD*, 711:2440, 2020.
- [5] P. Baumann, D.S. Hochbaum, and Y.T. Yang. A comparative study of the leading machine learning techniques and two new optimization algorithms. *European Journal of Operational Research*, 272(3):1041–1057, 2019.
- [6] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.
- [7] Jose Bernardes, Mateus Santos, Thiago Abreu, Lenio Prado, Dannilo Miranda, Ricardo Julio, Pedro Viana, Marcelo Fonseca, Edson Bortoni, and Guilherme Sousa Bastos. Hydropower operation optimization using machine learning: A systematic review. *AI*, 3(1):78–99, 2022.
- [8] Chiara Bordin, Hans Ivar Skjelbred, Jiehong Kong, and Zhirong Yang. Machine learning for hydropower scheduling: State of the art and future research directions. *Procedia Computer Science*, 176:1659–1668, 2020. Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 24th International Conference KES2020.
- [9] Statistic Canada. Hydroelectricity generation dries up amid low precipitation and record high temperatures: Electricity year in review 2023 - statistics canada, 2023.
- [10] Robert B Cleveland, William S Cleveland, Jean E McRae, Irma Terpenning, et al. Stl: A seasonal-trend decomposition. *J. Off. Stat.*, 6(1):3–73, 1990.
- [11] William S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368):829–836, 1979.

- [12] Hasan Huseyin Coban and Antans Sauhats. Optimization tool for small hydropower plant resource planning and development: A case study. *Journal of Advanced Research in Natural and Applied Sciences*, 8(3):391–428, 2022.
- [13] Magdalena Crisci and Rafael Terra. Valorization of irrigation water in a basin with large hydropower production through coupled hydrological and electric system modelling. *Water Resources Management*, 28(3):605–623, Feb 2014.
- [14] Pascal Côté and Richard Arsenault. Efficient implementation of sampling stochastic dynamic programming algorithm for multireservoir management in the hydropower sector. *Journal of Water Resources Planning and Management*, 145(4):05019005, 2019.
- [15] Maissa Daadaa, Sara Séguin, Kenjy Demeester, and Miguel F. Anjos. An optimization model to maximize energy generation in short-term hydropower unit commitment using efficiency points. *International Journal of Electrical Power & Energy Systems*, 125:106419, 2021.
- [16] Piyal Ekanayake, Lasini Wickramasinghe, J. M. Jeevani W. Jayasinghe, and Upaka Rathnayake. Regression-based prediction of power generation at samanawewa hydropower plant in sri lanka using machine learning. *Mathematical Problems in Engineering*, 2021:4913824, Jul 2021.
- [17] Danilo P. C. Filho, Erlon C. Finardi, and Antonio F. C. Aquino. Real-time dispatch for multi-unit hydroelectric plants with ac optimal power flow: The case of the santo antonio system. *IEEE Access*, 9:149322–149337, 2021.
- [18] Christodoulos A Floudas and Xiaoxia Lin. Mixed integer linear programming in process scheduling: Modeling, algorithms, and applications. *Annals of Operations Research*, 139:131–162, 2005.
- [19] Sudershan Gangrade, Dan Lu, Shih-Chieh Kao, and Scott L. Painter. Machine learning assisted reservoir operation model for long-term water management simulation. *JAWRA Journal of the American Water Resources Association*, 58(6):1592–1603, 2022.
- [20] Google. Hydropower plant Chute à la savane in Google Maps. <https://maps.app.goo.gl/GQHfUeRNthU8i8TUz6>, n.d. Spotted on April 08, 2025.
- [21] La Régie de l’énergie du Canada Gouvernement du Canada. Canada energy regulator / régie de l’énergie du canada, Jan 2024.
- [22] Ziming Guan, Ziad Shawwash, and Alaa Abdalla. Using sddp to develop water-value functions for a multireservoir system with international treaties. *Journal of Water Resources Planning and Management*, 144(2):05017021, 2018.
- [23] Ignacio Guisández, Juan Ignacio Pérez-Díaz, Wolfgang Nowak, and Jannik Haas. Should environmental constraints be considered in linear programming based water value calculators? *International Journal of Electrical Power & Energy Systems*, 117:105662, 2020.
- [24] Aurélien Géron. *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow*. O’Reilly Media, Inc., 2019.
- [25] Zhongzheng He, Chao Wang, Yongqiang Wang, Bowen Wei, Jianzhong Zhou, Hairong Zhang, and Hui Qin. Dynamic programming with successive approximation and relaxation strategy for long-term joint power generation scheduling of large-scale hydropower station group. *Energy*, 222:119960, 2021.
- [26] Arild Helseth, Birger Mo, Hans Olaf Hågenvik, and Linn E. Schäffer. Hydropower scheduling with state-dependent discharge constraints: An sddp approach. *Journal of Water Resources Planning and Management*, 148(11):04022061, 2022.
- [27] Hansika Hewamalage, Christoph Bergmeir, and Kasun Bandara. Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1):388–427, 2021.
- [28] Martin N. Hjelmeland, Jikai Zou, Arild Helseth, and Shabbir Ahmed. Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming. *IEEE Transactions on Sustainable Energy*, 10(1):481–490, 2019.
- [29] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [30] M. Jafari Aminabadi, S. Séguin, I. Fofana, S.-E. Fleten, and E. K. Aasgård. Short-term hydropower optimization in the day-ahead market using a nonlinear stochastic programming model. *Energy Systems*, Sep 2023.
- [31] Yueyang Ji and Hua Wei. An approximate dynamic programming method for unit-based small hydropower scheduling. *Frontiers in Energy Research*, 10, 2022.

- [32] Zhong kai Feng, Tao Luo, Wen jing Niu, Tao Yang, and Wen chuan Wang. A lstm-based approximate dynamic programming method for hydropower reservoir operation optimization. *Journal of Hydrology*, 625:130018, 2023.
- [33] Shengli Liao, Jie Liu, Benxi Liu, Chuntian Cheng, Lingan Zhou, and Huijun Wu. Multicore parallel dynamic programming algorithm for short-term hydro-unit load dispatching of huge hydropower stations serving multiple power grids. *Water Resources Management*, 34(1):359–376, Jan 2020.
- [34] Pavel Matrenin, Murodbek Safaraliev, Stepan Dmitriev, Sergey Kokin, Bahtiyor Eshchanov, and Anastasia Rusina. Adaptive ensemble models for medium-term forecasting of water inflow when planning electricity generation under climate change. *Energy Reports*, 8:439–447, 2022. 2021 The 8th International Conference on Power and Energy Systems Engineering.
- [35] Ministère de l’Environnement, de la Lutte contre les changements climatiques, de la Faune et des Parcs. Répertoire des barrages. <https://www.cehq.gouv.qc.ca/barrages/listebarrages.asp>, May 2018.
- [36] Florian Mitjana, Michel Denault, and Kenjy Demeester. Managing chance-constrained hydropower with reinforcement learning and backoffs. *Advances in Water Resources*, 169:104308, 2022.
- [37] Luca Moretti, Giampaolo Manzolini, and Emanuele Martelli. Milp and minlp models for the optimal scheduling of multi-energy systems accounting for delivery temperature of units, topology and non-isothermal mixing. *Applied Thermal Engineering*, 184:116161, 2021.
- [38] Prahlad Mundotiya, Parul Mathuria, and Harpal Tiwari. Mathematical Approach-Based Power System Analysis: A Review of Short-Term Hydro Scheduling, pages 85–98. Springer Nature Singapore, Singapore, 2022.
- [39] Jorge Nocedal and Stephen J Wright. Numerical optimization. Springer, 1999.
- [40] Digital Research Alliance of Canada. alliancecan.ca. <https://alliancecan.ca/en>, 2024. [Accessed 29-11-2024].
- [41] Furkan Özkan. Prediction of long-term streamflow by using adaptive neuro-fuzzy inference system ((anfis). Master’s thesis, Hasan Kalyoncu Üniversitesi, 2022.
- [42] Iram Parvez, Jianjian Shen, Mehran Khan, and Chuntian Cheng. Modeling and solution techniques used for hydro generation scheduling. *Water*, 11(7), 2019.
- [43] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [44] L. T. Pham, L. Luo, and A. Finley. Evaluation of random forests for short-term daily streamflow forecasting in rainfall- and snowmelt-driven watersheds. *Hydrology and Earth System Sciences*, 25(6):2997–3015, 2021.
- [45] Calcul Québec. calculquebec.ca. <https://www.calculquebec.ca/>, 2024. [Accessed 29-11-2024].
- [46] Murodbek Safaraliev, Natalya Kiryanova, Pavel Matrenin, Stepan Dmitriev, Sergey Kokin, and Firuz Kamalov. Medium-term forecasting of power generation by hydropower plants in isolated power systems under climate change. *Energy Reports*, 8:765–774, 2022. The 2022 International Conference on Energy Storage Technology and Power Systems.
- [47] Michelle Sapitang, Wanie M. Ridwan, Khairul Faizal Kushiari, Ali Najah Ahmed, and Ahmed El-Shafie. Machine learning application in reservoir water level forecasting for sustainable hydropower generation strategy. *Sustainability*, 12(15), 2020.
- [48] Marwah Sattar Hanoon, Ali Najah Ahmed, Arif Razzaq, Atheer Y. Oudah, Ahmed Alkhayyat, Yuk Feng Huang, Pavitra kumar, and Ahmed El-Shafie. Prediction of hydropower generation via machine learning algorithms at three gorges dam, china. *Ain Shams Engineering Journal*, 14(4):101919, 2023.
- [49] Linn Emelie Schäffer, Arild Helseth, and Magnus Korpås. A stochastic dynamic programming model for hydropower scheduling with state-dependent maximum discharge constraints. *Renewable Energy*, 194:571–581, 2022.
- [50] Alex Sherstinsky. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306, 2020.
- [51] Vineet Kumar Singh and S.K. Singal. Operation of hydro power plants-a review. *Renewable and Sustainable Energy Reviews*, 69:610–619, 2017.
- [52] Hans Ivar Skjelbred, Jiehong Kong, and Olav Bjarte Fosso. Dynamic incorporation of nonlinearity into milp formulation for short-term hydro scheduling. *International Journal of Electrical Power & Energy Systems*, 116:105530, 2020.

- [53] Heda Song, Isaac Triguero, and Ender Özcan. A review on the self and dual interactions between machine learning and optimisation. *Progress in Artificial Intelligence*, 8(2):143–165, Jun 2019.
- [54] Shan-e-hyder Soomro, Abdul Razzaque Soomro, Sahar Batool, Jiali Guo, Yinghai Li, Yanqin Bai, Caihong Hu, Muhammad Tayyab, Zhiqiang Zeng, Ao Li, Yao Zhen, Kang Rui, Aamir Hameed, and Yuanyang Wang. How does the climate change effect on hydropower potential, freshwater fisheries, and hydrological response of snow on water availability? *Applied Water Science*, 14(4):65, Mar 2024.
- [55] Shiliang Sun, Zehui Cao, Han Zhu, and Jing Zhao. A survey of optimization methods from a machine learning perspective. *IEEE Transactions on Cybernetics*, 50(8):3668–3681, 2020.
- [56] Sara Séguin, Charles Audet, and Pascal Côté. Scenario-tree modeling for stochastic short-term hydropower operations planning. *Journal of Water Resources Planning and Management*, 143(12):04017073, 2017.
- [57] Sara Séguin, Pascal Côté, and Charles Audet. Self-scheduling short-term unit commitment and loading problem. *IEEE Transactions on Power Systems*, 31(1):133–142, 2016.
- [58] Njogho Kenneth Tebong, Théophile Simo, Armand Nzeukou Takougang, and Patrick Herve Ntanguen. Stl-decomposition ensemble deep learning models for daily reservoir inflow forecast for hydroelectricity production. *Heliyon*, 9(6):e16456, June 2023.
- [59] TensorFlow. Time series forecasting, 2024. Accessed: 2024-05-30.
- [60] M. Thirunavukkarasu, Yashwant Sawle, and Himadri Lala. A comprehensive review on optimization of hybrid renewable energy systems using various optimization techniques. *Renewable and Sustainable Energy Reviews*, 176:113192, 2023.
- [61] Rio Tinto. Installations. <https://energie.riotinto.com/energie-electrique/installations/>, Mar 2021.
- [62] David Tsuanyo, Boris Amougou, Abdoul Aziz, Bernadette Nka Nnomo, Davide Fioriti, and Joseph Kenfack. Design models for small run-of-river hydropower plants: a review. *Sustainable Energy Research*, 10(1):3, feb 2023.
- [63] Yoan Villeneuve, Sara Séguin, and Abdellah Chehri. Ai-based scheduling models, optimization, and prediction for hydropower generation: Opportunities, issues, and future directions. *Energies*, 16(8), 2023.
- [64] Asphota Wasti, Patrick Ray, Sungwook Wi, Christine Folch, María Ubierna, and Pravin Karki. Climate change and the hydropower sector: A global review. *Wiley Interdisciplinary Reviews: Climate Change*, 13(2):e757, 2022.
- [65] Shuai Zhou, Yimin Wang, Hui Su, Jianxia Chang, Qiang Huang, and Ziyang Li. Dynamic quantitative assessment of multiple uncertainty sources in future hydropower generation prediction of cascade reservoirs with hydrological variations. *Energy*, 299:131447, 2024.
- [66] Mehdi Zolfaghari and Mohammad Reza Golabi. Modeling and predicting the electricity production in hydropower using conjunction of wavelet transform, long short-term memory and random forest models. *Renewable Energy*, 170:1367–1381, 2021.
- [67] Nicolas Zucchet and Antonio Orvieto. Recurrent neural networks: vanishing and exploding gradients are not the end of the story. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 139402–139443. Curran Associates, Inc., 2024.