Les Cahiers du GERAD ISSN: 0711–2440

Hydropower unit digital twin calibration using monitoring data and blackbox optimization

M. Fahed, A. Favrel, M. Gagnon, S. Alarie, M. Kokkolaras

G–2025–55

August 2025

Citation suggérée : M. Fahed, A. Favrel, M. Gagnon, S. Alarie, M. Kokkolaras (Août 2025). Hydropower unit digital twin calibration using monitoring data and blackbox optimization, Rapport technique, Les Cahiers du GERAD G– 2025–55, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2025-55) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Suggested citation: M. Fahed, A. Favrel, M. Gagnon, S. Alarie, M. Kokkolaras (August 2025). Hydropower unit digital twin calibration using monitoring data and blackbox optimization, Technical report, Les Cahiers du GERAD G–2025–55, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2025-55) 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
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Hydropower unit digital twin calibration using monitoring data and blackbox optimization

Melad Fahed [a, b]

Arthur Favrel [c]

Martin Gagnon [c]

Stéphane Alarie [b, c]

Michael Kokkolaras [a, b]

[a] *Department of Mechanical Engineering, McGill University, Montréal (Qc), Canada, H3A 0C3*

[b] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[c] *Hydro-Québec Research Institute (IREQ), Hydro-Québec, Varennes (Qc), Canada, J3X 1S1*

melad.fahed@mail.mcgill.ca
alarie.stephane@ireq.ca
michael.kokkolaras@mcgill.ca

August 2025
Les Cahiers du GERAD
G–2025–55

Abstract : One-dimensional models can enable the assessment of the dynamic behavior of hydropower units during transient operation with minimal computational resources. Since these models do not consider the full three-dimensional flow in hydraulic piping systems, their predictions rely on static performance characteristics obtained typically through reduced-scale measurements. These measurements usually cover only a small portion of the complete operating range of the machine, making it challenging to simulate transient events such as start-up, shutdown, or load rejections with high fidelity. To address this issue, we propose a novel approach for calibrating one-dimensional physics-based models of hydropower units by combining monitoring data with blackbox optimization. Specifically, the performance characteristic curves feeding the one-dimensional model of the power plant are represented by a polynomial whose parameters are optimized by minimizing the difference between simulation results and experimental data. We demonstrate that the proposed method is suitable for a 50 MW Kaplan turbine considering various startup scenarios, including different guide vane and blade opening sequences. The optimization was conducted using different combinations of training and test datasets to assess the validity of the calibrated models. Good agreement between simulation and experimental data was obtained using only a few startup sequences in the training dataset which demonstrate the robustness of the proposed methodology. This may pave the way for the calibration of hydropower unit digital twins for unit monitoring and anomaly detection.

Acknowledgements: This work will be published in the IOP Conference Series: Earth and Environmental Science following the presentation at the 10th IAHR Meeting of the WorkGroup on Cavitation and Dynamic Problems in Hydraulic Machinery and Systems.

1 Introduction

Digital twin (DT) technology has rapidly gained attention in the energy sector as a means to monitor, simulate, and optimize the performance of physical assets in real time. A DT is a virtual replica of a physical system continuously updated with field data, enabling anomaly detection, predictive maintenance and scenario analysis [12, 4]. In the context of hydropower, DTs promise enhanced energy production efficiency and reliability of turbine units by providing greater visibility into system dynamics under various operating conditions. However, achieving a truly predictive hydropower DT is challenging because the underlying simulation models are often imperfect. Meeting the high-fidelity demands of real-time DT simulations often requires personalized calibration, as traditional hydro turbine models alone may not capture the nuanced behavior of individual units across all operating conditions [12]. In particular, calibration of the model against monitoring data is essential to compensate for modeling errors and to ensure that the digital model accurately reflects the behavior of the physical turbine over its full operating range. Automated calibration strategies have thus become a focal point in the development of a DT [1].

Hydraulic turbines such as Kaplan units pose specific modeling challenges for DT development. Kaplan turbines are double-regulated machines with two independent control parameters: the wicket gate or guide vane opening (GVO) often denoted as α, and the runner blade pitch angle (often denoted β). Turbine manufacturers typically characterize steady-state performance via hill charts or characteristic surfaces that express the discharge coefficient (Q_{11}) and torque coefficient (T_{11}) of the turbine as functions of the speed coefficient (N_{11}), gate opening, and blade angle [8]. These $Q_{11}(N_{11}, \alpha, \beta)$ and $T_{11}(N_{11}, \alpha, \beta)$ surfaces are obtained from reduced-scale model testing under controlled conditions and then scaled up to the full prototype [8, 2, 7]. While such characteristic data provide a crucial foundation for physics-based 1D hydropower unit models [7], they have well-known limitations. Scale effects and site-specific factors (e.g. water temperature, wear, inflow conditions) mean that the raw model test curves do not always predict prototype performance with high accuracy [2]. Moreover, standard hill charts mainly cover operating regimes around the nominal conditions and do not capture the transient behavior during rapid events like start-ups, shutdowns and load rejections. In practice, full-scale Kaplan units experience dynamic regimes that were not fully examined in the laboratory [5]. For example, modern grid operation demands increasingly frequent start-stop cycles of hydro units to balance variable supply and demand, exposing turbines to off-design and highly transient sequences much more often than in the past [9]. These transients involve rapidly changing guide vane commands and unsteady flow, which can induce strong pressure fluctuations, vibrations, and component fatigue [5]. However, because transient operation is difficult to reproduce at scale and is not part of standard model testing protocols, the initial digital models often struggle to predict these behaviors. This gap underlines the need for calibrating the 1D physics-based model of the turbine using field measurements so that it can reliably simulate both steady-state and transient behaviors [13, 10].

To address this issue, we propose in this paper a novel approach for calibrating one-dimensional physics-based models of hydropower units by combining monitoring data with blackbox optimization. Specifically, the performance characteristic curves feeding a one-dimensional model of the power plant are represented by regression models whose parameters are optimized by minimizing the difference between simulation results and experimental data.

2 Methodology

Calibration of a model entails adjusting uncertain model parameters to minimize the discrepancy between simulated outputs and sensor data from the real asset. In a one-dimensional Kaplan turbine model, such parameters might include loss coefficients, inertias, empirical correction factors for the hill-chart surfaces, or other quantities that influence predicted flow, torque, and speed trajectories during transients. Manually tuning these parameters is labor-intensive and may not explore the complex, high-dimensional search space effectively. Instead, an automated calibration approach can be formulated as an optimization problem: find the set of model parameters that yields the best agreement between simulation and reality, typically by minimizing an objective function such as the root-mean-square error (RMSE) between measured and simulated signals. Because the physics-based 1D model is essentially a blackbox function of the calibration parameters, providing outputs (e.g. power, speed, pressure) without analytic gradients, it is natural to apply blackbox optimization (BBO) methods for this task. Blackbox (derivative-free) optimization algorithms are designed to handle expensive simulations and can navigate rugged error landscapes where gradient-based methods fail or are inapplicable. In particular, the Mesh Adaptive Direct Search (Mesh Adaptive Direct Search (MADS)) family of algorithms has proven effective for calibrating engineering models and other parameter-tuning problems across various domains [1]. In the proposed methodology, we employ the MADS-based Nonlinear Optimization by Mesh Adaptive Direct Search (NOMAD) optimization software version 3.9.1 [3, 6], to calibrate the 1D physics-based model of the Kaplan turbine automatically. The BBO has been successfully used to improve model fidelity in fields ranging from hydrological forecasting to nuclear plant simulation [12, 1], highlighting its suitability for our application.

This study aims to calibrate a one-dimensional physics-based 1D model of a Kaplan turbine to improve the accuracy of startup transient simulations. The calibration targets the torque response of the turbine, which is adequate to characterize the startup sequence. This relationship can be expressed as:

$$T_{11} = f(N_{11}, \alpha, \beta) \tag{1}$$

where T_{11} is the torque coefficient, expressed as a function of speed coefficient (N_{11}), GVO (α), and blade angle (β). The 1D physics-based model is implemented in SIMSEN, a simulation environment developed at the École Polytechnique Fédérale de Lausanne (EPFL) for modeling electromechanical transients in power systems [11].

The torque response function form is approximated using a third-order multivariate polynomial regression model. The initial model is trained on reduced-scale test data provided by the manufacturer, and a systematic model selection process is applied to identify the optimal polynomial structure. Candidate models are evaluated based on predictive error on a held-out test set, and statistically insignificant terms are removed using p-value screening. The final model, shown in Eq. 2, includes 17 symbolic coefficients ($p_1, p_2, \ldots, p_{17}$), which serve as calibration parameters.

$$\begin{aligned} T_{11}(N_{11}, \alpha, \beta) = {} & p_1 + p_2\beta + p_3\alpha + p_4N_{11} + p_5\beta^2 + p_6\beta\alpha + p_7\beta N_{11} \\ & + p_8\alpha N_{11} + p_9N_{11}^2 + p_{10}\beta^3 + p_{11}\beta^2\alpha + p_{12}\beta\alpha^2 \\ & + p_{13}\beta^2N_{11} + p_{14}\beta\alpha N_{11} + p_{15}\alpha^2N_{11} \\ & + p_{16}\beta N_{11}^2 + p_{17}\alpha N_{11}^2 \end{aligned} \tag{2}$$

During the optimization, the 17 polynomial coefficients are treated as decision variables. For each candidate parameter set, the torque surface is reconstructed and used to update the SIMSEN characteristic file. The reconstructed turbine characteristic surfaces are then compared to the original reduced-scale test data, and the discrepancy is quantified by the fit RMSE output (e_{fit}). The initial parameter vector (X_0), taken from a previously identified polynomial model, is used as the initial guess in NOMAD. SIMSEN simulates a set of predefined startup scenarios, and the resulting turbine speed profiles are compared to experimental prototype measurements. The difference is quantified by the simulation RMSE output (e_{sim}). The objective function minimized by NOMAD is the weighted sum of both RMSE errors denoted as e_{tot}, with 5% weight assigned to e_{fit} and 95% to e_{sim}, reflecting the priority of matching prototype startup sequences.

To assess generalization, eight experimentally recorded startup sequences are used, four for calibration (training set, comprising of sequences 1, 2, 5, and 7) and four for validation (testing set, comprising of sequences 3, 4, 6 and 8), represented in a 3D plot in Figure 1. The training set was selected based on a visual assessment to provide broad coverage of the operational space, with the aim of including sequences near the extreme boundaries of the observed domain. This approach was intended to expose the calibration process to the widest range of startup sequences. The remaining sequences, which lie more centrally within the domain, were used for validation to evaluate the model's ability to generalize to more intermediate sequences. In each NOMAD iteration, SIMSEN is executed once per training sequence, resulting in four simulations per iteration. Each simulation takes approximately 20 seconds. A maximum of 10,000 blackbox evaluations is allowed, with parameter bounds set between zero and 10,000. Calibration continues until the total weighted-sum error converges or the evaluation limit is reached. On a standard workstation, NOMAD typically converges within 3,000–4,000 iterations, requiring about one day of computation. The overall calibration methodology is illustrated in Figure 2.

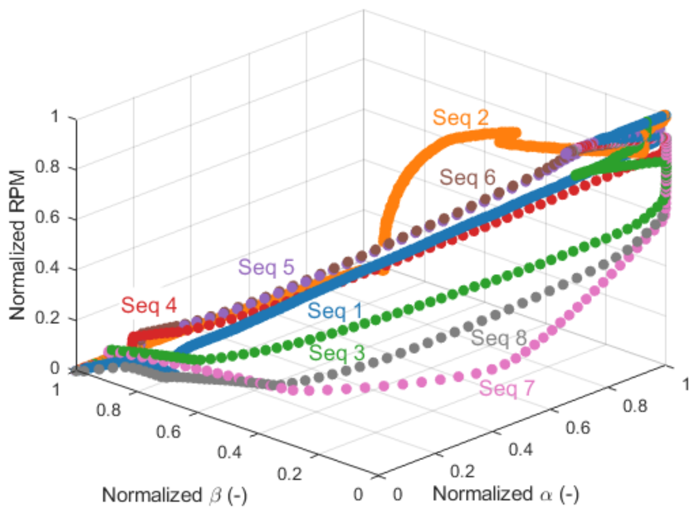

Figure 1: 3D representation of alpha, beta, and RPM across eight experimental startup sequences.

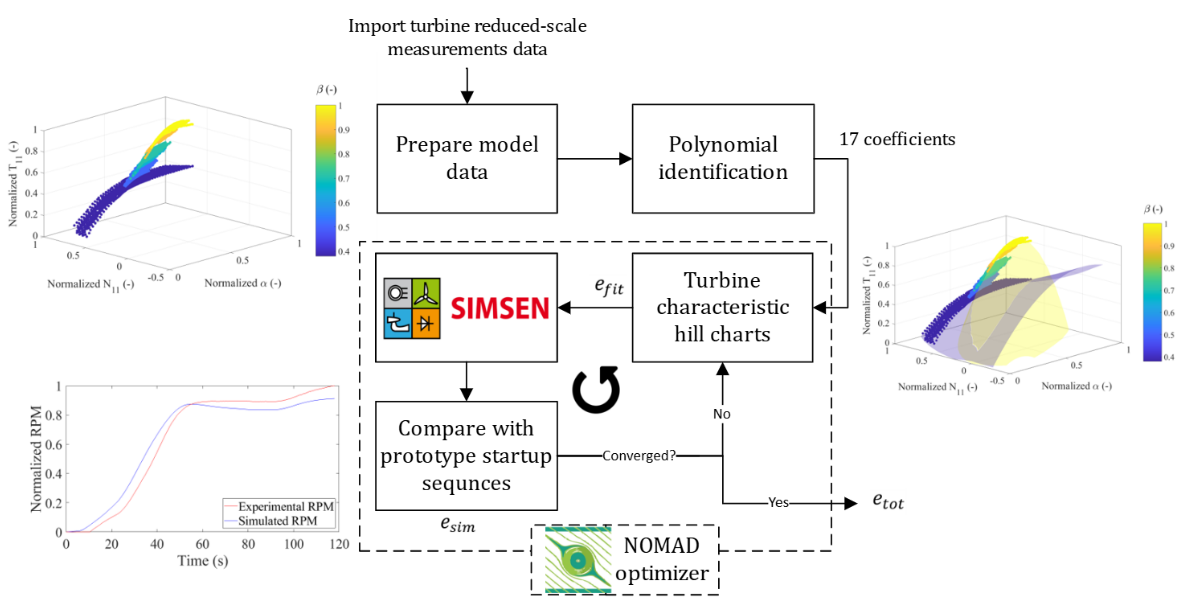

Figure 2: Workflow for calibrating turbine torque characteristics using data-driven optimization.

3 Results

The effectiveness of the proposed calibration methodology is evaluated by comparing the characteristic surfaces and startup sequence simulations of the turbine before and after optimization. Figure 3 illustrates the torque characteristic surfaces of the Kaplan turbine.Subfigure 3a shows the original

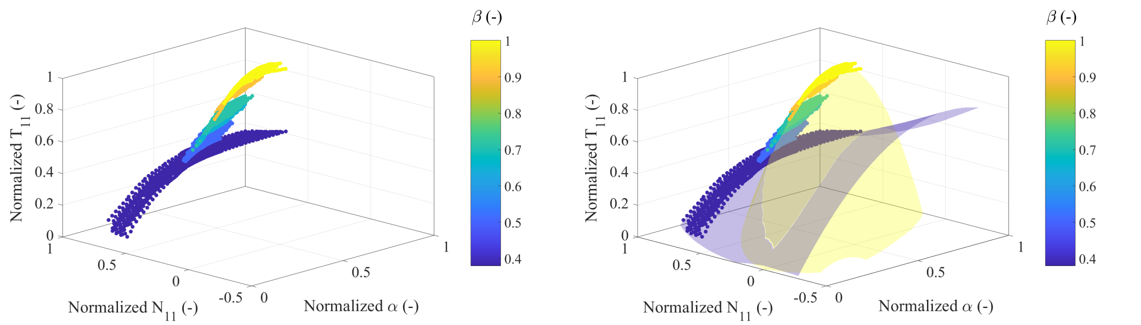

(a) Reduced-scale (model) measurement data (b) After extrapolation showing only two surfaces out of six

Figure 3: Kaplan turbine hillchart plots: (a) original reduced-scale (model) measurement data and (b) extrapolated surface showing only $\beta = 0.38$ (blue) and $\beta = 1.0$ (yellow) for visual clarity.

reduced-scale (model) measurement data, which covers only a limited portion of the operating range. After applying the calibration and extrapolation procedure, the reconstructed hillchart surface in subfigure 3b demonstrates a more complete and continuous representation of the characteristic of the turbine across a broader range of operating conditions. For clarity and to reduce visual clutter, the 3D surface is shown only for the minimum blade angle ($\beta = 0.38$), represented by the blue translucent surface, and the maximum blade angle ($\beta = 1.0$), represented by the yellow translucent surface. To set a benchmark, the dashed blue lines in Figure 4 represent the simulation results for the validation set (sequences 3, 4, 6 and 8) using the initial (uncalibrated) model. Notable discrepancies are observed between the simulated and measured rotational speed trajectories. After optimization, the same startup sequences were simulated using the calibrated model. As shown by comparing the solid black and green lines in Figure 4, the agreement between simulation and experimental data improves substantially after calibration. Results for the training set are omitted, as close agreement is expected for data seen during calibration and does not provide additional insight into the generalization capability of the model.

4 Discussion

Using a limited number of startup sequences, drawn from available monitoring data, were sufficient to tune the parameters of the model and capture the startup dynamics of a Kaplan turbine. This efficient use of data is facilitated by the ability of the BBO algorithm to converge to near-optimal parameter sets with relatively low simulation evaluations. Notably, even limited reduced-scale (model test) data proved effective for calibrating the full-scale (prototype) simulation model. This is important because comprehensive prototype performance maps are often unavailable due to testing constraints. The calibrated SIMSEN model is therefore able to interpolate between and extrapolate beyond the sparse test conditions as seen in subfigure 3b. Indeed, the model demonstrated robust generalization by accurately reproducing unseen startup sequences not included in the calibration set as seen in Figure 4, proving the effectiveness of the calibration in learning the underlying turbine dynamics. In addition, the proposed calibration framework is expected to generalize well to other turbine types such as Francis turbines, where the modeling complexity is lower due to fewer independent control variables. This suggests that the method could serve as a unified and scalable tool for hydropower DT calibration.

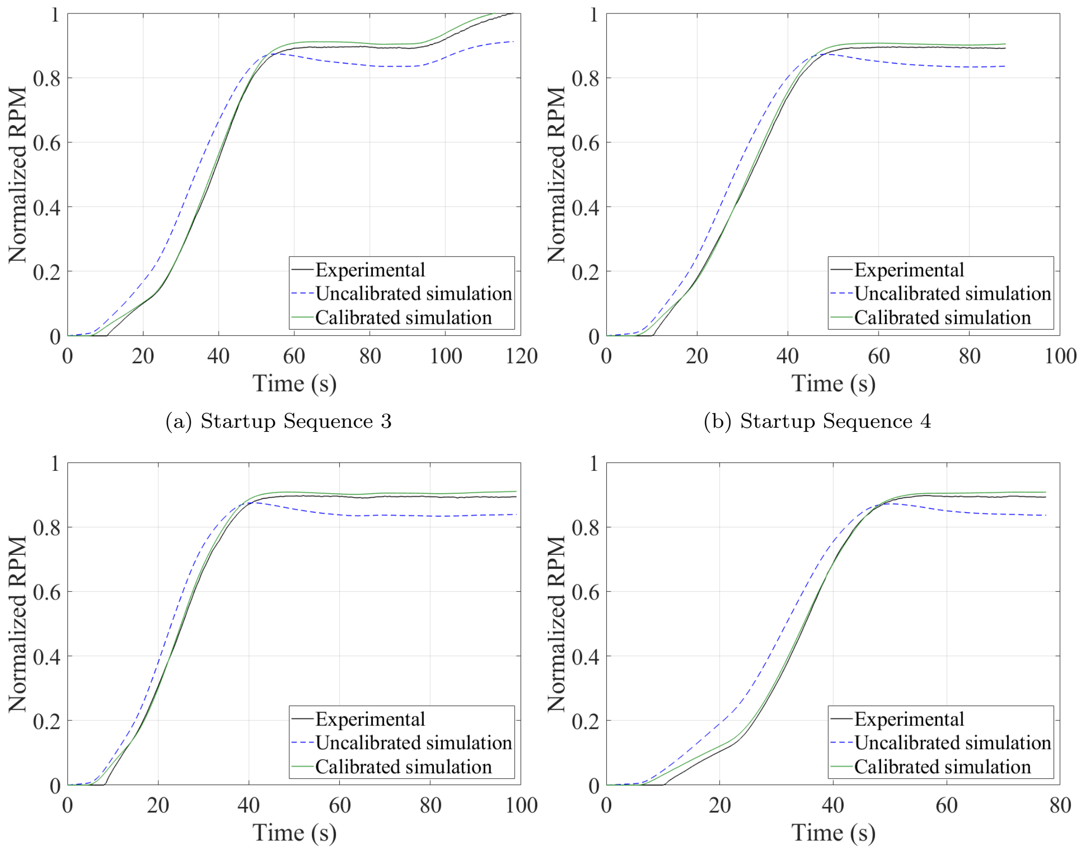

(a) Startup Sequence 3 (b) Startup Sequence 4

(c) Startup Sequence 6 (d) Startup Sequence 8

Figure 4: Normalized RPM profiles for the validation set. Each subplot compares experimental data (black solid), simulation before optimization (blue dashed), and after optimization (green solid) for: (a) startup sequence 3, (b) startup sequence 4, (c) startup sequence 6, and (d) startup sequence 8.

In future works, there are several promising avenues to enhance the capabilities of the DT. One improvement would be the integration of surrogate models as a proxy for the SIMSEN simulations. By training a surrogate on simulation input–output data, we can approximate the response surface of the turbine with a very small runtime. This surrogate could then be used in lieu of the full SIMSEN model for rapid what-if analyses or embedded real-time applications. A closely related direction is to implement active learning techniques to continually refine the calibration as new data become available. In an active learning approach, the DT would iteratively select new startup sequences to simulate or request data for, focusing on areas of the input space where the predictions of the model are most uncertain or where errors are potentially high. By updating the surrogate, and the underlying simulation model, with these targeted samples, the accuracy of the DT can be progressively improved over time. Such an approach would enable real-time or near-real-time calibration during short dedicated measurement campaigns. Ultimately, these enhancements aim to evolve the 1D hydropower unit model from an offline calibrated model to a live, self-updating DT.

5 Conclusion

The proposed approach for 1D hydropower unit model calibration enables the systematic elimination of inherent biases in the initial model, resulting in a self-correcting and highly representative simulation of the actual Kaplan turbine. Applied here to startup sequences, the calibrated model reliably reproduces the transient behavior of the turbine during these events with improved accuracy. This enhances confidence in using the model for predictive analysis and control optimization, and extends its potential value for decision support in plant operations. While the current study focused on startup transients, the same calibration framework could be extended to include other transient events to further improve the representation of the model across the full operating range. Overall, the automated calibration strategy addresses key limitations of conventional turbine models by integrating real-world data with physics-based simulations. It demonstrates how BBO techniques like NOMAD can effectively bridge the gap between reduced-scale model predictions and full-scale operational performance.

References

[1] Stéphane Alarie, Charles Audet, Aïmen E. Gheribi, Michael Kokkolaras, and Sébastien Le Digabel. Two decades of blackbox optimization applications. EURO Journal on Computational Optimization, 9(1):100011, 2021.

[2] Kaveh Amiri, Michel Jose Cervantes, and Berhanu Mulu. Experimental investigation of the hydraulic loads on the runner of a kaplan turbine model and the corresponding prototype. Journal of Hydraulic Research, 53(4):452–465, 2015.

[3] Charles Audet, Sébastien Le Digabel, and Christophe Tribes. NOMAD user guide. Technical Report G–2009–37, Les cahiers du GERAD, 2009.

[4] Manuela Böhm, Amy Robertson, Clemens Huebler, R. Rolfes, and Peter Schaumann. Optimization-based calibration of hydrodynamic drag coefficients for a semisubmersible platform using experimental data of an irregular sea state. In Journal of Physics: Conference Series, volume 1669, page 012023, 2020.

[5] Arash Soltani Dehkharqani, Fredrik Engström, Jan-Olov Aidanpää, and Michel Jose Cervantes. Experimental investigation of a 10 mw prototype kaplan turbine during start-up operation. Energies, 12(23):4582, 2019.

[6] Sébastien Le Digabel. Algorithm 909: Nomad: Nonlinear optimization with the mads algorithm. ACM Transactions on Mathematical Software, 37(4):44:1–44:15, 2011.

[7] Arthur Favrel, João Gomes Pereira Junior, Christian Landry, Sébastien Alligné, Loïc Andolfatto, Christophe Nicolet, and François Avellan. Prediction of hydro-acoustic resonances in hydropower plants by a new approach based on the concept of swirl number. Journal of Hydraulic Research, 58(1):87–104, 2020.

[8] Wencheng Guo and Zhiyuan Peng. Order reduction and dynamic performance of hydropower system with surge tank for grid-connected operation. Sustainable Energy Technologies and Assessments, 40:100777, 2020.

[9] Raluca Gabriela Iovănel, Arash Soltani Dehkharqani, and Michel Jose Cervantes. Numerical simulation of a kaplan prototype during speed-no-load operation. Energies, 15(14):5072, 2022.

[10] Khaled Mahmoud, Abdulrahman Alghamdi, and Ahmed Elkholy. Online calibration of digital twins using multi-source data for nuclear applications. Annals of Nuclear Energy, 173:109037, 2022.

[11] Christophe Nicolet, François Avellan, Philippe Allenbach, Alain Sapin, and Jean-Jacques Simond. New tool for the simulation of transient phenomena in francis turbine power plants. In Proceedings of the Hydraulic Machinery and Systems 21st IAHR Symposium, 2002.

[12] Houde Song, Meiqi Song, and Xiaojing Liu. Online autonomous calibration of digital twins using machine learning with application to nuclear power plants. Applied Energy, 326:119995, 2022.

[13] Kai Zeng, Xiaoyu Wang, and Rui Zhang. Data-driven calibration of turbine digital twins using hybrid surrogate models. Renewable Energy, 210:112947, 2025.