

Data-driven control of water reservoirs using an emulator of the climate system

Matteo Giuliani* Marta Zaniolo* Paul Block**
Andrea Castelletti*

* *Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milan, Italy (e-mail: matteo.giuliani@polimi.it, marta.zaniolo@polimi.it, andrea.castelletti@polimi.it).*

** *Department of Civil and Environmental Engineering, University of Wisconsin-Madison, Madison, WI (e-mail: paul.block@wisc.edu).*

Abstract: This study presents a novel approach to combine a data-driven control strategy with an emulator model of the climate system in order to make the optimal control of water systems more flexible and adaptive to the increasing frequency and intensity of extreme events. These latter are often associated with global climate anomalies, which are difficult to model and incorporate into optimal control algorithms. In this paper, we compare a traditional control policy conditioned only on the reservoir storage with an informed controller that enlarges the state space to include the emulated dynamics of global Sea Surface Temperature anomalies. The multi-purpose operations of Lake Como in Italy, accounting for flood control and water supply, is used as a case study. Numerical results show that the proposed approach provides a 59% improvement in system performance with respect to traditional solutions. This gain further increases during extreme drought episodes, which are influenced by global climate oscillations.

Copyright © 2020 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Optimal control of water resources systems; Data-driven control; Model reduction; Machine Learning; Direct Policy Search; Multi-objective optimal control

1. INTRODUCTION

The problem of designing optimal reservoir operations has been extensively studied since the seminal work by Rippl (1883) and a number of review papers have been published over the years to describe the evolving state-of-the-art in the field (e.g., Yeh, 1985; Labadie, 2004; Castelletti et al., 2008). Since the early applications in the 1960s by Hall and Buras (1961) and Maass et al. (1962), Dynamic Programming (DP) and its stochastic extension (SDP) are probably the most widely adopted methods in the reservoir operations literature. SDP formulates a Markov decision process (MDP) that requires sequential decisions at each time step that produce an immediate cost and affect the next system state, thereby affecting all the subsequent costs. The MDP formulation has been considered particularly suitable for modeling water resources systems as it removes the simplifying assumptions of linear models, quadratic costs, and white Gaussian disturbances required in the LQG framework (Baglietto et al., 2006).

Yet, the computation of the optimal value function substantially limits the adoption of SDP in complex real-world problems due to some limitations: (1) the well known curse of dimensionality (Bellman, 1957), which limits the dimension of the system to two or three reservoirs due to the exponential growth of computational cost with the number of state variables; (2) the curse of modeling (Tsitsiklis and Van Roy, 1996), which requires all variables used as input in the operating policy to be described by a dynamic model, thus contributing additional state variables; (3) the curse of multiple objectives (Powell, 2007),

which restricts the number of objective functions due to the single-objective nature of SDP that requires repeated scalarized single-objective optimizations for every Pareto optimal point, inducing a factorial growth of computation cost with the number of objectives (Giuliani et al., 2014). Given the presence of these curses, the majority of SDP applications cannot preserve the full complexity of the real water systems and, generally, restrict the problem formulation to focus solely on the dynamics of the reservoir system. The influence of the hydrologic and climate system on the reservoir dynamics is often neglected or overly simplified in a SDP-based formulation. Recent advances in Model Predictive Control (e.g., Raso et al., 2014; Galelli et al., 2015; Ficchi et al., 2016; Tian et al., 2019) provide a better characterization of inflow variability and uncertainty by means of ensemble forecast; however, forecast models may introduce modeling errors and forecast biases, which may negatively impact on the system performance (Giuliani et al., 2019).

In this paper, we contribute a different approach that, rather than relying on inflow forecasts, proposes to enlarge the state of the water system used for conditioning closed loop control policies beyond the reservoir storage. We argue that the observation of such enlarged state of the water system allows capturing hydroclimatic conditions, possibly including global climate oscillations, that are known to impact on current and future local hydrologic dynamics (e.g., Emerton et al., 2017). Specifically, we first search the Sea Surface Temperature anomalies at the global scale that are statistically correlated with local

hydroclimatic conditions. This teleconnection analysis is run separately depending on the phase of a climate signal (Lee et al., 2018), such as El Niño and La Niña for the El Niño Southern Oscillation (ENSO) and the positive and negative phases of the North Atlantic Oscillation (NAO), as measured by their corresponding indexes. Then, we perform a model reduction via Proper Orthogonal Decomposition (Chatterjee, 2000) to build an emulator of the selected climate variables. To directly use the output of the climate system emulator as input in the control policy design, we rely on the Evolutionary Multi-Objective Direct Policy Search method (Giuliani et al., 2016a), which implements a data-driven control strategy (Formentin et al., 2013) by integrating direct policy search, nonlinear approximating networks, and multi-objective evolutionary algorithms. The approach is demonstrated using the Lake Como basin as a case study. Lake Como is a regulated lake in northern Italy mainly operated for flood control and irrigation supply (Giuliani et al., 2016b).

2. PROBLEM FORMULATION

A generic water reservoir system can be described by an integrated dynamic model composed of the three main sub-systems illustrated in Figure 1, namely a hydroclimatic system, a reservoir system, and a downstream system. Each of the three systems has its own state \mathbf{x}_t^j ($j = H, R, D$) and the outputs of one system \mathbf{y}_{t+1}^j generally represent inputs to another system.

The dynamics of the composite system over time is affected by both control decisions \mathbf{u}_t and external disturbances $\boldsymbol{\varepsilon}_{t+1}$, i.e.

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \quad (1)$$

where the decisions at each time step are determined by a closed loop control rule $\mathbf{u}_t = \mu_t(\mathbf{x}_t)$, with the state of the system demonstrated to be sufficient for conditioning optimal operational decisions (Bertsekas, 1976). In the adopted notation, the time subscript of a variable indicates the instant when its value is deterministically known. An operating policy p is defined as a periodic sequence of control laws, $p \triangleq [\mu_0(\mathbf{x}_t), \dots, \mu_{T-1}(\mathbf{x}_{T-1})]$, where T is the period of the system (e.g., annual period).

The reservoir system (solid block in Figure 1) is the one directly controlled by the control decisions (i.e., the volume of water to be released from the dam) and that indirectly influences the downstream system (dotted block in Figure 1) by releasing water that supplies irrigation to the downstream agricultural district. Conversely, the hydroclimatic systems (dashed block in Figure 1) is uncontrolled and partially observable, but may still provide valuable information to the control rule.

For each of the Z stakeholders representing different water-related interests in the system (e.g., hydropower production, water supply, environmental protection), a specific objective function is formulated as a functional of the trajectory \mathcal{T} of relevant model variables over an evaluation horizon $[0, H]$ across an ensemble of K realizations of the system disturbances $\boldsymbol{\varepsilon}_{t+1}$

$$J^i = \Psi_{\boldsymbol{\varepsilon}^1, \dots, \boldsymbol{\varepsilon}^K} [\Phi_{t=1, \dots, h}(\mathcal{T})] \quad i = 1, \dots, Z \quad (2)$$

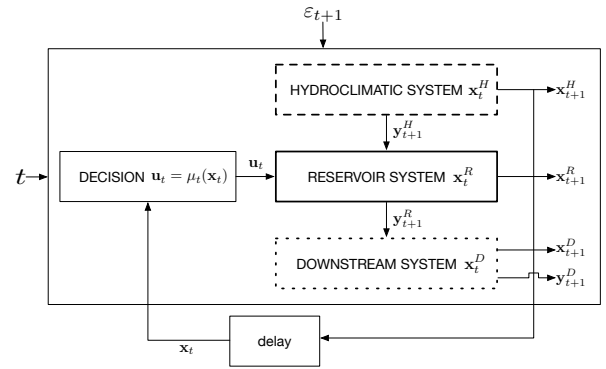


Fig. 1. Integrated dynamic model of a generic water system, composed of a hydroclimatic, reservoir, and downstream sub-systems.

where Φ_t is an operator for time aggregation (e.g., the cumulative cost or more complex cost functions that are not separable in time, such as duration curves), and Ψ_ε is a statistic used to filter the noise generated by the disturbances (e.g., expected value).

Given the dynamic model of the system in (1) and the objective functions in (2), the set of Pareto optimal policies p^* is designed by solving the following multi-objective optimal control problem (Castelletti et al., 2008):

$$p^* = \arg \min_p \mathbf{J}(\mathcal{T}) = |J^1(\mathcal{T}), \dots, J^Z(\mathcal{T})| \quad (3)$$

A policy is defined as Pareto-optimal if no other solution gives a better value for one objective without degrading the performance in at least one other objective.

3. METHODS

3.1 Climate system modeling

The model of the climate system focuses on the identification of Sea Surface Temperature (SST) anomalies that influence the current and future local hydroclimatic conditions according to global climate oscillations. The detection of these climate teleconnections is performed by means of the Niño Index Phase Analysis, a statistical framework originally developed by Zimmerman et al. (2016) for predicting seasonal precipitation conditioned on prior season atmospheric-oceanic variables.

The first step of the climate system modeling procedure generates correlation maps between seasonal SST anomalies Ω_t from a global monthly gridded dataset (NOAA's Extended Reconstructed SST Version 3b) and the next season mean local precipitation $P_{t+\tau}$. For each month, correlated grids ω_t at the 95% significance level are identified separately for the positive and negative phase of the considered climate signal (i.e., ENSO and NAO). The N selected SST anomalies therefore represents an observable sub-set of the hydroclimatic state (\mathbf{x}_t^H), i.e. $\omega_t \subseteq \Omega_t \subset \mathbf{x}_t^H$.

The second step of our procedure builds an emulator of the identified SST regions via Proper Orthogonal Decomposition. A Principal Component Analysis (PCA) is conducted on ω_t to search for linear combinations of these original

variables such that the coefficients of the output combinations (the principal vectors) form a low-dimensional subspace defined by directions explaining maximal variance in the original data (Jolliffe, 2002). PCA is performed via eigenvalue decomposition of the covariance matrix Q of the selected SST grids, i.e.

$$Q = \frac{1}{N_y} \sum_{j=1}^{N_y} (\omega_j)' \cdot \omega_j \quad (4)$$

where N_y is the number of years of available observations. The resulting eigenvectors ν^q (with $q = 1, \dots, N$) allow mapping the original SST grids into their principal components, i.e.

$$\varrho_t^q = \omega_t \cdot \nu^q \quad \text{with } q = 1, \dots, N \quad (5)$$

As in Zimmerman et al. (2016), only the first Principal Component (ϱ_t^1) for each climate signal is then considered, as it generally explains most of the variance in the selected SST predictor regions.

3.2 Evolutionary Multi-Objective Direct Policy Search

EMODPS is an approximate Dynamic Programming approach that relies on a functional optimization searching the optimal control policy within an infinite-dimensional space of functions (Baglietto et al., 2006). Specifically, EMODPS is based on the parameterization of the operating policies p_θ within a given family of functions and the exploration of the parameter space Θ to find a parameterized policy that optimizes the objective functions vector (Rückstieß et al., 2010), i.e.

$$p_\theta^* = \arg \min_{p_\theta} \mathbf{J}(\mathcal{T}) \quad (6)$$

subject to the dynamics of the system, where $\theta \in \Theta$.

Numerous policy search approaches have been developed in recent years for addressing a diverse suite of control problems (for a review, see Deisenroth et al., 2011, and references therein), with the adopted EMODPS method that can be classified as a stochastic, model-based, episode-based, multi-objective DPS approach. Since policy search methods do not provide any theoretical guarantee on the optimality of the resulting solutions, the choice of the parametric function for the control policy and the ability of the optimization algorithm used to search for the optimal control policy parameters emerged as two critical factors determining the ability of discovering high-quality solutions.

In this work, the control policies are defined as Gaussian radial basis functions (RBF, Busoniu et al., 2011), which have been demonstrated to be effective in solving these types of multi-objective control problems (Giuliani et al., 2016a). The k -th control decision in \mathbf{u}_t (with $k = 1, \dots, n_u$) is defined as:

$$u_t^k = \sum_{i=1}^B w_{i,k} \varphi_i(\mathcal{I}_t) + \alpha \quad (7)$$

where B is the number of RBFs $\varphi(\cdot)$ and $w_{i,k}$ the non-negative weight of the i -th RBF (i.e., $w_i \geq 0 \quad \forall i$), and α a constant parameter. The single RBF is defined as follows:

$$\varphi_i(\mathcal{I}_t) = \exp \left[- \sum_{j=1}^M \frac{((\mathcal{I}_t)_j - c_{j,i})^2}{b_{j,i}^2} \right] \quad (8)$$

where M is the number of input variables \mathcal{I}_t and $\mathbf{c}_i, \mathbf{b}_i$ are the M -dimensional center and radius vectors of the i -th RBF, respectively. The policy parameter vector θ is therefore defined as

$$\theta = [(c_1, \dots, c_M)_1^B, (b_1, \dots, b_M)_1^B, (w_1, \dots, w_{n_u})_1^B, \alpha]$$

and belongs to \mathbb{R}^{n_θ} , where $n_\theta = B(2M + n_u) + n_u$.

To cope with the high-dimensional space of policy parameters as well as nonlinear, noisy, and multimodal objective functions, we solved the EMODPS problem in (6) by using the self-adaptive Borg Multi-Objective Evolutionary Algorithm (Hadka and Reed, 2013). This latter employs multiple global probabilistic search operators for mating, selection, and mutation, and adjusts their probability of being selected during the search based on their demonstrated ability of generating new nondominated solutions. These two features contribute in making the Borg MOEA highly robust in solving multi-objective optimal control problems (Zatarain-Salazar et al., 2016), as they overcome the limitations of tuning the algorithm's parameters to the specific fitness landscape of the problem.

A key feature of EMODPS is the possibility of conditioning the policy through a data-driven controller tuning approach (Formentin et al., 2013), where the control decisions can be conditioned upon (partial) observations of the hydroclimatic system state (Giuliani et al., 2015, 2018). In this work, we consider two different EMODPS problem formulations. A *baseline* formulation in (9a), where the vector of control decisions \mathbf{u}_t is determined by a policy defined as a function of the day of the year d_t (to account for the time-dependency and cyclostationarity of the system and, consequently, of the control policy) and the reservoir storage x_t^R ; and a *climate-informed* formulation in (9b) that includes as additional policy inputs the emulated state of the hydroclimatic system represented by the first principal components of the selected SST grids for ENSO and NAO.

$$\mathbf{u}_t = p_\theta(d_t, x_t^R) \quad (9a)$$

$$\mathbf{u}_t = p_\theta(d_t, x_t^R, \varrho_t^{1, ENSO}, \varrho_t^{1, NAO}) \quad (9b)$$

4. CASE STUDY DESCRIPTION

Lake Como is a sub-alpine lake in the Italian lake district, northern Italy (Figure 2). The main tributary, and only emissary of the lake, is the Adda river, whose sublacual part originates in the southeastern branch of the lake and feeds four agricultural districts. The southwestern branch of the lake constitutes a dead end exposed to flooding events, particularly in the city of Como which is the lowest point of the lake shoreline. The lake regulation has the dual aim of guaranteeing flood protection to the lake shores and supplying water to downstream users.

Lake Como dynamics is modeled as a discrete-time, periodic, non-linear, stochastic MDP with the following features: the state variable x_t^R is the lake storage and the control u_t , determined by the control policy, is the release decision; the system is affected by a stochastic disturbance ε_{t+1}^R representing the lake's net inflow in the time interval $[t, t + 1)$; the state transition function $x_{t+1}^R = x_t^R +$




Fig. 2. Map of the Lake Como basin.

$\varepsilon_{t+1}^R - r_{t+1}$ describes the daily storage dynamics, with the actual released volume r_{t+1} is determined by a nonlinear, stochastic function of the control decision (Soncini-Sessa et al., 2007).

On the basis of previous works (e.g., Giuliani and Castelletti, 2016; Denaro et al., 2017) two operating objectives (both to be minimized) are formulated as follows: *i) flooding* J^F , defined in (10) as the average annual number of flooding days over the simulation horizon H with a lake level h_t higher than the flooding threshold $\bar{h} = 1.24$ m; *ii) water supply* J^S , defined in (11) as the daily average quadratic water deficit between the lake release r_{t+1} and the daily water demand of the downstream system w_t , subject to the minimum environmental flow constraint $q^{MEF} = 5$ m³/s, where a time-varying parameter β_t is used to penalize deficits occurring after the crops' germination until the beginning of phenological maturity.

$$J^F = \frac{1}{H/365} \sum_{t=1}^H \Lambda(h_t) \quad \text{where} \quad (10)$$

$$\Lambda(h_t) = \begin{cases} 1 & \text{if } h_t > \bar{h} \\ 0 & \text{otherwise} \end{cases}$$

$$J^S = \frac{1}{H} \sum_{t=1}^H \beta_t \times \max(w_t - \max(r_{t+1} - q^{MEF}, 0), 0)^2 \quad (11)$$

The EMODPS optimization is performed over the simulation horizon 1996-2008, which was selected because it shows good variability in the local hydrological conditions including some intense droughts events. To improve solution diversity and avoid dependence on randomness, the solution set from each formulation is the result of 20 random optimization trials. The final set of Pareto optimal policies for each experiment is defined as the set of non-dominated solutions from the results of all the optimization trials.

According to the Information Selection and Assessment framework (Giuliani et al., 2015), the improvement gen-

erated by the climate-informed formulation with respect to the baseline is quantified using as a reference a set of upper-bound solutions, designed assuming perfect foresight of future inflows. In particular, we used a set of control policies informed by a perfect forecast of the Lake Como inflow accumulated over a lead-time of 51 days, which was demonstrated by Denaro et al. (2017) to be a valuable information for improving the Lake Como operations.

Given the multi-objective nature of the Lake Como problem whose solution is a set of Pareto optimal (or approximate) policies, the benefit of adopting the climate-informed formulation in (9b) with respect to the baseline formulation in (9a) is measured in terms of hypervolume indicator (HV), a metric that captures both the proximity of the approximated set \mathcal{F} to the optimal front \mathcal{F}' and the distribution of the solutions for representing the full extent of tradeoffs. This indicator is formulated as the ratio of objective space's volumes dominated (\preceq) by the policy under evaluation (i.e., baseline or climate-informed) and the reference set obtained under the assumption of perfect foresight:

$$HV(\mathcal{F}, \mathcal{F}') = \frac{\int \alpha_{\mathcal{F}}(\mathbf{s}) d\mathbf{s}}{\int \alpha_{\mathcal{F}'}(\mathbf{s}) d\mathbf{s}} \quad (12a)$$

$$\alpha(\mathbf{s}) = \begin{cases} 1 & \text{if } \exists s' \in \mathcal{F} \text{ such that } s' \preceq \mathbf{s} \\ 0 & \text{otherwise} \end{cases} \quad (12b)$$

5. RESULTS

Figure 3 illustrates the performance of three Pareto approximate sets, namely the baseline (white circles), climate-informed (gray circles), and upper-bound control policies (black squares), over the time horizon (1996-2008). The two axes of the figure represent the operating objectives (to be minimized) and the black arrows show the direction of increasing preference, with the best solution located in the bottom-left corner of the figure.

Results provide a clear ranking of the three set of solutions, with the space between the baseline and upper-bound policies that quantify the potential room for improvement generated by the ideal perfect information of the future inflows. Inflows, nevertheless, depend on the hydroclimatic system dynamics (Figure 1) and the climate-informed policies completely dominate the baseline solutions by partially filling the gap with respect to the upper-bound policies. Yet, using an emulator of the SST dynamics provides a partial representation of the hydroclimatic system, yielding inferior solutions with respect to the upper-bound policies. It is worth noticing that both the climate-informed and the upper-bound policies improve mostly in terms of J^S . The seasonal information provided by both the pre-season SST (climate-informed policies) and 51 day ahead perfect inflow forecast (upper-bound policies) is indeed relevant for water supply operations, while flood control would need different information on a shorter lead time to capture the flood dynamics, which generally evolve over some hours to few days. Still, the downward shift of the Pareto front indirectly influence the performance in terms of J^F by allowing the identification of better compromise solutions.




Fig. 3. Performance obtained by the baseline, climate-informed and upper-bound control policies, with the arrows indicating the direction of increasing preference for the two objectives. Solutions A and B are selected as they attain the same performance of the historical lake regulation in terms of flood control.

A quantitative measure of this improvement is provided by the values of the hypervolume indicator in Table 1. Moving from the baseline solutions to the climate-informed policies generates an increase of HV equal to 0.19, corresponding to a 59% gain in system performance.

Table 1. Hypervolume indicator (HV) for the Pareto approximate sets in Figure 3.

Policies	HV	ΔHV
Baseline policies	0.32	-
Climate-informed policies	0.51	0.19
Upper-bound policies	1.00	0.68

To better understand the differences between the baseline and climate-informed policies, Figure 4 reports the simulated lake level under two distinct solutions (i.e., policies A and B in Figure 3) that attain the same performance of the historical regulation of the lake in terms of flood control. The rationale of this choice is to look at control policies able to improve J^S without degrading the performance in J^F . In terms of water supply deficit, policy B attains a reduction of J^S equal to 287 (m^3/s^2); this corresponds to an average deficit reduction of 1.42 m^3/s (i.e., $44.8 \times 10^6 \text{ m}^3/\text{year}$). All the trajectories show a common annual pattern, with the highest lake levels reached in late spring thanks to the contribution of the snowmelt, followed by a marked drawdown during the summer, when the downstream irrigation demand is at its maximum. The summer season is the period when the climate-informed policy (red lines) gain the most over the baseline one (blue lines) by delaying the drawdown to late July in order to save water for covering the irrigation demand in August. This difference is particularly evident during intense drought episodes, with a few blue lines that fall below 0.25 m in June, one month before the corresponding red trajectories.

6. CONCLUSION

The paper explores the potential benefit of including information from the hydroclimatic system as additional




Fig. 4. Yearly patterns of the Lake Como level over the simulation horizon (1996-2008) under the policies A (baseline) and B (climate-informed) marked in Figure 3.

input of the control policies used for the optimal operations of water reservoir systems. The multipurpose operations of Lake Como (Italy) is used as a case study.

The combination of a reduced model of the climate system capturing the dynamics of Sea Surface Temperature anomalies with a data-driven controller designed via Evolutionary Multi-Objective Direct Policy Search produces a 59% improvement in system performance. This improvement can be attributed to the information capturing global climate oscillations provided by the SST model emulator, which is influencing the Lake Como inflows especially during extreme drought events.

Future research efforts will focus on further improving the representation of the hydroclimatic system beside the SST anomalies for capturing additional processes that might be valuable for informing the control policy. Moreover, it will be interesting to generalize the proposed approach in different case studies featuring diverse hydroclimatic conditions and operating objectives.

REFERENCES

- Baglietto, M., Cervellera, C., Sanguineti, M., and Zoppoli, R. (2006). Water reservoirs management under uncertainty by approximating networks and learning from data. In *Topics on System Analysis and Integrated Water Resource Management*, 117–139. Elsevier Amsterdam, NL.
- Bellman, R. (1957). *Dynamic programming*. Princeton University Press, Princeton.
- Bertsekas, D. (1976). *Dynamic programming and stochastic control*. Academic Press, New York.
- Busoniu, L., Ernst, D., De Schutter, B., and Babuska, R. (2011). Cross-Entropy Optimization of Control Policies With Adaptive Basis Functions. *IEEE Transactions on systems, man and cybernetics–Part B: cybernetics*, 41(1), 196–209. doi:10.1109/TSMCB.2010.2050586.
- Castelletti, A., Pianosi, F., and Soncini-Sessa, R. (2008). Water reservoir control under economic, social and environmental constraints. *Automatica*, 44(6), 1595–1607. doi:10.1016/j.automatica.2008.03.003.

- Chatterjee, A. (2000). An introduction to the proper orthogonal decomposition. *Current science*, 808–817.
- Deisenroth, M., Neumann, G., and Peters, J. (2011). A Survey on Policy Search for Robotics. In *Foundations and Trends in Robotics*, volume 2, 1–142.
- Denaro, S., Anghileri, D., Giuliani, M., and Castelletti, A. (2017). Informing the operations of water reservoirs over multiple temporal scales by direct use of hydro-meteorological data. *Advances in Water Resources*, 103, 51–63. doi:10.1016/j.advwatres.2017.02.012.
- Emerton, R., Cloke, H., Stephens, E., Zsoter, E., Woolnough, S., and Pappenberger, F. (2017). Complex picture for likelihood of ENSO-driven flood hazard. *Nature communications*, 8, 14796.
- Ficchi, A., Raso, L., Dorchie, D., Pianosi, F., Malaterre, P., Van Overloop, P., and Jay-Allemand, M. (2016). Optimal operation of the multireservoir system in the seine river basin using deterministic and ensemble forecasts. *Journal of Water Resources Planning and Management*, 142(1), 05015005.
- Formentin, S., Karimi, A., and Savaresi, S. (2013). Optimal input design for direct data-driven tuning of model-reference controllers. *Automatica*, 49(6), 1874–1882.
- Galelli, S., Castelletti, A., and Goedbloed, A. (2015). High-Performance Integrated Control of water quality and quantity in urban water reservoirs. *Water Resources Research*, 51(11), 9053–9072.
- Giuliani, M. and Castelletti, A. (2016). Is robustness really robust? How different definitions of robustness impact decision-making under climate change. *Climatic Change*, 135, 409–424. doi:10.1007/s10584-015-1586-9.
- Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., and Reed, P. (2016a). Curses, tradeoffs, and scalable management: advancing evolutionary multi-objective direct policy search to improve water reservoir operations. *Journal of Water Resources Planning and Management*, 142(2). doi:10.1061/(ASCE)WR.1943-5452.0000570.
- Giuliani, M., Galelli, S., and Soncini-Sessa, R. (2014). A dimensionality reduction approach for Many-Objective Markov Decision Processes: application to a water reservoir operation problem. *Environmental Modeling & Software*, 57, 101–114. doi:10.1016/j.envsoft.2014.02.011.
- Giuliani, M., Li, Y., Castelletti, A., and Gandolfi, C. (2016b). A coupled human-natural systems analysis of irrigated agriculture under changing climate. *Water Resources Research*, 52(9), 6928–6947.
- Giuliani, M., Pianosi, F., and Castelletti, A. (2015). Making the most of data: an information selection and assessment framework to improve water systems operations. *Water Resources Research*, 51(11), 9073–9093. doi:10.1002/2015WR017044.
- Giuliani, M., Quinn, J.D., Herman, J.D., Castelletti, A., and Reed, P.M. (2018). Scalable multiobjective control for large-scale water resources systems under uncertainty. *IEEE Transactions on Control Systems Technology*, 26(4), 1492–1499. doi:10.1109/TCST.2017.2705162.
- Giuliani, M., Zaniolo, M., Castelletti, A., Davoli, G., and Block, P. (2019). Detecting the state of the climate system via artificial intelligence to improve seasonal forecasts and inform reservoir operations. *Water Resources Research*, 55(11), 9133–9147. doi:10.1029/2019WR025035.
- Hadka, D. and Reed, P. (2013). Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework. *Evolutionary Computation*, 21(2), 231–259.
- Hall, W. and Buras, N. (1961). The dynamic programming approach to water-resources development. *Journal of Geophysical Research*, 66(2), 517–520.
- Jolliffe, I. (2002). *Principal Component Analysis*. Springer, New York, N.Y.
- Labadie, J. (2004). Optimal operation of multireservoir systems: State-of-the-art review. *Journal of Water Resources Planning and Management*, 130(2), 93–111.
- Lee, D., Ward, P., and Block, P. (2018). Identification of symmetric and asymmetric responses in seasonal streamflow globally to ENSO phase. *Environmental Research Letters*, 13(4), 044031.
- Maass, A., Hufschmidt, M., Dorfman, R., Thomas Jr, H., Marglin, S., and Fair, G. (1962). *Design of water-resource systems*. Harvard University Press Cambridge, Mass.
- Powell, W. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*. Wiley, NJ.
- Raso, L., Schwanenberg, D., van de Giesen, N., and van Overloop, P. (2014). Short-term optimal operation of water systems using ensemble forecasts. *Advances in water resources*, 71, 200–208.
- Rippl, W. (1883). The capacity of storage reservoirs for water supply. In *Minutes of the Proceedings, Institution of Civil Engineers*, volume 71, 270–278. Thomas Telford.
- Rückstiegs, T., Sehnke, F., Schaul, T., Wierstra, D., Sun, Y., and Schmidhuber, J. (2010). Exploring parameter space in reinforcement learning. *Paladyn, Journal of Behavioral Robotics*, 1(1), 14–24.
- Soncini-Sessa, R., Castelletti, A., and Weber, E. (2007). *Integrated and participatory water resources management: Theory*. Elsevier, Amsterdam, NL.
- Tian, X., Guo, Y., Negenborn, R.R., Wei, L., Lin, N.M., and Maestre, J.M. (2019). Multi-scenario model predictive control based on genetic algorithms for level regulation of open water systems under ensemble forecasts. *Water Resources Management*, 33(9), 3025–3040.
- Tsitsiklis, J. and Van Roy, B. (1996). Feature-based methods for large scale dynamic programming. *Machine Learning*, 22(1-3), 59–94.
- Yeh, W. (1985). Reservoir management and operations models: a state of the art review. *Water Resources Research*, 21 (12), 1797–1818.
- Zatarain-Salazar, J., Reed, P., Herman, J., Giuliani, M., and Castelletti, A. (2016). A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances in Water Resources*, 92, 172–185. doi:10.1016/j.advwatres.2016.04.006.
- Zimmerman, B.G., Vimont, D.J., and Block, P.J. (2016). Utilizing the state of enso as a means for season-ahead predictor selection. *Water Resources Research*, 52(5), 3761–3774.