A Hybrid AI-Optimization Framework for Resilient Hydropower Operations to Support Grid Stability, Extreme Weather Management, and Large-Scale Industrial Loads

Quentin Ploussard[1], Arkaprabha Ganguli[2], and Jeremy Feinstein[3]

[1]Energy Systems and Infrastructure Assessment, Argonne National Laboratory
[2]Mathematics & Computer Science Division, Argonne National Laboratory
[3]Environmental Science Division, Argonne National Laboratory

Focal Areas:

This white paper addresses **water for energy** and intersects **other sectors** by supporting large, inflexible loads (e.g., data centers), proposing an integrated suite that fuses ML hydrologic prediction with stochastic scheduling to co-optimize hydropower for reliability and extreme-weather resilience.

Existing Challenge:

Hydropower operations sit between slow, point-based hydrologic forecasts that rarely quantify or propagate uncertainty and deterministic scheduling that assumes fixed inflows and favors steadiness over flexibility. As grid stress from extreme-weather shocks and continuous 24/7 loads (e.g., data centers, AI/ML) grows, this mismatch creates a forecast–operations disconnect (probabilistic information isn't produced or integrated fast enough), under-representation of extremes and compound events, scheduling that cannot hedge uncertainty or monetize flexibility, and operational inefficiencies such as spillage or missed reliability targets.

This operational silo limits fast, high-stakes decision-making and coordination across power and non-power objectives. This paper presents an opportunity to close these gaps and improve resilience and reliability by pairing probabilistic inflow prediction using generative AI with stochastic, multi-objective optimization for water power technology scheduling.

Near-Term Opportunity:

In the near-term (3–5 years), nationwide development and investment in a two-stage hybrid decision-support framework can break down the operational silo by integrating modern AI forecasting directly into advanced optimization approaches.

Stage 1: ML-based Lumped Watershed Modeling

A machine learning-based hydrologic model (e.g., those using CNNs, LSTMs, Transformers, or other neural networks) trained on decades of historical data serves as a high-speed, probabilistic "lumped watershed model" (Feinstein et al., 2023). It can assimilate real-time data (e.g., snowpack, soil moisture, precipitation forecasts) to generate probabilistic inflow forecasts (PIFs) for a reservoir or a cascading system. Unlike traditional, physics-based models, these ML approaches are purpose-built to capture the complex, nonlinear hydrologic responses characteristic of extreme weather events, and can provide operators with a clear, fast, and accurate understanding of "water-on-the-way".

Building on the lumped watershed model, another impactful development involves adding generative layers to learn the full probability distribution of inflow trajectories and synthesize new scenarios on demand. Modern diffusion and flow matching models reverse a controlled noise process to sample from complex data distributions; once trained, they can produce diverse, high-quality samples and be conditioned on hydrologic drivers to generate plausible inflow scenarios. Their generative nature allows operators to receive an ensemble of trajectories spanning typical to extreme conditions rather than a single forecast (Ou et al., 2025). However, naive applications of these models often fall short: they struggle to capture the tails of the inflow distribution and tend to behave as black boxes with little interpretability regarding known physical drivers. To address these challenges, generative modeling for hydrology can proceed in two directions:

1. **Extreme-aware generative modeling.** This direction customizes diffusion and flow matching models with heavy-tailed priors and extreme-value constraints to generate physically consistent inflow scenarios that represent the tails of the distribution, such as floods or prolonged droughts. Existing heavy-tailed diffusion models and generative hazard frameworks demonstrate how extreme-value theory improves the fidelity of rare events (Pandey et al., 2024). Extending these ideas to hydrology by training on decades of meteorological and inflow data would produce scenario ensembles that better inform risk-hedged scheduling.

2. **Disentangled and physics-guided generative modeling.** This direction builds interpretable generative models whose latent variables correspond to known physical factors. Recent generative rainfall-runoff models use low-dimensional latent variables to capture intrinsic catchment properties and climate forcing, enabling realistic runoff synthesis with minimal assumptions (Kratzert et al., 2018). By drawing on advances in disentangled representation learning (Ganguli et al., 2025; Wang et al., 2023), it becomes possible to develop models in which individual latent dimensions reflect topography, snowpack, soil moisture, temperature, or wind speed. Manipulating these latent factors would allow operators to explore "what-if" hydrologic scenarios in a transparent, physics-guided way.

Stage 2: Stochastic Mathematical Programming for Hydropower Scheduling

Probabilistic forecasts (developed in stage 1) can then be leveraged directly in a multi-stage stochastic mathematical programming (optimization) model. These types of decision models can co-optimize multiple, competing water uses simultaneously, with a clear objective to maximize total energy-water resilience. This has the potential to produce optimal, risk-hedged, day-ahead and intra-day generation and release schedules that:

- **Strengthen grid reliability:** Schedule ancillary services, ramping, and reserves to manage grid intermittency.
- **Manage extreme weather:** Use probabilistic inflows to hedge against floods (pre-position storage) and droughts (conserve water).
- **Serve new demands:** Guarantee firm, resilient power delivery to critical, high-demand loads like data centers.
- **Co-optimize other uses:** Explicitly include constraints for irrigation, water quality, and recreation.

Partnerships:

This opportunity requires a tight, cross-disciplinary partnership across Argonne National Laboratory's Energy Systems and Infrastructure Analysis (ESIA), Environmental Science (EVS), and Mathematics and Computer Science (MCS) Divisions, integrating mathematical optimization, grid security, hydrologic modeling, numerical simulation, generative AI, and uncertainty characterization; the U.S. Geological Survey (USGS) for co-development and validation of the Stage 1 ML watershed model and National Water Information System (NWIS) ground-truth data; and federal water and power operators as demonstration partners—Western Area Power Administration (WAPA) and the U.S. Bureau of Reclamation (USBR), with peer agencies/utilities such as the U.S. Army Corps of Engineers (USACE), Bonneville Power Administration (BPA), and the Tennessee Valley Authority (TVA) providing operational constraints, in-situ and hydro-asset data, and deployment pathways for resilient service to existing customers and emerging 24/7 loads.

Success Measures

Success can be measured quantitatively through tracking real operational gains and field validation: hydro units delivering higher ramping and regulation capacity (MW/min) with shorter forecast-to-dispatch latency; co-located data centers sustaining $\geq$99.99% service reliability; inflow-forecast error declining during declared extreme-weather periods relative to baseline; uneconomic or unmanaged spillage falling in high-flow events; and lower Value-at-Risk for missing flood-control or supply targets. Equally, qualitative measures of success include deploying and independently validating a replicable "soft tool" that natively links U.S. Geological Survey hydrologic science to utility scheduling, establishing a clear pathway for hydropower to provide firm, fast-ramping support for 24/7 loads and to catalyze regional economic development.

References

- Feinstein, J., Ploussard, Q., Veselka, T., Yan, E. (2023). Using Data-Prediction of Downstream 1D River Flow to Overcome the Challenges of Hydrologic River Modeling. *Water*, 15(21). https://doi.org/10.3390/w15213843

- Ganguli, A., Ramachandra, N., Bessac, J., & Constantinescu, E. (2025). Enhancing interpretability in generative modeling: Statistically disentangled latent spaces guided by generative factors in scientific datasets. *Mach. Learn.*, 114(9). https://doi.org/10.1007/s10994-025-06816-y

- Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018). Rainfall–runoff modelling using Long Short-Term Memory (LSTM) networks. *Hydrology and Earth System Sciences*, 22(11), 6005–6022. https://doi.org/10.5194/hess-22-6005-2018

- Ou, Z., Nai, C., Pan, B., Zheng, Y., Shen, C., Jiang, P., Liu, X., Tang, Q., Li, W., & Pan, M. (2025). Probabilistic Diffusion Models Advance Extreme Flood Forecasting. *Geophysical Research Letters*, 52(15), e2025GL115705. https://doi.org/10.1029/2025GL115705

- Pandey, K., Pathak, J., Xu, Y., Mandt, S., Pritchard, M., Vahdat, A., & Mardani, M. (2024). Heavy-Tailed Diffusion Models (arXiv:2410.14171). *arXiv*. https://doi.org/10.48550/arXiv.2410.14171

- Wang, Y., Schiff, Y., Gokaslan, A., Pan, W., Wang, F., De Sa, C., & Kuleshov, V. (2023). InfoDiffusion: Representation learning using information maximizing diffusion models. *Proceedings of the 40th International Conference on Machine Learning*.