2024 European Control Conference (ECC)
June 25-28, 2024. Stockholm, Sweden

Ensemble Forecasts with Blocked K-Fold Cross-Validation in Multi-Objective Water Systems Control*

Davide Spinelli[1], Matteo Giuliani[1] and Andrea Castelletti[1]

Abstract— **In this paper, we contribute the Parallel Ensemble foreCAst coNtrol (PECAN) algorithm to enhance multi-objective water systems control through the integration of Ensemble Forecast and data-driven control techniques. This integration allows evolving parallel system simulations for each forecast ensemble member to maximize the benefit provided by the probabilistic forecasts. To avoid potential overfitting and ensure the generalization capabilities of the designed solutions, we also implement a Blocked K-Fold cross-validation. Testing on the Lake Como water reservoir system shows that PECAN improves the controller performance by 8.2% with respect to traditional methods relying solely on forecast ensemble averages and by 26.8% over approaches that do not use any forecast. These results highlight the benefits of ensemble-based techniques for controlling water systems under highly variable hydroclimatic conditions.**

I. INTRODUCTION

Climate change has accelerated the hydrological cycle, leading to increased frequency and intensity of flooding and prolonged droughts [1]. In water system control, the primary approaches to address these challenges require expanding the capacity of the system or improving the flexibility and adaptability of control policies [2]. Lately, the latter solution is gaining prominence in the literature because of its cost-effectiveness. This strategy avoids excessive investment costs and takes advantage of the opportunities presented by the steady improvement in hydrological forecasting over the past decades, particularly over longer time frames [3]. Ensemble Forecasts (EF) have proven to be a valuable approach to assessing uncertainties in weather predictions, particularly for highly variable parameters such as precipitation [4]. Yet, it is still unclear how to maximize the benefit of these forecast products when controlling water reservoir systems.

Traditionally, water reservoir control was based on Stochastic Dynamic Programming (SDP) [5]. However, SDP's requirement to model exogenous signals can become impractical when dealing with EF. Real-time control methods like Model Predictive Control [6] allow for the use of forecasts, but their performance degrades significantly with longer lead times as forecasts become more biased and less accurate and they are directly used to solve the modeled mass-balance equation. Moreover, these methods require scalarization of the objective space to handle multiple objectives. An alternative approach is found in Reinforcement Learning methods that involve Approximation in Policy Space, such as Direct Policy Search (DPS) [7]. DPS is a simulation-based learning method that seeks to identify a policy that optimizes a specific objective function. It works by parameterizing the policy within a class of functions and performing the optimization directly in the policy parameter space. For the uncontrolled component of this partially data-driven approach, a model of the disturbance realizations is substituted with historical time series of observed or forecasted hydrometeorological variables. The parameterization allows DPS to accommodate a large number of time series as policy inputs, making it well-suited for effectively incorporating EF over extended lead times. However, the control policy is usually informed with deterministic forecasts or pre-defined statistics of the EF without taking full advantage of its probabilistic information.

Being a simulation-based approach, DPS relies on long simulations to approximate the expected value of the objective function over the probability distribution of disturbances using the average value from a sufficiently long time series of disturbance realizations. In contrast, the availability of skillful forecasts is restricted to recent years, limiting the essential validation of newly developed algorithms to a short time frame. Adopting DPS to design forecast-informed control policies in systems with limited forecast data is prone to the risk of overfitting, where policies are excessively tailored to training data, yielding favorable results but proving less effective when applied to different datasets.

In our study, we introduce a novel algorithm with the twofold goal of (*i*) making use of the full information provided by EF without reducing it to simpler statistics of the ensemble by simulating in parallel virtual systems for each ensemble member and (*ii*) implementing the Blocked K-Fold cross-validation [8] technique to reduce the risk of overfitting the designed control policies in order to ensure their generalization to unseen conditions. Our method is demonstrated using the real-world case study of the Lake Como system in Northern Italy, a multipurpose regulated lake operated for water supply and flood and drought control.

II. DATA AND MODELS

A. Case study

Lake Como, located in Northern Italy, is a regulated lake within the Adda River basin. It has a surface area of 145 km^2 and an active storage capacity of 246.5 Mm^3 that is operated

*This work was partially supported by the Italian Ministry of University and Research (MUR) under the National Recovery and Resilience Plan (NRRP), and by the European Union (EU) under the NextGenerationEU project, grant number 38-412-16-DOT1316508-789. This work has also been partially funded by the CLINT project under the H2020 framework programme of the European Union, grant number 101003876.

[1]Department of Electronics, Information and Bioengineering, Politecnico di Milano, Piazza Leonardo da Vinci 32, Milano, Italy
{davide.spinelli, matteo.giuliani, andrea.castelletti}@polimi.it

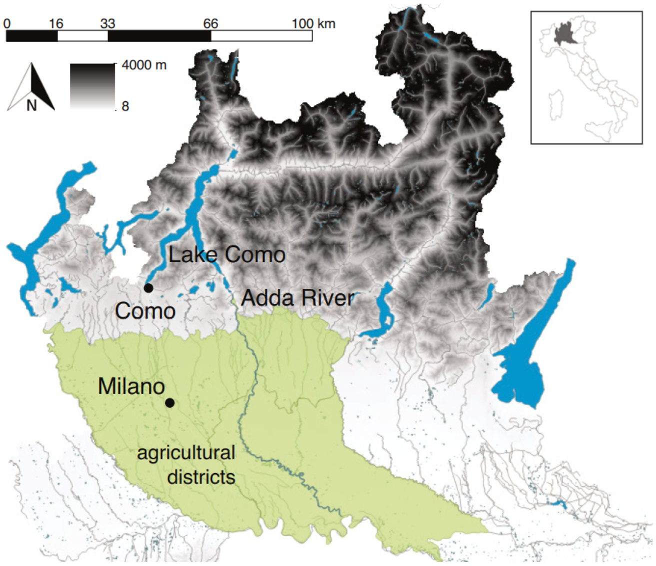

Fig. 1. Map of the Lake Como basin. (Reprinted, with permission, from [9] @ Author(s) 2020 under CC BY 4.0)

under the management of the Italian regional authority 'Consorzio dell'Adda.' The lake's basin covers a catchment area of 4552 km^2 and exhibits a characteristic subalpine hydrological flow pattern of glacial origin, featuring two distinct peaks, one resulting from spring snowmelt and the other from autumn precipitation.

The regulation of Lake Como serves three primary and competing objectives: it provides flood protection to the city center of Como; it supplies downstream water to support the 1400 km^2 irrigation district and nine run-of-river power plants; and it plays a crucial role in drought control, with implications for navigation, environmental considerations, and tourism. Additionally, the regulation must adhere to the preservation of a Minimum Environmental Flow (MEF) to safeguard the downstream ecosystem and environment.

B. Model, objectives, and problem formulation

The Lake Como system dynamics is modeled by a conventional mass balance equation of the lake water volume:

$$s_{t+1} = s_t + q_{t+1} - r_{t+1} \tag{1}$$

In this equation, s_t represents the lake storage at time t, q_{t+1} denotes the net inflow observed over a 24-hour period, and r_{t+1} represents the actual outflow obtained by a nonlinear and stochastic function, $R(\cdot)$, of the control decision u_t, i.e. $r_{t+1} = R(s_t, u_t, q_{t+1})$. The actual release may not be equal to the decision due to legal and physical constraints on the level and release of the reservoir, including spills when the reservoir exceeds the maximum capacity [10]. The control decision u_t is determined by the policy π following the control rule $u_t = \pi(t, s_t, \hat{q}_{t+1})$, where $\hat{q}_{t+1}$ represents the inflow prediction.

The objectives are modeled as follows:

- flood control: annual average excess of lake level above a flooding threshold $h^{flo} = 1.1m$:

$$J^{flo} = \frac{1}{H/T}\sum_{t=0}^{H-1} \max(0, h_{t+1} - h^{flo}) \tag{2}$$

- water supply: daily average deficit of water release compared to a cyclostationary daily downstream demand w_τ:

$$J^{def} = \frac{1}{H}\sum_{t=0}^{H-1} (\max(0, w_\tau - r_{t+1}))^{\beta_\tau} \tag{3}$$

  β_τ represents a time-varying exponent designed to increase the penalty during the summer period (April 1$^{\text{st}}$ to October 10$^{\text{th}}$), emulating dam operators' behavior.
- low-level control: annual average lake level scarcity compared to low-level threshold $h^{low} = -0.2m$:

$$J^{low} = -\frac{1}{H/T}\sum_{t=0}^{H-1} \min(0, h_{t+1} - h^{low}) \tag{4}$$

In these equations, the symbols represent the following: h, the lake level; H, the simulated horizon; T, the hydrological year; t, the day of the simulation, ranging from 0 to H; τ, the day of the year, with values between 1 and 365.

The problem of designing the optimal control policies for Lake Como can be formulated as a multi-objective, stochastic, periodic, non-linear, closed-loop control problem [11] defined as follows:

$$\pi^* = \arg\min_{\pi} \mathbf{J}(\pi) \tag{5}$$

where $\mathbf{J}(\pi) = \left|J^{flo}(\pi), J^{def}(\pi), J^{low}(\pi)\right|$ is the vector of the expected returns under policy π. The optimization problem is subject to the dynamics of the system (Eq. 1).

C. Data

Observational data for the lake have been provided by Consorzio dell'Adda and are available from 1946 at a daily resolution. The calculation of the total net inflow is derived from the inversion of the mass balance equation (1), using time series data for storage and release, which consider various tributaries and lake evaporation.

The ensemble inflow forecasts used in this study are sourced from the European Flood Awareness System (EFAS) [12], [13], a component of the Copernicus Emergency Management Service (CEMS) of the European Union. The hydrological forecasts generated by EFAS are produced by the European Center for Medium-range Weather Forecasts (ECMWF). Archived real-time forecasts are available from 2018 onward. However, due to the need for a longer historical context, we use reforecast products at both medium [12] and seasonal scales [13]. These re-forecasts span the period from 1999 to 2018 and are categorized into sub-seasonal and seasonal variants. The sub-seasonal re-forecasts have a maximum lead time of 46 days, consist of 11 ensemble members, and are initialized twice weekly. In contrast, seasonal reforecasts offer a maximum lead time of 215 days, comprise 25 members, and are initialized monthly.

We consider three specific lead times: 5 days, 21 days, and 180 days (equivalent to six months). This lead time selection was based on previous research on Lake Como conducted by Zanutto et al. [14]. In their study, these lead times were determined using the Input Selection and Assessment framework [15]. The selected inputs involve the aggregation of forecast averages up to the specified lead time. In our work, we employ three different combinations of lead times:

- 21d: a single-input experiment utilizing only the 21-day aggregated forecast;
- 5d-21d: a two-input experiment incorporating both the 5-day and the 21-day aggregated forecast;
- 5d-21d-180d: a three-input experiment integrating all three aggregated forecasts.

In addition to using real forecasts, we consider aggregated perfect forecasts for all three combinations of lead times to determine the maximum achievable results for the optimization in each scenario.

III. METHODS

A. Evolutionary Multi-Objective Direct Policy Search

PECAN is based on Evolutionary Multi-Objective Direct Policy Search (EMODPS) [16]. EMODPS is a Reinforcement Learning, Direct Policy Search approach known for its partially data-driven nature and effectiveness in addressing multiple competing objectives. Specifically, we adopted Gaussian Radial Basis Functions (RBF, [17]) for the parameterization of the control policy as a universal nonlinear approximating network. The policy is parameterized as a single hidden layer comprising N basis functions ($i = 1,\dots,N$), M inputs ($j = 1,\dots,M$), and K outputs ($k = 1,\dots,K$), representing the control decision u_t^k at time t in the vector $\mathbf{u}_t$. We incorporate three fixed inputs: two for the day of the year d_t, represented as $\sin(2\pi d_t/T)$ and $\cos(2\pi d_t/T)$ to capture the problem's periodicity and one for the current lake level representing the system's state. Additionally, we integrate one to three further inputs for the forecasted inflow values. We use $N = M+2$ for the number of nodes in the hidden layer. Following the parameterization introduced in [17], the control rule determining the control decision is thus expressed as

$$u_t^k = o^k + \sum_{i=1}^{N} w_{i,k}\phi_i(\chi_t) \tag{6}$$

Here, o^k is a linear parameter associated with each output, $w_{i,k}$ are non-negative weights ($w_{i,k} \geq 0, \forall i,k$) for the i-th RBF. The function $\phi_i(\chi_t)$ is the value of the i-th RBF, defined as:

$$\phi_i(\chi_t) = \exp\left[-\sum_{j=1}^{M} \frac{((\chi_t)_j - c_{i,j})^2}{b_{i,j}^2}\right] \tag{7}$$

In this equation, χ_t represents the tuple of inputs at time t, $c_{i,j}$ signifies the center of the RBF, and $b_{i,j}$ denotes its positive radius. The entire RBF parameter combination has a cardinality of $n_\theta = K + N(2M+K)$. In the case of the Lake Como system, there is a single control action (i.e. K=1).

Multi-Objective Evolutionary Algorithms (MOEAs) are optimization algorithms that evolve an entire set of Pareto-approximate solutions, mirroring the principles of natural genetic evolution. They are particularly well-suited for addressing the challenges posed by multi-objective problems [18]. Within the framework of EMODPS, the element in the population is the set of policy parameters denoted as θ. These parameters are evaluated based on the three objective functions (2,3,4) following the model evolution in (1). The search to identify the approximate Pareto front is performed by employing Borg MOEA [19]. This framework is chosen for its capacity to evade local minima and consistently promote progress and diversity during the optimization process [20]. In alignment with previous studies [21][22], we conducted 20 randomized optimization runs, each simulated over 2 million NFE (Number of Function Evaluations) with the Borg ε-box set to [5,50,10].

B. Cross-validation

In the design of the EMODPS policy, optimization is carried out using a designated data set as training data. When policy parameters are trained and evaluated on this same dataset, it can lead to overfitting. In such cases, the optimized policies perform well on the training data, but fail to generalize effectively to unseen data. This limitation is increasingly critical in the control of water systems due to the non-stationary effects of climate change.

Emerging approaches verify possible overfitting by splitting the available data into training and test sets, with the latter often artificially generated [23]. However, this approach can produce skewed results due to possible differences in training and test data. To address this, we propose to extend this approach by performing a Blocked K-Fold cross-validation.

In Blocked K-Fold cross-validation, the dataset is partitioned into K blocks of continuous observations, known as folds. Each fold is selected iteratively as the test set, while the remaining $K-1$ folds serve as the training set. After the training phase, the set of Pareto-efficient policies is evaluated on the test set, and the results of each iteration are averaged. This method optimizes the utilization of the limited dataset, as every data point can be included in the test set at some point. This approach helps mitigate the skewed results caused by variations in climatic conditions across different time periods.

Our dataset spans 20 years, from 1999 to 2018. We divided our data set into five folds, each covering four years, each fold beginning in January and concluding at the end of December of the fourth year. We opted to group multiple years together to prevent the need to calibrate a penalty at the end of each year. This approach allows us to leverage the water carryover from year to year, acting as a mitigating factor, thus eliminating the necessity for year-end penalties in the simulation.

C. Parallel Ensemble foreCAst coNtrol

In this section, we introduce the proposed PECAN (Parallel Ensemble foreCAst coNtrol) algorithm (see Alg. 1). PECAN leverages EF data to compute control decisions and incorporates constraint models into the system's regulation.

Algorithm 1: Parallel Ensemble foreCAst coNtrol

input : t day of the simulation; HH number of integration steps during the day; s_t state of the system at the beginning of day t; V number of ensemble members; $\hat{q}_t^v$ estimated lake inflow at time t for ensemble member v; $\mathbf{q}_t = \{q_{t,1}, \ldots, q_{t,HH}\}$ set of observed inflows at each integration step.

output: New state of the system $s_{t,HH} \equiv s_{t+1,0}$

for $v \in [1,V]$ **do**
    $u_t^v = \pi(t, s_t, \hat{q}_t^v)$
    $s_{t,0}^v = s_t$
end
for $h \in [0, HH-1]$ **do**
    for $v \in [1,V]$ **do**
        $r_{t,h+1}^v = R(s_{t,h}^v, u_t^v, q_{t,h+1})$
        $s_{t,h+1}^v = s_{t,h}^v + q_{t,h+1} - r_{t,h+1}^v$
    end
    $\overline{u_{t,h}} = \frac{1}{V}\sum_{v=1}^{V} r_{t,h+1}^v$
    $r_{t,h+1} = R(s_{t,h}, \overline{u_{t,h}}, q_{t,h+1})$
    $s_{t,h+1} = s_{t,h} + q_{t,h+1} - r_{t,h+1}$
end
return $s_{t,HH}$

The algorithm operates at each daily time step t of the simulation, divided into $HH = 24$ hourly integration steps to emulate the actions of dam operators at 1-hour intervals.

The novelty of our approach lies in the central role of the forecasts' ensemble nature, integrated into our algorithm in their entirety. In previous studies [14], forecasts were either a single deterministic value or derived from statistical measures of the ensemble (such as average, minimum, maximum, quantile, etc.). These statistical measures were treated as deterministic inputs to the policy, which computed a single control decision and applied it to the system.

In contrast, at each time step, PECAN computes an independent control decision for each member of the Ensemble Forecasts. It then initializes a virtual system for each ensemble member based on the current state and applies each release decision to the respective virtual system. Keeping the virtual systems independent, PECAN simulates each system's transition for the integration step duration. This results in an ensemble of actual releases incorporating both the system and the constraint models. Finally, it applies the average of these releases to the original system, repeating this process for each integration step while maintaining the separation of the virtual systems throughout the day.

PECAN begins by computing the control decision u_t^v for each member of the EF, following the policy π. This results in a set of control decisions denoted as $\mathbf{u}_t = \{u_t^1, \ldots, u_t^V\}$, where V represents the number of ensemble members. Then, the system state s_t at time t is replicated for each control decision, creating the set of virtual lakes $\mathbf{s}_{t,0} = \{s_{t,0}^1, \ldots, s_{t,0}^V\}$ where $(t,0)$ represents time t at the integration step 0 and $s_{t,0}^v = s_t \; \forall v \in [1,V]$. Then, the algorithm performs the hourly integration of the state transition function (eq. 1) independently for each ensemble member by calculating the constrained release volume for each control decision u_t^v, i.e., $r_{t,h+1}^v = R(s_{t,h}^v, u_t^v, q_{t,h+1})$.

After performing these operations for all members of the ensemble, the algorithm generates V hourly releases and updates the state of the system accordingly, resulting in potentially diverse states depending on the member of the considered ensemble. The constrained releases are finally averaged into $\overline{u_{t,h}}$ and, after a final check against the system constraints, we obtain the hourly actual release $r_{t,h+1} = R(s_{t,h}, \overline{u_{t,h}}, q_{t,h+1})$. The actual release is then applied to the system state $s_{t,h}$ resulting in the new state of the actual lake $s_{t,h+1} = s_{t,h} + q_{t,h+1} - r_{t,h+1}$. This process is iterated for each subsequent integration step until the end of the day, ensuring that the set of virtual lakes remains independently updated. This approach results in the application of constraints in various ways in different integration steps based on the diverse forecasts within the ensemble. These differences are preserved thanks to averaging all decisions, ultimately impacting the controlled system dynamics.

D. Benchmarking framework

We employ two distinct sets of policies to facilitate a comparative analysis of the different methods: the Basic Operating Policy (BOP) and the Perfect Operating Policy (POP). The BOP serves as the foundational framework for the current system, covering policies derived using EMODPS with solely the fixed inputs, without consideration of any forecasts for future inflows. In contrast, the POP consists of optimal policies achieved through Deterministic Dynamic Programming, employing the weighting method [5]. This approach assumes perfect knowledge of inflows for the entire operational horizon.

We employ the Hypervolume indicator (HV) [24] to evaluate the Pareto solutions obtained with Borg. This metric is commonly used in multi-objective optimization and quantifies the volume dominated by the Pareto front within the Q-dimensional space of the objectives. As the reference point for the HV, we chose a combination within the objective space representing the worst result among the BOP policies. The Hypervolume is calculated relative to the optimal Pareto front of the POP, with $HV^{POP} = 1$.

IV. RESULTS

A. Improvements comparison

Figure 2 provides a comparative analysis of the quality of the Pareto approximate sets for different solution methods evaluated over the test set. When incorporating Blocked K-Fold cross-validation, PECAN achieves a significant +26.75% improvement over the Basic Operating Policy (BOP), which does not use forecasts. By leveraging ensemble information, PECAN consistently outperforms policies designed using the average of the EF across all experiments, as used in previous works [9], showing an average improvement of +8.23% compared to relying solely on the ensemble average. Furthermore, using perfect information at the same

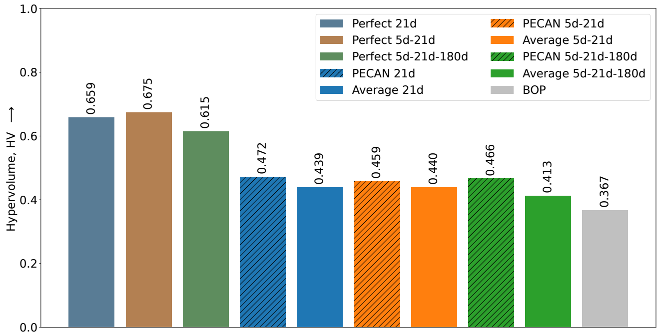

Fig. 2. Hypervolume values obtained by the different methods. The plain bars represent the performance when only the ensemble average is used, where blue signifies policies with single-input 21-day forecasts, orange with two-input 5-day and 21-day forecasts, and green with three-input 5-day, 21-day, and 180-day forecasts. The diagonal hatched bars represent the performance achieved by PECAN, employing the same color coding. For reference, we provide the results of the Basic Operating Policy (BOP) in grey, and of policies using perfect forecasts at the same lead times. These reference bars are benchmarks for understanding the minimum and maximum achievable results given the specified lead times.

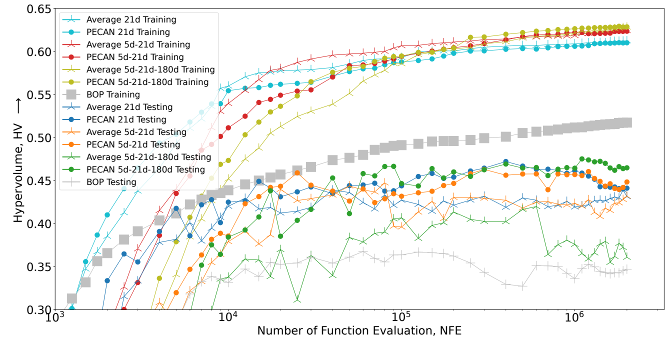

Fig. 3. Evolution of HV for the different methods evaluated, both on the training and test sets.

lead times presents an average improvement of +76.84% over the BOP. These results underscore the complexity of the problem and the potential for improvement of novel methods, like PECAN, that harness the uncertainty of the entire EF, achieving better results than using only a statistic of the ensemble.

B. *Cross-validation analysis*

The results highlight the superior generalization capabilities of PECAN, which become particularly evident when comparing the outcomes with and without cross-validation. Figure 3 illustrates the HV progression over the number of function evaluations (NFE) for policies evaluated on both the training set and the test sets using Blocked K-Fold cross-validation. This analysis clearly demonstrates the presence of overfitting in existing methods and PECAN, underscoring the substantial impact of overfitting on the results and highlighting the importance of cross-validation. When policies are evaluated on the training set, both the ensemble average and PECAN exhibit similar performance after 2 million NFE. Interestingly, PECAN performs slightly worse when using the 21-day or 5-day and 21-day forecast inputs. A more in-depth examination of the training set results reveals that policies continue to learn beyond 2 million NFE, although the rate of improvement significantly diminishes after exceeding 1 million NFE. On the contrary, when policies are evaluated on the test set, the learning process stops after approximately 150,000 NFE, on average. Notably, this stopping point occurs later when more forecast inputs are included, reflecting the added complexity of optimizing policies with additional parameters. Beyond this stage, the HV starts to decline, with occasional spikes that have minimal impact on the set of solutions. This analysis highlights the value of cross-validation, as it not only evaluates the generalization capabilities of solutions but also does so without increasing the computational resources required. Without cross-validation, the NFE required to reach stable solutions is approximately one order of magnitude higher than the NFE needed to reach the stopping point when using cross-validation. Despite the additional optimization runs, cross-validation results in an earlier stopping point and reduces the overall computational time.

C. *Pareto front distributions*

The application of cross-validation with multiple folds enabled us to conduct an in-depth analysis of policy behavior under diverse validation scenarios and assess the effectiveness and distribution of their outcomes within the three-dimensional objective space. Figure 4 presents the Pareto front results across the five folds for both PECAN and average methods, considering only the 21-day input. Notably, the coverage of the objective space is most prominent in folds 3 and 5, indicating PECAN's ability to provide a wide range of high-quality solutions. However, specific challenges arise, as demonstrated by the limited solutions obtained in fold 2, characterized by an extended period of exceptionally low water availability. The observed variations across different folds emphasize the importance of employing Blocked K-Fold cross-validation. This approach ensures unbiased and reliable assessments, especially when dealing with varying climatological conditions.

V. CONCLUSIONS

In this study, we introduced the novel algorithm PECAN (Parallel Ensemble foreCAst coNtrol) to enhance the performance of multi-objective policy optimization for reservoir control. PECAN leverages the full information of Ensemble Forecasts (EF) and incorporates system and constraint models, aiming to improve the quality and robustness of policy solutions. Our findings demonstrate that PECAN consistently outperforms using the average of the EF, with an average improvement of 8.23% in our case study, signifying its promise for enhanced water resource management. We also highlighted the importance of cross-validation in assessing the generalization capabilities of the policies, particularly under varying climatic conditions, to ensure unbiased results.

Future work in this field could explore the optimal selection of lead times based on ensemble uncertainty and the

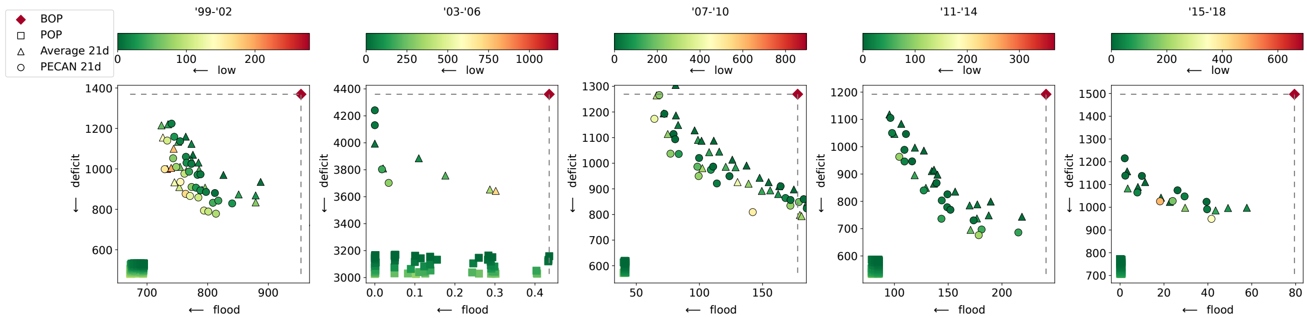

Fig. 4. Performance of Pareto optimal policies informed by the average and PECAN methods across the five folds, indicated on top with the four-year period used as the test set. The axes represent the flood and deficit objectives, while the color encodes the low-level objective based on the provided colormap. The diamond symbolizes the reference point for the HV, squares denote the set of best solutions from the Perfect Operating Policy (POP), and triangles and circles represent the solutions generated by the average and PECAN methods, respectively.

development of fold-specific parameters to ensure comprehensive coverage of solutions in the objective space across all folds. Additionally, given the room for improvement demonstrated by the results obtained with perfect future information, we conclude that there is ample scope for advancements in harnessing forecast ensemble data effectively.

ACKNOWLEDGMENT

The authors would like to thank Consorzio dell'Adda, the European Centre for Medium-Range Weather Forecasts (ECMWF), and the Copernicus Emergency Management (CEMS) Service for providing the data used in this study.

REFERENCES

[1] Stevenson, S., Coats, S., Touma, D., Cole, J., Lehner, F., Fasullo, J. & Otto-Bliesner, B. Twenty-first century hydroclimate: A continually changing baseline, with more frequent extremes. *Proceedings Of The National Academy Of Sciences*. **119**, e2108124119 (2022,3)

[2] Gleick, P. & Palaniappan, M. Peak water limits to freshwater withdrawal and use. *Proceedings Of The National Academy Of Sciences*. **107**, 11155-11162 (2010,6)

[3] Min, Y., Ham, S., Yoo, J. & Han, S. Recent Progress and Future Prospects of Subseasonal and Seasonal Climate Predictions. *Bulletin Of The American Meteorological Society*. **101**, E640-E644 (2020,5)

[4] Hailing, Z. & Pu, Z. Beating the Uncertainties: Ensemble Forecasting and Ensemble-Based Data Assimilation in Modern Numerical Weather Prediction. *Advances In Meteorology*. **2010** (2010,1)

[5] Bellman, R. Dynamic Programming. (Princeton University Press,1957)

[6] Castelletti, A., Ficchì, A., Cominola, A., Segovia, P., Giuliani, M., Wu, W., Lucia, S., Ocampo-Martinez, C., De Schutter, B. & Maestre, J. Model Predictive Control of water resources systems: A review and research agenda. *Annual Reviews In Control*. **55** pp. 442-465 (2023,1)

[7] Schmidhuber, J. Sequential Decision Making Based on Direct Search. *Sequence Learning: Paradigms, Algorithms, And Applications*. pp. 213-240 (2001)

[8] Cerqueira, V., Torgo, L. & Soares, C. Model Selection for Time Series Forecasting: Empirical Analysis of Different Estimators. (arXiv,2022,2)

[9] Giuliani, M., Crochemore, L., Pechlivanidis, I. & Castelletti, A. From skill to value: isolating the influence of end user behavior on seasonal forecast assessment. *Hydrology And Earth System Sciences*. **24**, 5891-5902 (2020,12)

[10] Soncini-Sessa, R., Weber, E. & Castelletti, A. Integrated and Participatory Water Resources Management - Theory. (Elsevier Science,2007,10)

[11] Castelletti, A., Pianosi, F. & Soncini-Sessa, R. Water reservoir control under economic, social and environmental constraints. *Automatica*. **44** pp. 1595-1607 (2008,6)

[12] Barnard, C., Krzeminski, B., Mazzetti, C., Decremer, D., Wiart, C., Harrigan, S., Blick, M., Ferrario, I., Wetterhall, F., Thiemig, V., Salamon, P. & Prudhomme, C. Reforecasts of river discharge and related data by the European Flood Awareness System. (ECMWF,2020), https://cds.climate.copernicus.eu/doi/10.24381/cds.c83f560f

[13] Wetterhall, F., Arnal., L., Barnard, C., Krzeminski, B., Ferrario, I., Mazzetti, C., Thiemig, V., Salamon, P. & Prudhomme, C. Seasonal forecasts of river discharge and related data by the European Flood Awareness System. (ECMWF,2020), https://cds.climate.copernicus.eu/doi/10.24381/cds.eb224b0e

[14] Zanutto, D. Multi-timescale forecast-based direct policy search for water system control. (2022), https://hdl.handle.net/10589/190646

[15] Giuliani, M., Pianosi, F. & Castelletti, A. Making the most of data: An information selection and assessment framework to improve water systems operations. *Water Resources Research*. **51**, 9073-9093 (2015)

[16] Giuliani, M., Herman, J., Castelletti, A. & Reed, P. Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management. *Water Resources Research*. **50** (2014,4)

[17] Busoniu, L., Ernst, D., De Schutter, B. & Babuska, R. Cross-Entropy Optimization of Control Policies With Adaptive Basis Functions. *IEEE Transactions On Systems, Man, And Cybernetics, Part B (Cybernetics)*. **41**, 196-209 (2011,2)

[18] Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A., Karamouz, M., Minsker, B., Ostfeld, A., Singh, A., Zechman, E. & Null, N. State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management. *Journal Of Water Resources Planning And Management*. **136**, 412-432 (2010,7)

[19] Hadka, D. & Reed, P. Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework. *Evolutionary Computation*. **21** (2012,3)

[20] Zatarain Salazar, J., Reed, P., Herman, J., Giuliani, M. & Castelletti, A. A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances In Water Resources*. **92** pp. 172-185 (2016,6)

[21] Giuliani, M., Castelletti, A., Pianosi, F., Mason, E. & Reed, P. Curses, Tradeoffs, and Scalable Management: Advancing Evolutionary Multiobjective Direct Policy Search to Improve Water Reservoir Operations. *Journal Of Water Resources Planning And Management*. **142**, 04015050 (2016,2)

[22] Denaro, S., Anghileri, D., Giuliani, M. & Castelletti, A. Informing the operations of water reservoirs over multiple temporal scales by direct use of hydro-meteorological data. *Advances In Water Resources*. **103** pp. 51-63 (2017,5)

[23] Quinn, J., Reed, P., Giuliani, M. & Castelletti, A. Rival framings: A framework for discovering how problem formulation uncertainties shape risk management trade-offs in water resources systems. *Water Resources Research*. **53**, 7208-7233 (2017)

[24] Zitzler, E. & Thiele, L. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Transactions On Evolutionary Computation*. **3**, 257-271 (1999,11)