Technical Session 2B (14:00 - 15:10)	**Heuristics and Machine Learning applied to Hydropower Planning**	**Chair: Jiehong Kong, SINTEF**	
	An Explainable Physical-Machine Learning Hybrid Approach for Reservoir Inflow Forecasting	Rahmathulla Parambath	NTNU
	Applied solution techniques for bi-level problems in hydropower equivalent estimation	Uli Max Rahmlow	KTH Royal Institute of Technology
	Predicting Short-Term Hydropower Production using a Long Short-Term Memory (LSTM) Model	Sara Séguin	Univ. Quebec

An Explainable Physical-Machine Learning Hybrid Approach for Reservoir Inflow Forecasting

Rahmathulla Madathi Parambath[a], **Jayaprakash Rajasekharan**[a], **Hossein Farahmand**[a]

[a] *Norwegian University of Science and Technology (NTNU), Trondheim, 7034, Norway*

1. Introduction

Accurate and efficient inflow forecasting is crucial for optimizing hydropower scheduling. A range of physical, conceptual, and empirical models have been developed to address this need [1]. These models support effective water resource management and help mitigate risks related to water availability and flow [2]. Inflow forecasting models can be broadly categorized into (1) physical and conceptual models, (2) empirical or data-driven models, and (3) hybrid models.

Physical models use equations to describe water system dynamics, while conceptual models simplify these dynamics with assumptions. Both are complex but provide accurate, interpretable results when calibrated. Empirical models, including statistical and machine learning approaches, forecast inflows using historical data and can capture complex patterns but lack physical interpretability. Post hoc methods can offer explanations, but these models don't have a true physical basis for their predictions.

2. Objective

Inflow forecasting has become more complex due to climate change and human activities [3], making simple empirical models inadequate and requiring frequent recalibration of physical models. Advanced machine learning models can handle these complexities but often lack explainability. There is a trade-off between accuracy, efficiency, and explainability in forecasting models, and for high-stakes applications, both accuracy and interpretability are essential. Therefore, it's important to develop models that balance capturing complex dynamics with being interpretable.

The primary goal of this study is to develop and validate a novel hybrid model for reservoir inflow forecasting that integrates the strengths of physical hydrological principles with advanced machine learning techniques. This model aims to overcome the inherent trade-offs between accuracy, efficiency, and explainability that characterise existing forecasting approaches. By leveraging the robust theoretical underpinnings of physical models and the dynamic pattern recognition capabilities of machine learning, our hybrid model seeks to provide highly accurate forecasts while maintaining the interpretability essential for strategic water resource management. This approach addresses the increasing complexity and uncertainty in inflow patterns driven by climate variability and anthropogenic factors, offering a scalable solution adaptable to diverse hydrological contexts. The ultimate objective is to enhance decision-making in hydropower scheduling and water management through improved forecast reliability and deeper insights into the forecasting process itself.

3. Methodology

In the proposed hybrid architecture, the physical component is represented by the HBV model [4] [5], while the machine learning component is implemented using Long Short-Term Memory (LSTM) layers. The HBV model, a widely used conceptual hydrological model in Scandinavia, is employed to simulate the runoff process. Among the various versions of the HBV model, this study utilizes the lumped HBV version. Since the hybrid model is developed in Python, a surrogate version of the HBV model is used to facilitate integration.

The HBV model consists of several submodules, each representing a different hydrological process, including the snow routine, soil moisture routine, response function, and routing routine. The model takes inputs such as precipitation, temperature, and evaporation, processes them through hydrological equations in each module, and outputs the runoff corresponding to the subsequent time step. Each submodule contains parameters that are estimated through model calibration.

This HBV-ML hybrid approach allows for the development of various hybrid architectures by selecting different modules for replacement and tuning the neural network components. Through detailed analysis, the optimal architecture is selected—one that effectively balances model performance and explainability, ensuring both accurate predictions and interpretability of the underlying processes. In the final proposed hybrid architecture, the snow module and routing routine of the HBV model are replaced with LSTM and CNN layers. The soil moisture module is kept as such as in the HBV model, which models the process inside the soil system physically and contributes to the physical explainability of the hybrid model.

In our review of published literature on hybrid inflow forecasting methods, most existing hybrid approaches either utilize simulated data from physical models to train empirical models or employ machine learning techniques to correct errors in physical model predictions. This highlights the novelty of our approach, which uniquely combines physical hydrological information with neural network layers within a unified model architecture. By structuring the model into distinct modules and training it as a single unit, we achieve integration of physical and data-driven methodologies.

An additional key feature of this approach is the flexibility to define the parameters of the physical module based on user input, rather than exclusively calibrating them from data. This allows the model to quickly and efficiently adapt to new systems with known physical parameters, offering a significant advantage when transferring the model to different hydrological contexts or regions.

The hybrid model is trained using daily inflow data collected from January 1982 to October 1994. The dataset is split into 70% for training, 10% for validation, and 20% for testing. Once trained, the hybrid model is evaluated, and its performance is compared against two benchmarks: simulations from the HBV Light model and predictions from a standalone LSTM model.

4. Results

All models were evaluated using NSE, MAE, and RMSE. The hybrid model significantly outperformed the physical HBV Light model and showed slightly better performance than the standalone LSTM model. While increasing complexity could further improve accuracy, the focus here was on balancing performance with interpretability, so simpler neural network layers were used. Compared to the pure LSTM model, the current hybrid model achieved similar or slightly better results, with the added benefit of physical interpretability, making it more practical for hydrological applications. Test predictions are plotted against observed data in the figure below, and evaluation metrics are summarized in the table. Ongoing work on hybrid model optimization may further enhance performance.

4.1 Figures and tables

The model performance comparison is summarized in Figure 1 and Table 1. Figure 1 illustrates the comparison of forecasted inflow by the physical, LSTM, and hybrid models on the test dataset, while Table 1 presents the NSE (Nash–Sutcliffe model efficiency coefficient), MAE (Mean absolute error), and RMSE (Root mean square error) metrics for each model on the test data.

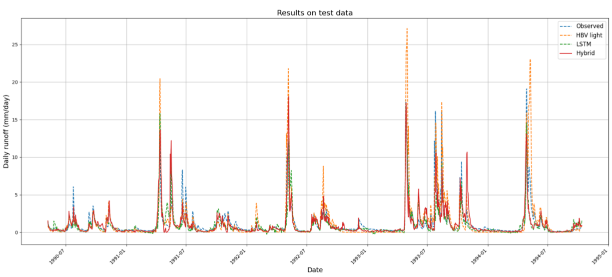

Fig. 1: Comparison of the hybrid model's predictions on the test dataset with observed values, the physical HBV Light model, and the LSTM model

Table 1: Comparison of model evaluation metrics on test dataset

Model	NSE	MAE	RMSE
HBV light	0.3191	0.7077	1.6494
LSTM	0.6187	0.5511	1.2342
Hybrid	**0.6254**	**0.5485**	**1.2233**

References

[1] Sarmad Dashti Latif and Ali Najah Ahmed. A review of deep learning and machine learning techniques for hydrological inflow forecasting. 25(11):12189–12216.

[2] Shi-Mei Choong and A. El-Shafie. State-of-the-art for modelling reservoir inflows and management optimization. 29(4):1267–1282.

[3] Shengli Liao, Zhanwei Liu, Benxi Liu, Chuntian Cheng, Xinfeng Jin, and Zhipeng Zhao. Multistep-ahead daily inflow forecasting using the ERA-interim reanalysis data set based on gradient-boosting regression trees. 24(5):2343–2363. Publisher: Copernicus GmbH.

[4] Development and application of a conceptual runoff model for scandinavian catchments | SMHI.

[5] Göran Lindström, Barbro Johansson, Magnus Persson, Marie Gardelin, and Sten Bergström. Development and test of the distributed HBV-96 hydrological model. 201(1):272–288.

Applied solution techniques for bi-level problems in hydropower equivalent estimation

Large energy system models are trying to simulate each component as detailed as possible to understand the system behaviour. Hydropower plays an important role in the energy system, as it is able to supply short and long-term flexibility. However, modelling hydropower can be challenging, as power stations and reservoirs in big river systems are often arranged in cascades, and relevant data may not always be publicly accessible. To address this, some models use a so-called "hydropower equivalent," which simplifies and aggregates hydropower within a specific region.

There are different approaches to calculating an equivalent. One method relies on the geographical location[1] of hydropower stations and reservoirs, while another methods using mathematical optimization to determine the best parameter settings through bi-level optimization[2]. In this case, the upper level of this problem aims to minimize the difference in power production between a given hydropower schedule and the equivalent schedule which is calculated in the lower problem. Particle swarm optimization (PSO) has been shown to be effective in solving the bi-level problem. The idea behind particle swarm optimization is to explore solution space using multiple particles, which then move toward the best-found solution in every iteration. Additional measures are implemented to prevent the algorithm from moving only into a local minimum, helping it converge toward a global optimum. A disadvantage is that a huge number of runs can be required to receive a sufficient good results. Additionally in earlier research, it has been shown that adjusting the first parameter set from the PSO can improve high price performances, which will increase even more the simulation time [3,4].

In this paper, a hyperparameter optimization algorithm is applied to determine the optimal parameters for a hydropower equivalent. Existing software frameworks, such as Optuna [5], can be used for hyperparameter optimization. This framework includes various sampling algorithms, all of which were tested for their applicability to this problem. In total, two different algorithms (Tree-structured Parzen Estimator algorithm, Gaussian process-based Bayesian optimization) are analysed in greater depth, as they appeared to demonstrate the best performance and applicability. Additionally to using Optuna also a combined method is used were the idea is to only use the PSO within a more limited space with already given or pre-estimated solutions from the hyperparameter optimisation.

Optuna's algorithms have demonstrated a significant reduction in computation time compared to the PSO algorithm. Additionally, Optuna provides a feature to assess the importance of variables, helping to understand the impact of newly introduced parameters. The results further show that simple models can achieve the same accuracy as the PSO algorithm. However, for more complex equivalent systems, the PSO algorithm offers better accuracy. In these cases, a combined method indicates to be the optimal choice.

In conclusion, the PSO algorithm still demonstrates the highest accuracy in more complex equivalent systems. However, Optuna and its hyperparameter optimization algorithms offer a practical alternative for computing hydropower equivalents, especially for simpler systems.

This approach also comes with the advantage of lower simulation time. Furthermore, Optuna can identify the most important parameters, leading to a better understanding of the equivalent system while also evaluating newly introduced parameters. For example, in a multiple station equivalent, it can identify which station parameter are more sensitive to the solution, providing new insights for improving equivalent system.

References

[1] P. Härtel and M. Korpås, "Aggregation Methods for Modelling Hydropower and Its Implications for a Highly Decarbonised Energy System in Europe," Energies, vol. 10, no. 11, Art. no. 11, Nov. 2017, doi: 10.3390/en10111841.

[2] E. Blom and L. Soder, "Computation of Multi-Scenario Hydropower Equivalents Using Particle Swarm Optimization," in 2020 IEEE International Conference on Environment and Electrical Engineering and 2020 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe), Madrid, Spain: IEEE, Jun. 2020, pp. 1–6. doi: 10.1109/EEEIC/ICPSEurope49358.2020.9160649.

[3]E. Blom and L. Söder, "Comparison of Different Computational Methods and Formulations for Hydropower Equivalents," 2022 IEEE 7th International Energy Conference (ENERGYCON), Riga, Latvia, 2022, pp. 1-6, doi: 10.1109/ENERGYCON53164.2022.9830320.

[4] U. M. Rahmlow and M. Amelin, "Comparison of Different Hydropower Equivalent Formulations to Improve High and Low Price Performance," 2024 20th International Conference on the European Energy Market (EEM), Istanbul, Turkiye, 2024, pp. 1-6, doi: 10.1109/EEM60825.2024.10608912.

[5]T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," in Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, in KDD '19. New York, NY, USA: Association for Computing Machinery, Jul. 2019, pp. 2623–2631. doi: 10.1145/3292500.3330701.

Predicting Short-Term Hydropower Production using a Long Short-Term Memory (LSTM) Model

Y. Villeneuve,[1,2] S. Séguin,[1,2] A. Chehri,[3] K. Demeester [4]

[1] Department of Computer Science and Mathematics, Université du Québec à Chicoutimi, Saguenay (QC), Canada
[2] Group for Research in Decision Analysis (GERAD), Montréal (QC), Canada
[3] Department of Mathematics and Computer Science, Royal Military College of Canada, Kingston (ON), Canada
[4] Rio Tinto, Saguenay (QC), Canada

Introduction

Hydropower is a clean and renewable energy that allows to produce energy with no air emission polluants. In the province of Québec, Canada, 99% of the energy produced is from hydropower, making the province a world-wide leader in hydropower production. The market is regulated by a government-owned corporation whom regulates the market and the independent producers. Producers such as Rio Tinto own their hydropower production system for their private benefit and need to manage wisely their systems in order to produce the most out of the available water, which accounts for 90 % of their total energy needs. To do so, many levels of optimization models are used. On an operational basis, mid-term and short-term models allow to manage the reservoir levels and the actual production at each of the hydropower plants. Producers have access to large amounts of historical data on the operations of their plants. Traditionally, mathematical optimization [2] is used to manage efficiently hydropower plants, but given the large amounts of data hydropower producers collect, this project aims at predicting the operations of the power plants based on these data sets.

Objectives

Machine learning has gained attention in recent years in the field of hydropower optimization, especially for the mid- and long-term models. Few papers address the short-term problem and the main goal of this project is to predict future hourly water discharges for a system composed of two powerplants. Results are then compared with a traditionnal Mixed-Integer Linear Problem (MILP) used to solve the short-term hydropower optimization model based on the efficiency points of the hydropower production functions [1]. One of the main motivations for this work is that the authors are interested to determine if predicting the operations of the power plants based on past operations could be viable or realistic, but also to assess the limitations of such an approach for the operations of a hydropower plant. The LSTM model is chosen since it is effective at capturing long-term dependencies.

Methodology

Two data sets composed of twelve years of data from Rio Tinto are available for this project. The hydropower system is composed of two powerplants with five turbines each. The data sets contains the following characteristics: natural inflows, volume, water discharged, water spilled, energy produced and turbine states. In order to predict future hourly water discharges the following sub-objectives are defined:

1. **Prepare the input data** The input data is prepared for usage in the LSTM model. Data is available every 2 minutes and must be transformed into hourly data. Also, since the goal of the model is to predict future water discharges, inflows must be shifted in the future, since the inflows are unknown at the time of making a prediction. As the powerplants are cascaded, water balance constraints need to be accounted for and this is done by creating a new feature in the data set that contains water inflows, but also water discharged and spilled from the upstream plant, for the downstream plant.

2. **Develop a framework to formulate the short-term hydropower problem as a LSTM model.** As the decision from the upstream plant affects the decision for the downstream plant, a framework is developed to account for the interconnection of the plants. The prediction from the upstream plant is made, then this decision is given as an input to the second LSTM model for the downstream plant. The LSTM is an autoregressive Recurrent Neural Network (RNN) that requires training to adjust different parameters representing the trade-off between short-term and long-term memory. In order to predict a forecast, a rolling-horizon framework is used to predict a series of future hourly water discharges.

3. **Compare results with a MILP.** The MILP in [1] is used to compare results with the prediction from the LSTM framework. As the short-term hydropower problem is formulated using the efficiency points of the hydropower production functions and that model is validated, comparing the actual operations of the plants with the two different approaches will allow to assess the limitations of the predictive approach.

4. **Compare results with the actual historical operations of the hydropower plants.** The results from the MILP were compared with the historical decisions, therefore, comparing the results of the prediction with the historical decisions will reinforce the analysis.

Results

The twelve year data set is split in an 80% training set and a 20% validation set. Hence, 84,000 hourly instances are used for the training set and 20,968 hourly instances for validation. Preliminary results require around 10 hours to train the model.

Preliminary results of the water discharge, obtained from the validation set are shown in Fig.1.

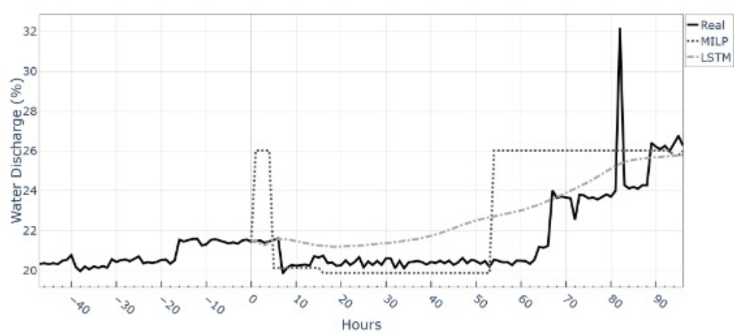

Figure 1: Water discharges for upstream power plant

Results show that the MILP solution, the LSTM prediction and the actual historical decisions are quite similar. These results are encouraging since they show that the prediction is not too far off from the historical decisions and the optimization solution. In this case, the final volumes are also quite close, but in practice, it is impossible to impose a final volume for the LSTM model. Therefore, further results will allow to assess the impact of such a limitation on the predictive model.

References

[1] M. Daadaa, S. Séguin, K. Demeester, and M. 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.

[2] S. Séguin, S.-E. Fleten, P. Côté, A. Pichler, and C. Audet. Stochastic short-term hydropower planning with inflow scenario trees. *European Journal of Operational Research*, 259(3):1156–1168, 2017.