Ph.D. in Information Technology
Thesis Defense

January 26th, 2026
at 5:00 p.m.
Room Alpha - Building 24

Davide SPINELLI – XXXVIII Cycle

A reinforcement learning framework for multi-objective forecast-informed reservoir operation
Supervisor: Prof. Andrea Francesco Castelletti

Abstract:

The operational management of water reservoirs can greatly benefit from increasingly accurate hydrological forecasts. Modern optimization techniques, including machine learning, enable control policies to utilize this predictive information to better address the critical challenges posed by greater hydro-climatic variability and complex, conflicting stakeholder demands. However, for dam operators to successfully adopt these advanced strategies, the resulting methods must demonstrate robustness, select the most valuable forecast information, and provide clear and understandable decisions. This dissertation advances the state-of-the-art in Forecast-Informed Reservoir Operation (FIRO) by developing a comprehensive suite of methodologies designed to create control systems that are simultaneously high-performing, robust, automated, and interpretable. This dissertation organizes the research into three primary contributions, each addressing a distinct and critical gap in modern water systems control.

First, this work confronts the dual challenge of effectively utilizing probabilistic forecast information while ensuring policy robustness against overfitting. While ensemble forecasts are the standard for representing uncertainty, their rich informational content is often discarded by reducing them to a single deterministic statistic, such as the mean. To address this, this research introduces the PECAN (Parallel Ensemble foreCAst coNtrol) algorithm, a novel method that enables running parallel simulations of the system's response natively, one for each forecast ensemble member. The results demonstrate that PECAN consistently and substantially outperforms policies based on the conventional deterministic method, capturing a wider range of solutions and achieving more effective trade-offs across the competing operational objectives. To ensure the out-of-sample validity of these data-driven policies, this research implemented a rigorous Blocked K-Fold Cross-Validation framework. This proved essential for preventing overfitting to limited historical forecast records and for identifying policies with superior generalization capabilities, while also enhancing the computational efficiency of the design process.

Second, this dissertation addresses two core challenges in the application of artificial intelligence to control problems: automated feature selection and model interpretability. Existing methods for designing control policies typically require a pre-selected set of inputs, a significant limitation when faced with an abundance of potential data sources. This work extends the NEMODPS (Neuro-Evolutionary Multi-Objective Direct Policy Search) algorithm to create an integrated framework for the concurrent optimization of a policy's inputs, internal architecture, and parameters. This novel approach automates the discovery of the most valuable information, yielding policies that match

the performance of state-of-the-art methods without requiring prior expert knowledge of which inputs to use. To render the resulting complex "black-box" policies transparent, this work leverages an adapted Time-Varying Sensitivity Analysis (TVSA). This diagnostic tool provides clear visualizations of how, when, and for which objectives the system learns to rely on different information sources. The analysis revealed sophisticated, objective-specific strategies, such as how policies focused on flood control automatically learn to rely on short-range forecasts to manage immediate risks, while policies designed to mitigate agricultural deficits learn to incorporate longer-range seasonal forecasts for strategic planning.

Finally, this research examines basin-scale governance conflicts within the heavily regulated Adda River basin in Northern Italy. Using an expanded basin-wide model that incorporates additional upstream hydropower plants, the analysis quantifies the potential benefits of ideal cooperation. Simulating a fully cooperative policy reveals that such coordination could more than halve the agricultural deficit compared to the historical baseline, at a quantifiable cost representing only a small fraction of the total potential revenue for upstream hydropower operators. Building on this benchmark, the core contribution is the design and evaluation of a novel, parametrized restitution rule, proposed as a practical, semi-cooperative instrument suitable for stakeholder negotiation. The analysis demonstrates its capacity to generate "win-win" solutions, substantially improving Lake Como's objectives at a minimal cost to upstream hydropower revenue. This work offers a starting point to guide the forthcoming renegotiation of hydropower concessions and provides a pathway for bridging the gap between advanced technical modeling and effective water governance.

PhD Committee

Prof. Matteo Giuliani, **Politecnico di Milano**

Prof. Julianne Quinn, **University of Virginia at Charlottesville**

Prof. Jonathan Herman, **University of california, Davis**