

Water Resources Research



REVIEW ARTICLE

10.1029/2021WR029927

A State-of-the-Art Review of Optimal Reservoir Control for Managing Conflicting Demands in a Changing World

M. Giuliani¹ , J. R. Lamontagne² , P. M. Reed³ , and A. Castelletti¹ 

¹Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milano, Italy, ²Department of Civil and Environmental Engineering, Tufts University, Medford, MA, USA, ³School of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA

Key Points:

- We review the evolution of water reservoir operation studies driven by emerging societal challenges
- Simulation-based control methods facilitate realistic problem formulations that are able to handle system uncertainties
- Research opportunities remain to combine policy design with more information sources and emulation modeling

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

A. Castelletti,
andrea.castelletti@polimi.it

Citation:

Giuliani, M., Lamontagne, J. R., Reed, P. M., & Castelletti, A. (2021). A state-of-the-art review of optimal reservoir control for managing conflicting demands in a changing world. *Water Resources Research*, 57, e2021WR029927. <https://doi.org/10.1029/2021WR029927>

Received 6 MAR 2021

Accepted 7 NOV 2021

Abstract The state of the art for optimal water reservoir operations is rapidly evolving, driven by emerging societal challenges. Changing values for balancing environmental resources, multisectoral human system pressures, and more frequent climate extremes are increasing the complexity of operational decision making. Today, reservoir operations benefit from technological advances, including improved monitoring and forecasting systems as well as increasing computational power. Past research in this area has largely focused on improving solution algorithms within the limits of the available computational power, using simplified problem formulations that can misrepresent important systemic complexities and intersectoral interactions. In this study, we review the recent literature focusing on how the operation design problem is formulated, rather than solved, to address existing challenges and take advantage of new opportunities. This paper contributes a comprehensive classification of over 300 studies published over the last years into distinctive categories depending on the adopted problem formulation, which clarifies consolidated methodological approaches and emerging trends. Our analysis also suggests that control policy design methods may benefit from broadening the types of information that is used to condition operational decisions, and from using emulation modeling to identify low-order, computationally efficient surrogate models capturing realistic representations of river basin systems' complexity in order to isolate key decision-relevant processes. These advances in reservoir operations hold significant promise for better addressing the challenges of conflicting human pressures and a changing world, which is particularly important, given the renewed interest in dam construction globally.

1. Introduction

Reservoir operations are increasingly important in the water cycle (Hanasaki et al., 2006; Padowski et al., 2015; Wada et al., 2017) and support regional growth and development by increasing water availability for various economic sectors, contributing renewable electricity production, and reducing flood risks (Billington & Jackson, 2017). Recent estimates (Grill et al., 2019) suggest that existing dams control around 50% of the all rivers globally. This figure is expected to grow rapidly following the renewed interest in dam construction (World Bank, 2009), which poses the challenge of designing their future operations (e.g., Bertoni et al., 2019; Geressu & Harou, 2015; Mortazavi-Naeini et al., 2014) to secure water and energy supplies in fast developing African and Asian countries (Zarfl et al., 2015). Rapid changes in climate and society suggest an urgent need to reoperate existing infrastructures (Benson, 2016), particularly in systems that are failing to produce the expected benefits that motivated their construction (Ansar et al., 2014; Sovacool et al., 2014).

The problem of designing optimal reservoir operations has been extensively studied since the seminal works by Rippl (1883), Hazen (1914), and Varlet (1923). Over the years, several review papers (e.g., Dobson et al., 2019; Klemesš, 1987; Labadie, 2004; Macian-Sorribes & Pulido-Velazquez, 2020; Simonovic, 1992; Stedinger et al., 2013; Wurbs, 1993; Yakowitz, 1982; Yeh, 1985) have described the evolving state of the art in the field. Changes in societal perceptions of natural resources and increasing environmental awareness are modifying and enlarging the number of objectives considered (e.g., Schmitt et al., 2018; Wild et al., 2019; Winemiller et al., 2016). In addition, changing climate extremes and societal demands amplify and reshape uncertain stressors, ultimately altering decision makers' preferences and risk perception (e.g., AghaKouchak, 2015; Hall et al., 2014; Mallakpour et al., 2019; Mora et al., 2018; Poff et al., 2015). Despite these new demands, today we also have numerous opportunities for improving reservoir operations, including more information from pervasive monitoring and remote sensing (e.g., Hart & Martinez, 2006; Lettenmaier et al., 2015), more skillful forecasts (e.g., Bauer et al., 2015; Brunet et al., 2015; Cloke et al., 2017), and more computing power (e.g., Kollet et al., 2010; National Research

© 2021. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Council, 2005; Reed & Hadka, 2014). In this review, we contribute an updated analysis of emerging methodological trends with a focus on optimal control methods, which support the design of closed-loop policies that use feedback loops to compensate for system uncertainties (Bennett, 1996). Notably, we discuss how the operation design problem is formulated, rather than solved, in order to highlight how recently developed approaches are evolving to address global challenges, take advantage of recent innovations, and offer new opportunities.

For decades, analysts have had to frame reservoir operation problems in simplified ways to comply with the constraints of specific solution techniques and computational limits. These simplifications include approximated models that decompose the reservoir system into smaller and simpler subsystems or aggregate the reservoir system, or part thereof, into a single equivalent reservoir to derive composite solutions to be then disaggregated into individual policies (e.g., Nandalal & Bogardi, 2007; Pereira & Pinto, 1985; Saad et al., 1992, 1994). Simplified models are also adjusted to meet the mathematical requirements of specific solution methods, such as Linear Programming (e.g., Reville et al., 1968; Rheinheimer et al., 2014; Zhang et al., 2016; Zmijewski et al., 2016), Linear Quadratic Gaussian control (i.e., a linear system with quadratic costs and Gaussian disturbances; e.g., Georgakakos, 1989; Ghorbanidehno et al., 2017), or analytical optimization methods (e.g., Clark, 1950; Ding et al., 2017; Hui et al., 2016; Hui & Lund, 2015; Johnson et al., 1991; Lund, 2000; Lund & Guzman, 1999; Zhao et al., 2019). Although analytical approaches show promise for gaining useful insights to derive reservoir operating rules for hypothetical water systems, as noted by Labadie (2005), these methods are primarily relegated to academic exercises because the complexity and highly regulated nature of real-world systems strongly limit their practical application. Besides, the simplified access to increasing computational power (Rupp, 2020) has now enabled the adoption of more and more realistic and decision-relevant problem formulations.

Another common simplification is the definition of search space in terms of either a “rule curve” specifying a reference trajectory for the reservoir storage under normal hydroclimatic conditions (e.g., Ahn & Lyu, 2017; Liu et al., 2015; Whateley et al., 2014; Zhang, Li, et al., 2017) or an open-loop sequence of release decisions, corresponding to a sequence of decision variables one for each time step under the strong assumption implicit to a deterministic problem formulation where future inflows are known at the moment when operational decisions must be made (e.g., Bahrami et al., 2017; Bozorg-Haddad et al., 2015; Li et al., 2015; Schardong & Simonovic, 2015; Yang et al., 2015). It is, however, important to consider that rule curves, despite being the preferred tool by reservoir operators (Loucks & Sigvaldason, 1982; Teegavarapu & Simonovic, 2001), are inherently limited when the system deviates from the typical conditions assumed in their design (Howard, 1999), as they would suggest the operator to spill water when the storage exceeds the level specified by the rule, and to hope for rain when it falls below (Maass et al., 1962). Under increasingly variable hydrologic regimes (e.g., Dai, 2011; Trenberth, 2011), these deviations from the typical conditions are larger and more frequent, amplifying errors in the open-loop decisions determined by the rule curves. Moreover, the lack of feedback information generates rules that largely ignore the potential for adaptive and coordinated strategies, especially in multireservoir systems managing multiple conflicting objectives (Quinn et al., 2019). An open-loop sequence of decisions provides optimal releases only under the strong assumption of a deterministic problem formulation, which implies that the resulting solutions are impossible to implement in the real world (Faber & Stedinger, 2001), and can only be considered in benchmarking exercises as a reference performance that is unattainable without perfect knowledge of the future (Koutsoyiannis & Economou, 2003b). In addition, traditional open-loop intertemporal optimizations over long time horizons yield problematic large decision spaces. Finally, the weak system responses to late period decisions strongly limit the possibility of generating high-quality solutions, unless the system is simplified as a linear model (Quinn, Reed, & Keller, 2017; Ward et al., 2015).

Stochastic Dynamic Programming (SDP) has been considered, since the '60s (Hall & Buras, 1961; Maass et al., 1962), the best method to preserve a realistic problem structure due to mild requirements on systems representation (Esogbue, 1989), favoring its adoption in practical applications (e.g., Côté & Leconte, 2015; Dias et al., 1985; Kelman et al., 1988) and its systematic use in the reservoir operation literature (e.g., Anghileri et al., 2013; Cheng et al., 2017; Davidsen et al., 2014; Delipetrev et al., 2015; Mabaya et al., 2017; Turner & Galelli, 2016). Yet, SDP has several limitations that constrain problem framing in addressing emerging challenges in water system operations: (a) the well-known curse of dimensionality (Bellman, 1957) 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; (b) the curse of modeling (Tsitsiklis & Van Roy, 1996) requires all variables used as input in the operating policy to be described by a dynamic model, contributing additional state variables; (c) the curse of

multiple objectives (Powell, 2007) 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, Galelli, & Soncini-Sessa, 2014; Reed & Kollat, 2013).

Growing human pressure along with evolving natural and socioeconomic conditions are making these three curses particularly restrictive. The complexity of a real water system, which is generally highly nonlinear, spatially heterogeneous, and comprises several processes with multiscale time dynamics such as flash floods (e.g., Hapuarachchi et al., 2011) and multiannual droughts (e.g., Howitt et al., 2017), poses a major challenge for the identification of reliable simulation models relying on a limited number of state variables. The increased variability and uncertainty of climatic and socioeconomic drivers motivate the need for more anticipatory capacity, which can be provided by broadening the array of information sources that can be used for conditioning reservoir operations with additional observational or forecast data (e.g., Giuliani et al., 2015). The transition from multiobjective problem formulations including two or three objectives (e.g., irrigation supply, flood protection, and hydropower production) to many-objective problems, where the number of objectives is four or more (Fleming et al., 2005), is accelerating given emerging issues related to environmental flow regulations, water quality targets, recreational interests, and energy markets. Many objective formulations also allow for rigorous evaluations of multiple rival framing for component objectives, such as minimizing average costs versus minimizing worst-case costs, thus addressing the issue of epistemic uncertainty in the formulations of river basin management problems (Quinn, Reed, Giuliani, & Castelletti, 2017). Moreover, there is a growing recognition of the need for reservoir operations to account for a broader array of uncertainties and their projected long-term changes, particularly in terms of extreme events. As the breadth of uncertainties that must be considered grows, there is a higher demand to explore multiple objectives capturing potentially evolving and heterogeneous stakeholder preferences and attitudes toward risk and uncertainty (McPhail et al., 2018). Lastly, the growing interest in implementing participatory approaches to water resources planning and management (e.g., Soncini-Sessa, Cellina, et al., 2007) discourages using mathematically binding problem formulations (e.g., linear or time-separable objectives) that may not always allow the representation of the interests elicited from stakeholders and policy makers (i.e., the freedom to formulate more credible and relevant versions of problems), for example, when they are interested in the duration of a low-flow or high-flow event as required in the Indices of Hydrologic Alteration (Richter et al., 1996).

Over the last years, researchers have developed numerous attempts to support real-world applications by overcoming SDP's curses and reducing simplifications to reservoir control problems' formulation, which pose the risk of misrepresenting important systemic properties, dynamics, and intersectoral interactions. Defining a problem formulation that (a) is able to better reproduce the relevant complexities of water system dynamics and variability in space and time, (b) ensures a realistic representation of decisions and objectives for discovering key trade-offs between multisectoral demands, and (c) can capture a range of heterogeneous stakeholder attitudes toward risk and uncertainty is paramount to the design of decision-relevant alternatives as well as to the selection of the most appropriate solution method. We argue these challenges are critical for reservoir operation design problems, while acknowledging real-time operational contexts face other difficulties in implementing short-term decisions (e.g., Cheng et al., 2020; Séguin et al., 2016). Optimal control methods reduce restrictions on the formulation of the reservoir operation problems supporting the design of closed-loop policies. A closed-loop policy (see Figure 2) determines the operational decisions at each time step on the basis of the observed conditions of the system (e.g., the reservoir storage), which depend on the decisions made in previous time steps and therefore introduce a recursive loop; an open-loop policy instead specifies operational decisions a priori for each time step, regardless of the conditions that might occur. Following the methodological classification proposed by Bertsekas (2019) and Powell (2019), these optimal control methods can be categorized into two main classes: (a) *approximation in value space* (AVS), which searches an approximation of the value function (Bertsekas, 2005); (b) *approximation in policy space* (APS), which first defines the operating policy within a restricted class of parameterized functions and then explores the policy parameter space to optimize the system performance (Deisenroth et al., 2011).

Our review contributes a systematic classification of 337 recent studies on reservoir operation design into six distinct categories depending on the adopted problem formulation (Figure 1): Mathematical Programming, Open-Loop decisions, Rule Curves, Dynamic Programming, Approximation in Value Space, and Approximation in Policy Space. These papers were obtained by searching for topic equal to "optimal water reservoir operations" in ISI Web of Knowledge (see Figure S1 in Supporting Information S1 for details on the technical principles

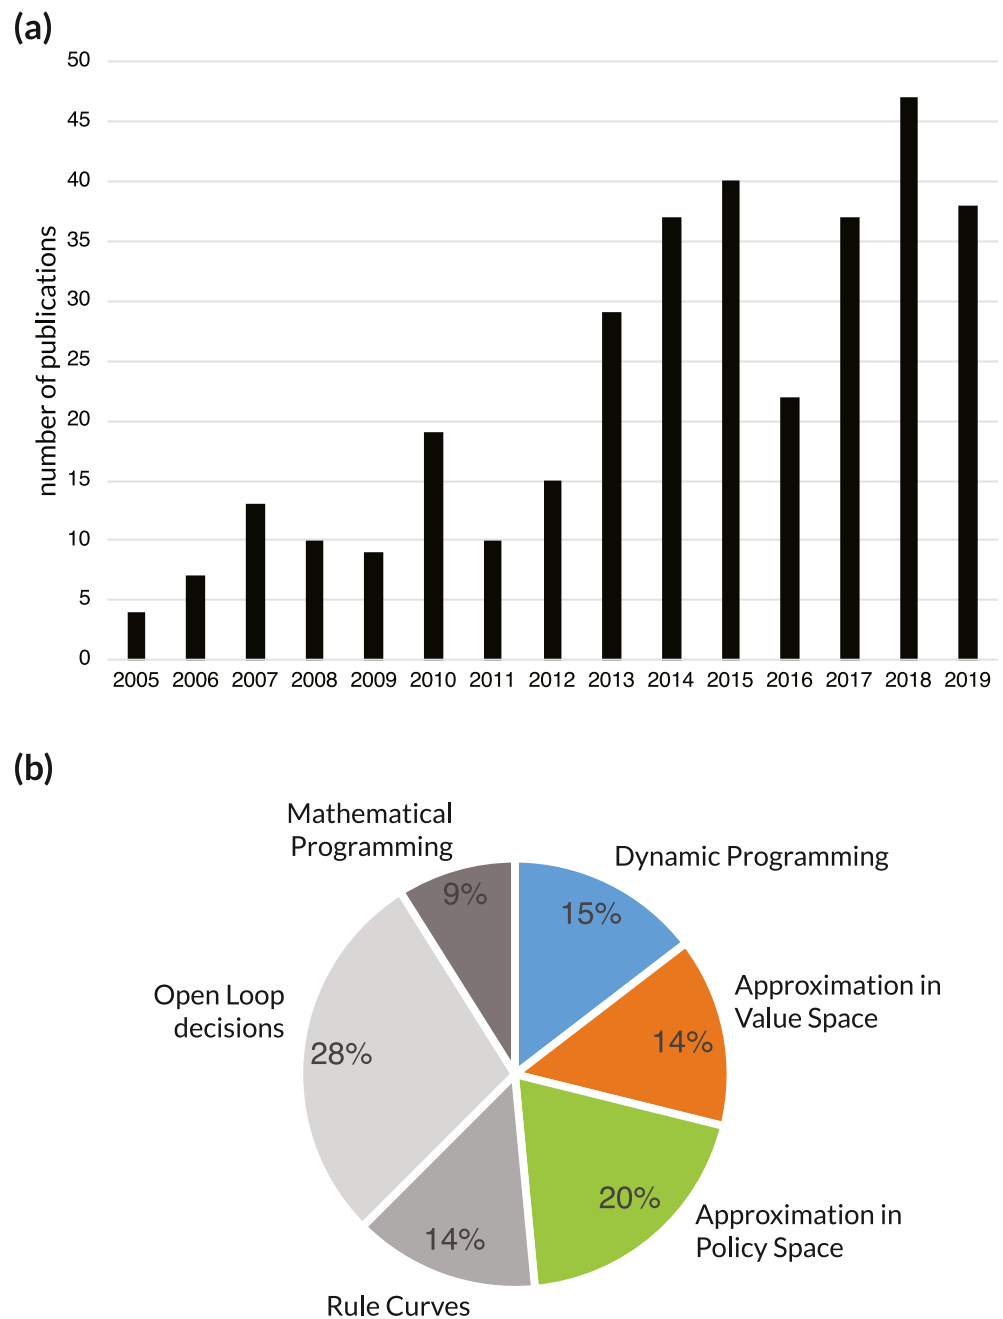


Figure 1. Annual counts of the 337 publications reviewed in this study (a) and their classification depending on the adopted problem formulation (b).

for systematic reviews and meta-analyses that have informed the literature search). In particular, we provide a critical analysis of the major emerging challenges for optimal reservoir control methods that support the design of closed-loop policies, as derived from the 115 publications that use either AVS or AVP methods. We do not discuss in detail the other methods (i.e., Mathematical Programming, Open-Loop decisions, and Rule Curves) that rely on simplified problem formulations and that have already been well covered in existing review papers (e.g., Labadie, 2004; Macian-Sorribes & Pulido-Velazquez, 2020).

The rest of the paper is structured as follows: The next section introduces a traditional formulation of the optimal reservoir operation problem, followed by the presentation of Stochastic Dynamic Programming. We then provide a critical analysis of the 115 studies, which attempt to go beyond SDP limitations by using either AVS

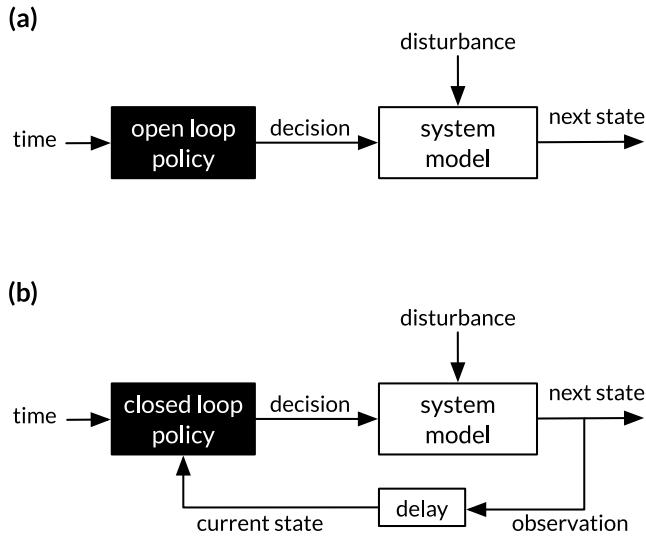


Figure 2. Illustration of open-loop (a) and closed-loop (b) policies.

or APS, and suggest a possible road map for future research directions. The last section reports final remarks by clarifying consolidated approaches and emerging trends from our analysis.

2. Traditional Problem Formulation

A generic model of a water reservoir system relies on a mass-balance equation describing the dynamics of the water stored in the reservoir:

$$s_{t+1} = s_t + q_{t+1} - r_{t+1}(s_t, u_t, q_{t+1}) \quad (1)$$

where s_t is the reservoir storage at time t , q_{t+1} and r_{t+1} are the net inflow (i.e., inflow and direct precipitation minus evaporation and seepage losses) and the actual release volume in the time interval $[t, t + 1)$, respectively, and u_t is the release decision. In the adopted notation, the time subscript of a variable indicates the time instant when its value is deterministically known. The reservoir storage is measured at time t and thus is denoted as s_t , while the net inflow and the actual release are denoted as q_{t+1} and r_{t+1} , respectively, because they can be known only at the end of the time interval. Note that in the case of multireservoir systems, all the variables in Equation 1 become vectors and the network topology can be represented by an incidence matrix.

The release volume r_{t+1} is determined by a nonlinear, stochastic function of the release decision u_t (Soncini-Sessa, Castelletti, & Weber, 2007), because between the time t at which the decision is taken and the time $t + 1$ at which the release is completed, the uncertain inflow is affecting the reservoir storage. Thus, the actual release may not be equal to the decision due to legal and physical constraints on reservoir level and release, including spills when the reservoir exceeds maximum capacity.

The net inflow q_{t+1} is often modeled as a system disturbance (i.e., $q_{t+1} = \varepsilon_{t+1}$), which aggregates multiple sources of uncertainty, such as main tributaries, distributed inflow runoff, evaporation, and precipitation over the reservoir. These processes can also be separately modeled as distinct disturbances. Alternatively, the reservoir inflow and the persistence of hydrologic processes can be described using dynamic, process-based, spatially distributed hydrological models. However, the resulting complexity of these models may lead to serious computational problems, favoring the adoption of simpler data-driven statistical models (e.g., Salas et al., 1985; Wang et al., 2009).

Even though the majority of these processes act in continuous time, their representations are often discretized to match the decision process made by human operators. According to the sampling theorem (Shannon, 1949), this discretization occurs without losing information only if the selected time step is smaller than the time constant of the linearized reservoir model by more than one order of magnitude. This implies that the duration of the time step should be short enough to timely adapt the decision to a variation of the system conditions (system representability) (Soncini-Sessa, Castelletti, & Weber, 2007). However, it should also be long enough to keep the decision unchanged and observe its impact on the system dynamics (social acceptability). For example, a monthly time step could be ineffective to design flood control reservoirs operations as floods generally have much faster dynamics requiring daily or hourly decisions. Conversely, a short time step may be beneficial for following fast inflow or electricity price dynamics, but at subhourly time steps, increasingly detailed hydraulic representation is needed, which may be irrelevant to the management decision at hand.

Reservoirs are operated to achieve some objective(s), such as hydropower production, flood control, and water supply. An objective function J^i (with $i = 1, \dots, M$) is defined to measure the performance in each of the M objectives (assumed to be a cost) over an evaluation horizon $[0, h]$ across an ensemble of K realizations of the system disturbances ε_{t+1} , that is,

$$J^i = \Psi_{\varepsilon^1, \dots, \varepsilon^K} [\Phi_{t=0, \dots, h} (g_0^i(s_0, u_0, q_1), \dots, g_h^i(s_h))] \quad i = 1, \dots, M \quad (2)$$

where $g_t^i(\cdot)$ is the i -th immediate and time-separable cost function associated to the time transition from t to $t + 1$, $g_h^i(s_h)$ is a penalty function over the storage reached at the end of the evaluation horizon, Φ_t is an operator for time

aggregation (e.g., the cumulative cost), and Ψ_ϵ is a statistic used to filter the noise generated by the disturbances (e.g., expected value).

The horizon h can be either finite or infinite. In case of a finite horizon, the system is assumed to cease existing after h . Since this is unlikely for most environmental systems, setting a finite horizon generally requires defining a penalty function $g_h^i(s_h)$ accounting for the future cost after h . Conversely, the long-term performance can be considered by setting an infinite evaluation horizon. However, to guarantee the convergence of the cumulated cost for $h \rightarrow \infty$, it is necessary either to discount future costs using a discounting factor $\gamma \in (0, 1]$ or to estimate the average immediate cost (Soncini-Sessa, Castelletti, & Weber, 2007).

The operator Ψ_ϵ reflects the attitude of the water operator in dealing with the uncertainties affecting the system, which can be either described by a probability density function (i.e., $\epsilon_{t+1}^k \sim \phi_t$ for $t = 0, \dots, h - 1$ and $k = 1, \dots, K$) or in a deterministic and set-membership-based fashion (i.e., $\epsilon_{t+1}^k \in \Xi_t$ for $t = 0, \dots, h - 1$ and $k = 1, \dots, K$) (Milanese et al., 1996). In this second case, the uncertainty is considered “deep” as the definition of appropriate prior probability distributions describing ϵ_{t+1} remains contested (Lempert, 2002). The most common formulations of the Ψ_ϵ operator are the expected value (Laplace, 1951) and the worst case (Wald, 1950) across the K realizations of ϵ_{t+1} , but other alternatives can be considered including higher-order moments, such as variance and skewness (Kwakkel et al., 2016), regret-based metrics (Savage, 1951), or satisficing metrics (Simon, 1956) (for a review, see McPhail et al., 2018, and references therein).

3. Stochastic Dynamic Programming

Stochastic Dynamic Programming 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 subsequent costs. The problem is formulated as a series of single-objective problems, where the M objectives are aggregated by means of a scalarization function (e.g., convex combination; Gass & Saaty, 1955) and solved via SDP by computing the long-term cost of a policy for each possible state of the system \mathbf{x}_t (generally the reservoir storage) at time t by means of the Bellman equation:

$$H_t(\mathbf{x}_t) = \min_{\mathbf{u}_t} \Psi_{\epsilon_{t+1}}[\Phi_t[G_t(\mathbf{x}_t, \mathbf{u}_t, \epsilon_{t+1}), H_{t+1}(\mathbf{x}_{t+1})]] \quad (3)$$

where $H_t(\cdot)$ is the optimal cost-to-go function defined over a discrete grid of states for the scalarized objective and $G_t(\cdot)$ represents the corresponding scalarized immediate and time-separable cost function. Equation 3 is restricted to only specific combinations of operators Φ_t and Ψ_ϵ . In particular, if the objective in Equation 2 computes the average cost over the time horizon h (i.e., $\Phi_t = \sum_t/h$), then the disturbances' uncertainty must be filtered by computing the expected value (i.e., $\Psi_\epsilon = E_\epsilon$). Conversely, the computation of the maximum cost in time (i.e., $\Phi_t = \max_t$) requires using the maximum also for filtering the disturbances' uncertainty (i.e., $\Psi_\epsilon = \max_\epsilon$).

Given the recursive formulation of Equation 3, its solution is obtained through a backward-looking optimization over the period $T - 1, \dots, 0$, which first initializes $H_T(\mathbf{x}_T)$ and then iteratively computes $H_t(\mathbf{x}_t)$ using Equation 3 until a suitable termination test is satisfied. From the resulting optimal value function $H_t^*(\mathbf{x}_t)$, the optimal policy is derived as the optimal sequence of control laws $p^* \triangleq [\mu_0(\mathbf{x}_0), \dots, \mu_{T-1}(\mathbf{x}_{T-1})]$, where the operating rule $\mathbf{u}_t = \mu_t(\mathbf{x}_t)$ formalizes mapping between the current system conditions represented by the state vector into release decisions and T is the period of the system (e.g., annual period), by solving the following problem

$$\mu_t^*(\mathbf{x}_t) = \arg \min_{\mathbf{u}_t} H_t^*(\mathbf{x}_t) \quad (4)$$

Note that the operating rule can also be defined as a set-value rule $\mathbf{u}_t = \mathcal{M}_t(\mathbf{x}_t)$ that provides at each time instant a set of decisions providing the same performance in the long term (Orlovski et al., 1983). This formulation, despite being seldom adopted, looks promising as it does not fully replace the human operators who remain in charge of selecting which decision to implement among the set of equally optimal options returned by the rule.

The MDP formulation has been considered particularly suitable for modeling water resource systems to capture nonlinear system dynamics and objective functions, as well as closing the loop between operating decisions and evolving system conditions. The MDP requirements of discrete domains of state, decision, and disturbance variables, along with the time separability of objective functions and constraints, can indeed be applied to many water systems by properly enlarging the state space (Castelletti, Pianosi, & Soncini-Sessa, 2012). SDP is expected to outperform both Approximation in Value and Policy Space methods for a small system (e.g., one or two reser-

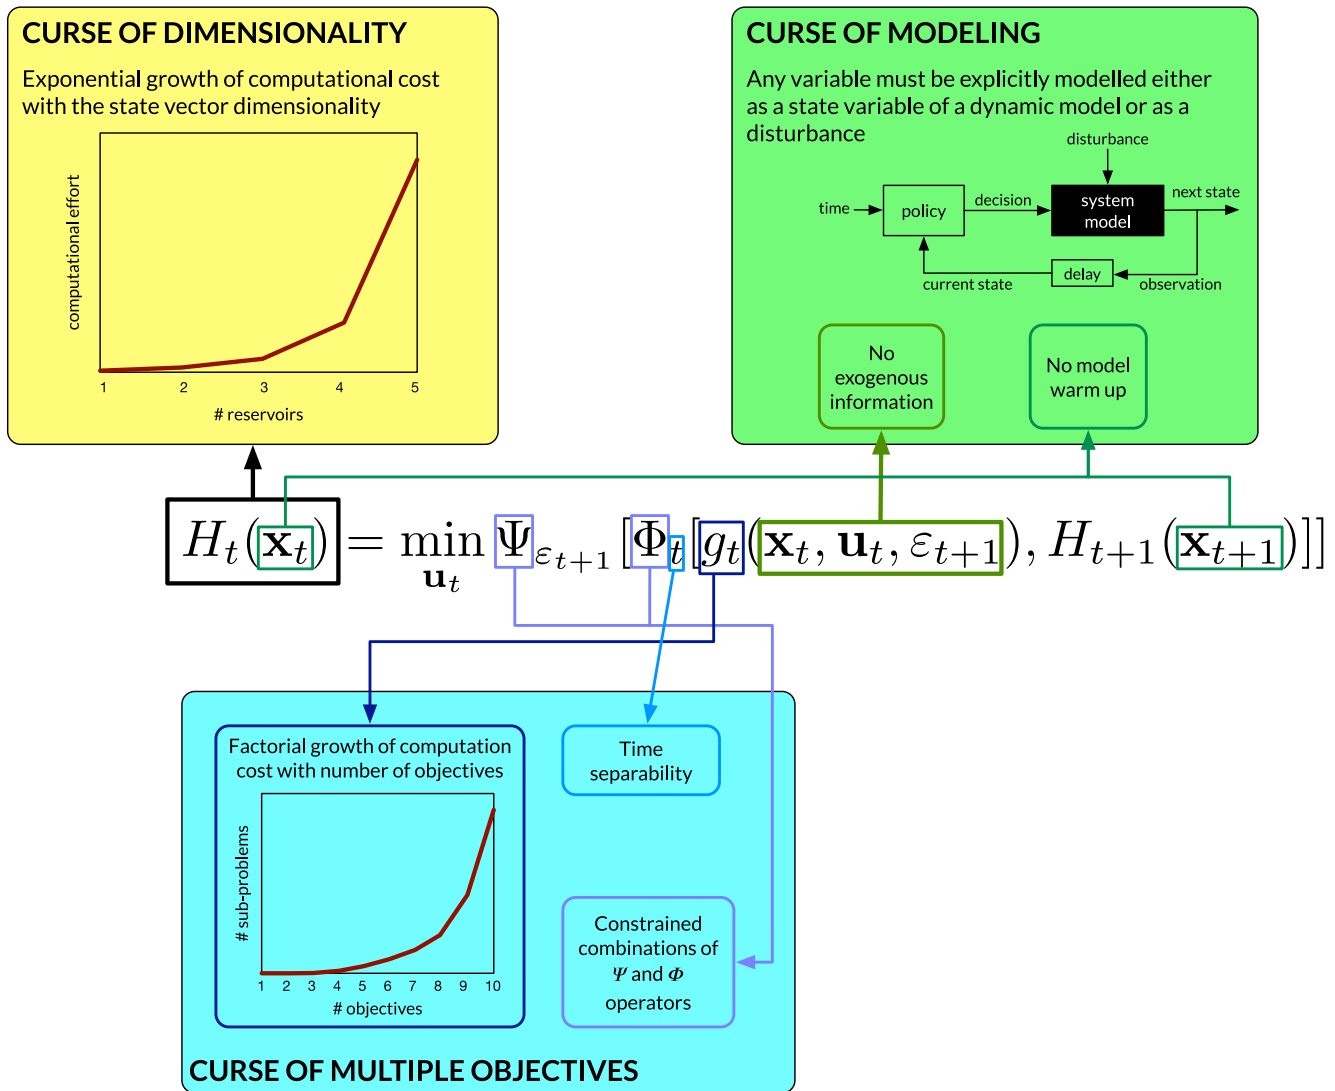


Figure 3. Schematization of the three curses of SDP with respect to the Bellman equation.

voirs) where it is possible to use a dense discretization of the modeled variables and rely on an adequate model of the reservoir inflows. Yet, the computation of the optimal value function $H_t^*(\mathbf{x}_t)$ using Equation 3 substantially limits the adoption of SDP in complex real-world problems due to limitations that here we group into three major curses of dimensionality, modeling, and multiple objectives (see Figure 3).

The *curse of dimensionality* (Bellman, 1957) means that the computation cost of the optimal value function grows exponentially with the state vector dimensionality and SDP becomes intractable when the dimensionality of the system exceeds 2 or 3 reservoirs (Loucks et al., 2005). This curse also arises for a single reservoir system where inflows are described by dynamic models or when additional variables should be included to account for water quality aspects (e.g., Kerachian & Karamouz, 2006; Loucks et al., 1985; Orlob & Simonovic, 1982; Rinaldi et al., 1979). Specifically, let n_x be the number of state variables with N_{x_i} the number of elements in the discretized domain of the i -th state variable, the computational complexity of SDP is proportional to $\prod_{i=1}^{n_x} N_{x_i}$.

The *curse of modeling* (Tsitsiklis & Van Roy, 1996) refers to the SDP requirement that any variable must be explicitly modeled to fully predict the one-step lookahead transition used for the estimation of the value function (Equation 3), either as a state variable of a dynamic model or as a stochastic disturbance. In particular, although reservoir inflows are often both spatially and temporally correlated and largely driven by the climate system dynamics, properly accounting for their temporal correlation or explicitly representing rainfall-runoff hydrolog-

ical processes requires the use of dynamic models, which contribute additional state variables and exacerbate the curse of dimensionality. As a consequence, this curse tends to cause DP application to adopt simplified representations of the uncontrolled part of the water system upstream of the reservoir. No exogenous information, defined as variables that are observable but are not explicitly modeled in the MDP formulation, can therefore be directly used for conditioning the operating policy (Desreumaux et al., 2018). Furthermore, SDP cannot be combined with high-fidelity, distributed parameter, process-based simulation models (e.g., hydrodynamic and ecologic models Ostfeld & Salomons, 2005), which require a warm-up period to remove the effects of the initialization of unobservable state variables. In fact, simulation models that require initial state variables to undergo a warm-up period cannot be employed in the one-step lookahead optimization routine of SDP. This issue has strongly limited the integration of water quantity and quality targets into optimal reservoir operations (e.g., Castelletti et al., 2014; Chaves & Kojiri, 2007; Rieker & Labadie, 2012).

The *curse of multiple objectives* is related to the generation of the full set of Pareto optimal (or approximate) solutions to support a posteriori decision making (Cohon & Marks, 1975) in many-objective control problems. In this study, we formalize this curse into three distinct limitations: time separability, number of objectives, and objective formulations. First, the MDP formulation adopted by SDP requires time-separable objective functions, which severely constrains the type of objectives that can be formulated and prevents the use of duration curves or resiliency as operating objectives. The second limitation is related to the single-objective nature of SDP, which must be repeated for every Pareto optimal point by using different scalarization values (Powell, 2007), such as changing the weights used in the convex combination of the objectives (Gass & Saaty, 1955). The overall cost of SDP to obtain an approximation of the Pareto optimal set scales factorially with the number of objectives (Giuliani, Galelli, & Soncini-Sessa, 2014; Reed & Kollat, 2013). Moreover, using a convex weighting scheme for aggregating the multiple objectives often yields gaps in concave regions that are common in nonlinear problems. Although concave regions can be explored by using alternative scalarization functions, such as the ϵ -constraint method (Haimes et al., 1971), this approach cannot be applied in the SDP framework because it violates the requirement of time separability. Third, the formulation of the Bellman equation (Equation 3) allows only specific combinations of operators Φ_t and Ψ_e , thus making it impossible to explore rival problem framings reflecting different risk attitudes with respect to different objectives (e.g., by calculating some objectives in expectation and others using a worst-case formulation as in Quinn, Reed, Giuliani, & Castelletti, 2017).

4. Beyond Stochastic Dynamic Programming

In this section, we classify and analyze 115 recent publications (see Table S1 in Supporting Information S1), which attempt to overcome some or all of the SDP curses by adopting a less restrictive problem formulation for better reproducing the full complexity of natural and multisector dynamics that characterize most water reservoir systems. Following Bertsekas (2019) and Powell (2019), we distinguish two main approaches: (a) Approximation in Value Space and (b) Approximation in Policy Space. The mathematical details of these methods are reported in the Supporting Information S1, while a list of code repositories providing their implementation is reported in Table 1.

Both AVS and APS methods are often considered as suboptimal control schemes due to their use of approximations for either the value function or the operating policy that are optimal for the formulated Markov Decision Process but may not be optimal for the real system. Caution should be taken by considering how the modeling simplifications required to attain classical definitions of optimality translate to real-world implications. When the SDP curses prevent formulating an MDP that can represent key complexities of real water systems, the formulated problem might generate solutions that are mathematically optimal but not decision relevant in the real system. Approximate solution techniques that permit the use of detailed simulations, broader sources of information, and a wider range of candidate objectives in problem formulations that better capture operational contexts may be closer to the real, perhaps unknowable optimum, and thus more valuable for the real operators.

4.1. Approximation in Value Space

The underlying principle of AVS methods is to approximate the optimal value function $H_t^*(\cdot)$ on the right-hand side of the Bellman Equation 3. There are two main approaches for estimating an approximated value function (Bertsekas, 2005): (a) *explicit* AVS, which includes continuous function approximation, Reinforcement Learning,

Table 1
Code Repositories Implementing AVS and APS Methods Along With Other Useful References for the Broader Problem of Reservoir Operations

Repository	Methods	Language	Website
M3O	AVS and APS	MATLAB	http://mxgiuliani00.github.io/M3O-Multi-Objective-Optimal-Operations/
SDDP.jl	AVS	JULIA	https://github.com/odow/SDDP.jl
Optimist	AVS	PYTHON	https://github.com/luciofaso/optimist
Policy Tree	APS	PYTHON	https://github.com/jdherman/ptrecopt
VIC-ResOpt	APS	PYTHON	https://github.com/thanhiwer/VICResOpt
LakeComo EMODPS	APS	CPP	https://github.com/mxgiuliani00/LakeComo
LakeProblem DPS	APS	CPP	https://github.com/julianneq/Lake_Problem_DPS
iRONS	APS	PYTHON	https://github.com/AndresPenuela/iRONS
reservoir_operation_optimization_examples	-	MATLAB	https://github.com/barneydobson/reservoir_operation_optimization_examples
VIC-Res	-	PYTHON	https://github.com/thanhiwer/VICRes
ReservoirManagement	-	JULIA	https://github.com/dourouc05/ReservoirManagement.jl
Reservoir	-	R	https://cran.r-project.org/web/packages/reservoir/index.html
CALFEWS	-	PYTHON	https://github.com/hbz5000/CALFEWS
PySedSim	-	PYTHON	https://github.com/FeralFlows/pysedsim

and Sampling Stochastic Dynamic Programming; (b) *implicit* AVS, which solves the problem online (e.g., Model Predictive Control). Both explicit and implicit AVS methods require that problems must be formulated as MDPs and are often considered approximate dynamic programming methods (Powell, 2007).

4.1.1. Continuous Function Approximation

The most common explicit AVS approach is a continuous function approximation. Rather than computing the optimal value function for all N_x values of state variables' discretized domain, this approach approximates it by means of a continuous function interpolated over a smaller number of value function evaluations $\tilde{N}_x < N_x$ (see the Supporting Information S1 for details about the mathematical formulation). Different continuous function approximations have been proposed, including piecewise linear polynomials (e.g., Tsitsiklis & Van Roy, 1996), Hermite polynomials (Foufoula-Georgiou & Kitanidis, 1988), splines (e.g., Chen et al., 1999; Johnson et al., 1993; Lamontagne, 2015), or artificial neural networks (e.g., Castelletti et al., 2007).

In particular, the piecewise linear approximation of the value function is the core of Stochastic Dual Dynamic Programming (SDDP, see Pereira & Pinto, 1991), a technique that approximates the value function through sampling and decomposition (e.g., Goor et al., 2011; Macian-Sorribes et al., 2017; Pereira-Cardenal et al., 2016; Raso et al., 2017; Tilmant et al., 2007). Specifically, SDDP computes an approximation of the value function over a state-space domain that is no longer discretized but rather sampled in the form of Bender cuts, corresponding to linear approximations of the value function derived from the primal and dual information of the one-stage optimization step at time $t + 1$ (Shapiro, 2011). This approach is structured into two phases. First, a backward optimization estimates an upper bound of the true value function for the sampled points. In the second phase, the system is simulated forward on a series of synthetic and/or historical hydrologic scenarios to check that the estimated Bender cuts approximation is acceptable, otherwise the two phases are iterated.

The main advantage of continuous function approximation with respect to SDP is the reduced computational complexity when computing the value function on $\tilde{N}_x < N_x$ state values. This enables the solution of much larger problems where the number of reservoirs is higher than 10 (Goor et al., 2010; Macian-Sorribes et al., 2017) and where the state vector dimensionality is even higher to include irrigation districts (Tilmant et al., 2008) and hydrologic information through the use of autoregressive models (Goor et al., 2011; Tilmant et al., 2010),

which allows for extension to include snow- and precipitation-related variables (Pina et al., 2017). SDDP requires the optimization problem to be modeled as linear since linearity ensures cost-to-go function convexity (Raso et al., 2017). Finally, all continuous function approximation approaches remain largely constrained by the curse of multiple objectives; all of the reviewed studies (Table S2 in Supporting Information S1) consider a single (or a priori aggregated) time-separable objective and none explore mixed formulations capturing different risk attitudes across different objectives.

4.1.2. Reinforcement Learning

Reinforcement Learning (RL) is another explicit AVS approach that relies on the same principle of continuous function approximation to reduce the number of value function evaluations needed to explore the state-decision space performed by SDP. Specifically, RL is based on the idea of designing the optimal operating policy through a trial-and-error process and learning from experience (Sutton & Barto, 1998). The learning experience is gained either in the real environment or generated using a system model or historical observations. While the first option is clearly impractical for real water reservoir systems as it may result in unacceptable social costs (e.g., bad decisions producing high damages must be experienced) and the learning process would be too long, offline model-based learning offers the opportunity to simulate the system dynamics and generate knowledge of the stochastic environment, including the possibility of avoiding any explicit probabilistic characterization of the system disturbances that is instead learned via experience (Castelletti et al., 2002; Lee & Labadie, 2007).

The most popular RL algorithm is Q-learning (Watkins & Dayan, 1992), which optimizes the action-value function version of the Bellman equation through incremental updates during the learning phase (see the Supporting Information S1 for details about the mathematical formulation). In contrast to SDP, Q-learning works online and the next state is not estimated by a model, but observed at time $t + 1$ in the real system or in an external model simulation. The action-value function is then updated balancing its current value and the recently experienced transition. This procedure alleviates the SDP curse of dimensionality, as the optimization is not exploring the entire state-decision space and instead is performed only for the visited states by testing alternative sequences of decisions to be generated with a proper design of experiments. Yet, the action-value function still requires a discretization of the state-decision space, which yields a computational complexity that grows exponentially with the number of state and decision variables.

A more recent approach called fitted Q-iteration (FQI, see Ernst et al., 2005 and Castelletti et al., 2010) combines the RL idea of learning from experience with the continuous function approximation scheme illustrated in the previous section. This method relies on the experience contained in a sample data set \mathcal{F} previously collected either from system observations or model simulations, which is used for estimating an approximation of the value function by means of tree-based regressors (Geurts et al., 2006). FQI replaces the recursive solution of the Bellman equation with a sequence of nonlinear regressions over the sample data set \mathcal{F} . While Q-learning converges only when the value function updates are performed incrementally, FQI processes the full data set \mathcal{F} in a batch mode for updating the value function, which speeds up algorithm convergence (Kalyanakrishnan & Stone, 2007). An extension of this method to multiobjective problems is the Multi-Objective FQI proposed by Pianosi et al. (2013) and Castelletti et al. (2013). While using traditional FQI in multiobjective problems would require repeated optimization runs using different scalarizations of the objectives, MOFQI enlarges the continuous approximation of the value function to the weight space by including a new state variable (i.e., the weight used in the objectives' aggregation) within the arguments of the value function. Including the weight space enables the generation of the full Pareto optimal set in a single optimization run while slightly increasing the computational cost of a single FQI run due to an enlarged state space (Castelletti et al., 2013).

The continuous tree-based approximation of the value function used by FQI and MOFQI mitigates the SDP curse of dimensionality by allowing the use of a coarse grid of state and decision domains. The use of deep RL (Lillicrap et al., 2015), albeit still in its infancy, could further reduce the limitations of the curse of dimensionality (e.g., Xu et al., 2020, 2021). In addition, the batch-mode learning from the sample data set generated offline combining historical observations and model simulations aids RL in overcoming the two limitations related to the curse of modeling: First, operating policies can be conditioned upon exogenous information of the uncontrolled part of the system or forecast information (Rieker & Labadie, 2012); second, it avoids the one-step lookahead SDP optimization and therefore can be combined with any simulation model, including complex process-based models for controlling water quality targets (e.g., Castelletti et al., 2014; Giuliani, Galelli, & Soncini-Sessa, 2014). However, RL approaches (Table S2 in Supporting Information S1) remain largely constrained by the curse of multiple

objectives, where the underlying MDP formulation requires time-separable objective functions and the explicit approximation of the value function prevents adopting mixed combinations of operators for time aggregation and uncertainty filtering in the definition of the objectives. Lastly, most RL applications are either single objective (e.g., Castelletti, Galelli, Restelli, & Soncini-Sessa, 2012; Lee & Labadie, 2007; Madani & Hooshyar, 2014; Rieker & Labadie, 2012) or rely on the same scalarization techniques adopted by SDP for exploring trade-offs across competing objectives (e.g., Castelletti et al., 2014, 2010; Giuliani, Galelli, & Soncini-Sessa, 2014), with the notable exception of multiobjective RL algorithms (Castelletti et al., 2013). While the use of multiobjective Reinforcement Learning in the reservoir control domain is still limited, there is a growing opportunity to test emerging methods developed in other domains (for a review, see Hayes et al., 2021, and references therein). However, the increasing state-space dimensionality to include the weights used in the objective aggregation suggests the need to test the scalability of multiobjective RL to many-objective policy design problems.

4.1.3. Sampling Stochastic Dynamic Programming

Sampling Stochastic Dynamic Programming (SSDP) is an alternative explicit AVS approach that uses an approximation of the system's disturbances rather than (or in addition to) an approximation in the state space. SSDP replaces the SDP closed-form probabilistic description of system disturbances (e.g., reservoir inflows) with multiple scenarios that form an empirical distribution of the disturbances (see the Supporting Information S1 for details about the mathematical formulation). This can better represent flow persistence and spatial correlation than in traditional SDP, as intact hydrographs are used rather than explicit probabilistic models (Faber & Stedinger, 2001). The scenarios can be either multiple historical time series (Kelman et al., 1990) or ensemble streamflow projections (e.g., Côté & Leconte, 2015; Faber & Stedinger, 2001; Kim et al., 2007), with the latter providing frequently updated information about the current system conditions.

The use of scenario ensembles primarily combats the curse of modeling (Kim et al., 2007). Traditional SDP requires an explicit probabilistic model of stochastic state transitions, which for large, multireservoir systems would necessitate either the addition of many stochastic state variables or the introduction of simplifying representation of the inflow processes. In the SSDP formulation, instead, a single state variable implicitly captures all relevant stochastic variables and their complex, nonlinear spatial and temporal dependency structures. Furthermore, if historical flows are used as scenarios, the historic relationship between flows and water/energy demands and prices can also be included in the optimization (Lamontagne, 2015). For SSDP to perform well, the scenario ensemble must provide a good representation of relevant future uncertainties. If this is not the case, for instance, due to nonstationarity or inadequate coverage of extremes, then SSDP may perform poorly.

While SSDP addresses the curse of modeling through a compact representation of the stochastic state, the curse of dimensionality still limits its applicability to large-scale multireservoir systems. SSDP can be readily paired with continuous function approximations to reduce the computational burden of solving large-scale problems (e.g., Côté & Leconte, 2015; Lamontagne, 2015; Vicuna et al., 2010). However, this combination of methods is prone to the limitations of SDDP described above. Finally, SSDP is also constrained by the curse of multiple objectives; to our knowledge, no SSDP study (Table S2 in Supporting Information S1) has considered multiple objectives and none have explored nontime separable or mixed objective formulations.

4.1.4. Model Predictive Control

Model Predictive Control is an implicit AVS approach based on the sequential, online resolution of multiple open-loop control problems defined over a finite, receding time horizon (Bertsekas, 2005; Scattolini, 2009). At each time t , a forecast over the finite future horizon $[t, t + \kappa]$ of the external disturbances, generally the inflow to the reservoir $\hat{q}_{t+1}, \dots, \hat{q}_{t+\kappa}$, is generated by a model predictor that uses all information available at time t (e.g., precipitation, snowpack, and inflow at previous time). The corresponding sequence of optimal decisions is then obtained by solving a mathematical programming problem (Maciejowski, 2002), assuming that the realization of the disturbances will be equal to the predicted value (i.e., deterministic MPC; see the Supporting Information S1 for details about the mathematical formulation). Only the first control is actually applied and, at time $t + 1$, a new problem is formulated over the horizon $[t + 1, t + 1 + \kappa]$ on the basis of the updated information available, namely the state of the system at time $t + 1$ as well as updated forecasts (Mayne et al., 2000). This sequential resolution of the problem generates a feedback between the operational decisions and the updated state, which can partially

compensate for the effects of the disturbances as it is unlikely that the actual realization of the inflow equals the predicted value, with the system actually not evolving as expected.

The availability of good forecasts reduces the distance between expected and actual conditions, resulting in near-optimal decisions. However, this deterministic formulation, also referred to as naive feedback control (Bertsekas, 1976), ignores uncertainty and thus may fail to hedge against risk appropriately. Stochastic MPC extensions allow some forms of explicit probabilistic characterization of forecast uncertainty. The open-loop feedback control (OLFC) describes the future disturbances according to their probability distribution and computes the objectives through some function to filter the disturbances (e.g., expected values). The OLFC performance can be further improved by adopting a partial open-loop feedback control (POLFC) formulation (e.g., Castelletti et al., 2008; Faber & Stedinger, 2001; Howard, 1992; Pianosi & Soncini-Sessa, 2009), which explicitly assumes that in the future the state of the system will be measured and a new problem will be reformulated. The POLFC problem, therefore, computes at each time step the optimal release decision for the first step reflecting first-step uncertainty and the optimal operating policy for the following steps.

Another approach to capture forecast uncertainty is the Tree-Based MPC proposed by Raso et al. (2014). In this case, the forecast is provided in the form of an ensemble generated by running the forecast model multiple times to account for major sources of forecast uncertainty (Gneiting & Raftery, 2005). This ensemble is then transformed into a tree where similar ensemble members are bundled together into one trajectory (branch) up to the point when some of them start to significantly diverge from the others. The tree structure is then used to optimize a control tree defining a distinct control sequence for each branch. Control sequences are constrained to be the same up to the time when two ensemble members diverge.

MPC provides two main advantages with respect to Stochastic Dynamic Programming: it allows handling multireservoir systems with more than three reservoirs (e.g., Ficchi et al., 2016; Karimanzira et al., 2016; Kistenmacher & Georgakakos, 2015), because it avoids computation of the value function by searching the sequence of optimal controls over a finite horizon; its structure exploits real-time exogenous information, particularly inflow forecasts from hourly (e.g., Galelli et al., 2014, 2015) to seasonal or longer (e.g., Anghileri et al., 2016; Raso & Malaterre, 2016) lead times. Moreover, different from other approaches applied at daily or monthly time scales, MPC is used to control reservoirs with subdaily decisions in 30% of the reviewed studies (Table S3 in Supporting Information S1). The application of MPC in multiobjective problems suffers the same limitations as SDP; MPC is a single-objective approach and relies on the same scalarization techniques adopted by SDP for exploring trade-offs across competing objectives (e.g., Giuliani & Castelletti, 2013; Karimanzira et al., 2016; Pianosi & Soncini-Sessa, 2009), the objectives are formulated as time-separable objective functions, and none of the reviewed studies explored mixed formulations to reflect different risk attitudes with respect to different objectives. Most MPC studies (Table S3 in Supporting Information S1) consider either a single objective (e.g., Ficchi et al., 2016; Xu et al., 2015) or a priori weight-based aggregation of objectives (e.g., Kistenmacher & Georgakakos, 2015; Raso et al., 2014; Uysal et al., 2018; Wang et al., 2014).

4.2. Approximation in Policy Space

The main idea of APS methods is to replace the traditional SDP approach based on the computation of the value function with a policy design that directly operates in the policy space. The main APS approach is Direct Policy Search (DPS), or *explicit* APS, a simulation-based optimization also known as parameterization-simulation-optimization in the water resources literature (Koutsoyiannis & Economou, 2003a). An alternative approach is the Implicit Stochastic Optimization (ISO), or *implicit* APS, which first solves a deterministic problem and then fits a parameterized policy on the trajectory of optimal releases of the deterministic solution.

4.2.1. Direct Policy Search

Direct Policy Search (Rosenstein & Barto, 2001; Schmidhuber, 2001) is based on the parameterization of the operating policy p_θ within a given family of functions (e.g., linear, piecewise linear, nonlinear network of sigmoid or radial basis functions, etc.) and the exploration of the parameter space Θ to find a parameterized policy that optimizes the objective functions vector \mathbf{J} (Rückstieß et al., 2010) (see the Supporting Information S1 for details about the mathematical formulation).

As noted in Nalbantis and Koutsoyiannis (1997), DPS can be seen as an optimization-based generalization of well-known simulation-based, single-purpose heuristic operating rules (Lund & Guzman, 1999), such as the New York City rule (Clark, 1950), the spill-minimizing “space rule” (Clark, 1956; Johnson et al., 1991), or the widely used Standard Operating Policy (Draper & Lund, 2004). Selecting a suitable class of function for the parameterization of the operating policy is crucial for the discovery of high-performing solutions because this choice might restrict the policy search to a subspace of the decision space that might not include the optimal set of solutions. Many studies (see Table S4 in Supporting Information S1) apply linear (e.g., Ashrafi & Dariane, 2017; Fallah-Mehdipour et al., 2015; Feng et al., 2017) or piecewise linear (e.g., Hu et al., 2016; Pan et al., 2015; Tan et al., 2017) policy functions, which likely constrain the possibility of conditioning the release decisions on multidimensional state vectors (e.g., making the decision dependent on more information than reservoir storage, including the time variability of the policy). A nonlinear multi-input multi-output function, such as a nonlinear approximating network (Tikk et al., 2003), can provide flexibility to the shape of the operating policy, which is particularly important in multireservoir systems that require coordinated operation across the reservoirs (e.g., Giuliani, Anghileri, et al., 2016; Biglarbeigi et al., 2018).

DPS offers some substantial advantages over Stochastic Dynamic Programming: First, it avoids the computation of the value function and allows handling multireservoir systems larger than three reservoirs (e.g., Geressu & Harou, 2015; Karami & Dariane, 2018). Second, DPS allows the direct use of exogenous information through a data-driven controller tuning approach (Formentin et al., 2013; Hou & Wang, 2013), where release decisions are directly conditioned upon nonmodeled observational data from the climate or hydrologic systems to better capture the variability of stochastic processes that cannot be accurately modeled or forecasted. Examples in Table S4 in Supporting Information S1 include the use of inflow forecasts (e.g., Adams et al., 2017; Herman & Giuliani, 2018; Xu et al., 2017), observations of previous day inflow (e.g., Bozorg-Haddad et al., 2008; Giuliani et al., 2015), Snow Water Equivalent (e.g., Denaro et al., 2017; Desreumaux et al., 2018), Sea Surface Temperature (Giuliani et al., 2019), ENSO indexes (Libisch-Lehner et al., 2019), and crowdsourced information (e.g., Giuliani, Castelletti, Fedorov, & Fraternali, 2016). Moreover, the simulation-based optimization allows the use of models requiring a warm-up period (e.g., Castelletti, Pianosi, Quach, & Soncini-Sessa, 2012; Giuliani, Castelletti, Pianosi, et al., 2016), which conflicts with the step-based optimization mode of SDP. Third, coupling DPS with multiobjective evolutionary algorithms enables the generation of a Pareto approximate set of solutions in a single optimization run, which allows broadening the number and complexity of objectives that can be resolved. This feature favors transitioning from multiobjective problems including two or three objectives (e.g., Ahmadi et al., 2014; Ashofteh et al., 2015; Labadie & Wan, 2010) to many-objective problems where the number of objectives is four or more (e.g., Giuliani, Herman, et al., 2014; Hurford & Harou, 2014; Libisch-Lehner et al., 2019). In addition to exploring multisectoral trade-offs, including a broader suite of objective functions in the policy search enables the implementation of a rival framing approach to explore multiple competing problem formulations that capture a range of stakeholder attitudes toward risk and uncertainty (e.g., Giuliani & Castelletti, 2016; Quinn, Reed, Giuliani, & Castelletti, 2017). As a simulation-based optimization, DPS is not restricted to time-separable cost functions and allows formulating more complex objectives such as the resiliency of a hydropower system (e.g., Ahmadi et al., 2014) or a flood protection system (e.g., Quinn, Reed, Giuliani, & Castelletti, 2017). Avoiding the restriction to a single value function also relaxes the constraint on the combinations of operators for time aggregation and uncertainty filtering in Equation 3, opening up the possibility of mixed objective formulations, such as average in time and worst case over disturbances' scenarios (e.g., Giuliani & Castelletti, 2016; Giuliani, Herman, et al., 2014; Libisch-Lehner et al., 2019) or maximum cost over time and expected value over disturbances' scenarios (e.g., Biglarbeigi et al., 2018; Quinn, Reed, Giuliani, & Castelletti, 2017).

However, some challenges still limit the adoption of DPS approaches. One key issue is the approximation of the objective function via simulation of system dynamics over a sufficiently long time series of disturbance realizations (Pianosi et al., 2011). This approach can overfit the policy parameters to the particular stochastic realizations experienced during the simulation-based optimization, potentially yielding impressive calibration results that can largely degrade when tested on out-of-sample observations. This issue can be overcome by splitting the available observations (or an ensemble of synthetically generated records) into two statistically equivalent data sets to optimize the policy parameters on the first set of data and test policy performance on the second data set (Brodeur et al., 2020). Another often overlooked aspect of DPS is the a priori definition of the policy architecture, generally based on intuition, analytical methods, or on few trial-and-error experiments. Yet, the functional class

Table 2

Summary of the 115 Applications of Approximation in Value Space and Approximation in Policy Space With Respect to the Three Curses of Stochastic Dynamic Programming

SDP curse	Metric	Explicit AVS	Implicit AVS	Explicit APS	Implicit APS
Dimensionality	Average number of reservoir	3.56	2.95	1.65	1.60
Modeling	% Studies using exogenous information	25%	100%	58%	60%
	% Studies using models warm-up	7%	5%	5%	60%
Multiple objectives	Average number of objectives	1.4	1.5	2.5	1.2
	% Studies using time-separable objectives	100%	100%	87%	100%
	% Studies using mixed objective formulations	0%	0%	15%	0%
Number of publications (share of the total)		28 (24.3%)	20 (17.4%)	62 (53.9%)	5 (4.3%)

should be tailored to the problem at hand and its selection is crucial, as it determines the search space within which solutions can be found as well as their attainable control dynamics (Zaniolo et al., 2021).

4.2.2. Implicit Stochastic Optimization

Implicit Stochastic Optimization is an implicit APS approach originally proposed by Young (1967) and Karamouz and Houck (1987). ISO is structured in two steps (see the Supporting Information S1 for details about the mathematical formulation): first, it solves a deterministic problem to find optimal reservoir releases under several inflow scenarios, which are generally synthetically generated to cover a wide range of different hydrologic conditions; the ensemble of optimal release trajectories is then used as a target in a regression problem to identify a parameterized policy approximating the deterministic solution.

ISO applications use different optimization algorithms for solving the deterministic problem, including deterministic Dynamic Programming (e.g., Zhou et al., 2019), quadratic programming (e.g., Celeste & Billib, 2009), and genetic algorithms (e.g., Sangiorgio & Guariso, 2018). Alternative parameterized functions depending on a variable number of policy inputs are used in the policy identification phase, including piecewise linear functions (e.g., Zhou et al., 2019), power laws (e.g., Celeste & El-Shafie, 2018), artificial neural networks (e.g., Sangiorgio & Guariso, 2018), and extreme learning machines (e.g., Feng et al., 2019).

The main advantage of ISO with respect to Stochastic Dynamic Programming is in terms of mitigating the curse of modeling because the identification of the parameterized policy via regression allows the direct use of exogenous information among the policy inputs (e.g., Celeste & Billib, 2009; Sangiorgio & Guariso, 2018) as well as the use of model warm-up. ISO approaches remain, however, constrained by both the curse of dimensionality and the curse of multiple objectives; most ISO works (Table S4 in Supporting Information S1) consider single reservoir systems and a single, time-separable objective and despite the use of inflow scenarios, none of these studies explores mixed objectives' formulations.

5. Discussion and Future Research Opportunities

A summary of the critical analysis conducted over the 115 studies reviewed in the previous section is reported in Table 2 (for details, see Tables S1–S4 in Supporting Information S1), which illustrates a set of metrics reflecting how Approximation in Value Space and Approximation in Policy Space applications cope with the three SDP curses and generally outperforms SDP also in terms of computational requirements (e.g., Castelletti et al., 2010; Côté & Leconte, 2015; Desreumaux et al., 2018; Galelli et al., 2014; Giuliani, Castelletti, Pianosi, et al., 2016). The distribution of the publications across the four approaches suggests that DPS is emerging as the most widely adopted method to advance water reservoir control. The table also shows that large multireservoir systems have been mostly solved using explicit AVS methods and particularly SDDP that scales far beyond SDP limits. The use of exogenous information is instead predominant for implicit AVS methods, which require inflow forecasts to solve the multiple open-loop control problems over a finite, receding time horizon; conversely, only SDDP among the explicit AVS approaches relies on inflow forecasts for informing the operating policy. Moreover, about 60% of APS applications condition operational decisions on exogenous information, mostly in the form of observed hydrological variables while the use of forecasts is still limited. It is worth mentioning that MPC

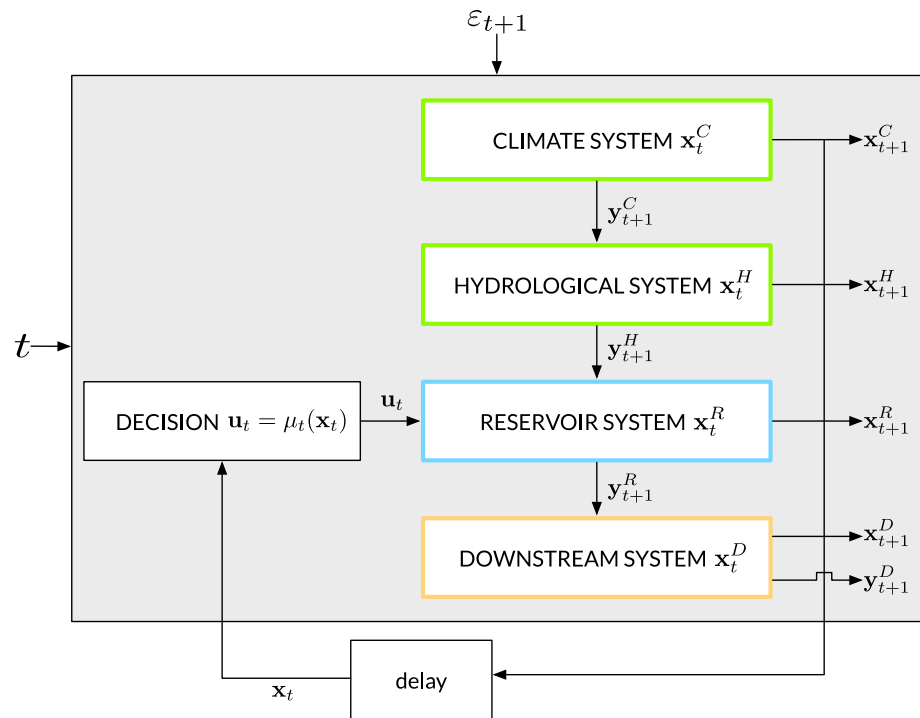


Figure 4. Illustration the four main dynamic components of a water reservoir system.

introduces stronger requirements than APS regarding forecast skill, as the predicted inflows must be sufficiently accurate for simulating the reservoir dynamics over the considered lead time. Conversely, APS methods can be informed by biased forecasts as they are more sensitive to forecast anomalies rather than absolute values (Giuliani et al., 2020). This difference becomes particularly relevant when considering seasonal forecasts that provide skillful information in terms of long-term hydroclimatic anomalies (e.g., Crochemore et al., 2020). Focusing on the second limitations related to the curse of modeling, we observe that few studies (e.g., Castelletti et al., 2014) used high-fidelity, distributed parameter, process-based simulation models that require initial state variables to undergo a warm-up period.

A marked difference between DPS and other methods emerges with regard to the number of objectives considered in the reviewed applications: AVS methods typically solve problems with 1 or 2 objectives, whereas the average number of objectives resolved via DPS is equal to 2.5. Moreover, DPS methods provide the most flexibility in terms of objective formulations (concerning both nontime separable and mixed formulations), while all AVS applications and also ISO methods tend to adopt traditional time-separable objectives evaluated as expected average costs. The maximization of expected performance generates solutions that may perform poorly under extreme flow quantiles (floods and droughts), while a reduction in average system performance is often considered acceptable if accompanied by a sufficient reduction in variance (Beyer & Sendhoff, 2007).

It should be noted that addressing many-objective problems to support the exploration of multidimensional trade-offs is critical in operational design studies, while it becomes less compelling in real-time operational settings. We recommend using DPS to negotiate trade-offs and priorities offline, while MPC becomes more appropriate for short-term operational decisions. How to bridge the two time scales is still an unsolved problem, which could be tackled by using a multiscale MPC approach with a cascade of controllers working at different time scales (see Lu, 2015, and references therein).

The results of our analysis suggest a possible road map of the main open research directions to be investigated over the next years for further improving existing problem formulations. We argue that, from a systems point of view, the traditional formulation focused on the dynamics of the reservoir storage should be enlarged to capture other systemic components, including the climate system, the hydrological system, and the downstream system as illustrated in Figure 4. Each of the four dynamic systems in Figure 4 has its own state \mathbf{x}_t^j ($j = C, H, R, D$) and

the outputs of one system \mathbf{y}_{t+1}^j generally represent inputs to another system. An example of these interactions is the following: the climate conditions (e.g., air, sea, land temperature, humidity, etc.) produce as output of the climate system the temperature and precipitation in a river basin, which represent the main drivers of the hydrological processes generating the streamflow at the basin outlet. This hydrological system output is then considered as the inflow to the reservoir system, which produces the reservoir release. Finally, the volume of water released can be used for irrigation in a downstream agricultural district, which returns in outputs the crops' yield at the end of the agricultural season. It is worth mentioning that while the reservoir system state is the storage, the hydrological and climate states are more complex, spatially distributed, and might be only partially observable. Moreover, the downstream system often includes complex multisector dynamics providing diverse and potentially competing water-related services.

The composite system evolves over time with dynamics affected by release decisions \mathbf{u}_t and external disturbances $\boldsymbol{\varepsilon}_{t+1}$, that is,

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \quad (5)$$

where the release decisions at each time step are determined by an operating rule $\mathbf{u}_t = \mu_t(\mathbf{x}_t)$, which formalizes mapping between the current system conditions represented by the state vector into release decisions, with the state of the system demonstrated to be sufficient for conditioning optimal operational decisions (Bertsekas, 1976). 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). Given the dynamic model of the system in Equation 5 and the resulting trajectory over time τ of states, decisions, and outputs variables, the set of Pareto optimal policies \mathcal{P}^* is designed by solving the following problem that preserves a more realistic formulation by removing the MDP requirements on system representations:

$$\min_p \mathbf{J}(\tau) = |J^1(\tau), \dots, J^M(\tau)| \quad (6)$$

According to Figure 4, the reservoir system (cyan block) is the one directly controlled by the release decisions and that indirectly influences the downstream system (orange block) by releasing water that supplies irrigation to the downstream agricultural district. Conversely, climate and hydrological systems (green blocks) are uncontrolled and partially observable, but may still provide valuable information to the control rule. For example, the snow water equivalent is a component of the hydrological system state that can be extremely valuable in snow-dominated catchments (e.g., Denaro et al., 2017; Desreumaux et al., 2014). Despite some pioneering studies (e.g., Georgakakos et al., 1998; Guariso et al., 1984; Tejada-Guibert et al., 1995) showed the value of inflow forecasts, most commonly in the literature to date (e.g., Bozorg-Haddad et al., 2008; Giuliani et al., 2018; Karamouz et al., 2009; Labadie & Wan, 2010; Macian-Sorribes et al., 2017; Raso et al., 2017) the hydrological state is approximated as the inflow at time $t - 1$, while the use of the climate state is still largely unexplored.

Today, recent advances in monitoring and forecasting systems are making available a wider range of data that represent an untapped resource to better understand water systems' complexity and to improve the predictability of their future evolution, ultimately increasing our ability to make better operational decisions. Advances in data analytics and machine learning (e.g., Watt et al., 2016) can provide valuable support to the identification of the most valuable information, including observations (e.g., Giuliani et al., 2015), medium- to long-term forecasts (e.g., Anghileri et al., 2016; Kistenmacher & Georgakakos, 2015), and climatic trends (e.g., Herman & Giuliani, 2018). Diversifying the suite of information used to condition operational decisions by capturing the state of the hydrologic system (e.g., Denaro et al., 2017; Pina et al., 2017) and, possibly, the climate state (e.g., Giuliani et al., 2019; Libisch-Lehner et al., 2019) can be an asset for improving reservoir operations, particularly under increasingly variable conditions and more frequent climate extremes. Yet, most of the existing approaches are effective under stationary hydroclimatic conditions that allow the characterization of the uncertainty of the system's disturbances. The resulting solutions are instead vulnerable in a nonstationary environment that would require to timely adapt the system operation to the evolving conditions (e.g., Cohen & Herman, 2021; Feng et al., 2017; Herman & Giuliani, 2018; Zhang, Liu, et al., 2017). MPC can mitigate this risk as the online approach allows the adaptation of the operation to the changing conditions, assuming the updated inflow forecasts are able to capture the underlining nonstationary trends. Besides, the nonstationarity and growing uncertainties of future conditions will require a careful formulation of the operating objectives (e.g., McPhail et al., 2018; Zhang et al., 2018).

The enlarged formulation proposed in Figure 4, however, features evident modeling challenges that suggest the need of combining policy design methods with emulation modeling. An emulator is a low-order, computationally efficient surrogate model identified from an original, complex process-based model. The emulator is then used to replace the original model for computationally intensive applications, such as sensitivity analysis, scenario analysis, evolutionary algorithms, and optimal control (Castelletti, Galelli, Ratto, et al., 2012; Razavi et al., 2012). The rationale for developing emulators is that while all models inevitably simplify complex systems, we need modeling strategies that simplify judiciously with a deeper understanding of the system and the problem at hand because not all the process details in the original model are equally important and relevant to the dynamic behaviors that affect operational decisions (Castelletti, Galelli, Restelli, & Soncini-Sessa, 2012). This step becomes crucial for addressing the balance in the credibility, salience, and legitimacy of the final decision support (Cash et al., 2003). There are two broad families of emulator models: response surfaces and dynamic emulators. Response surface emulators employ static data-driven function approximation techniques to empirically approximate the performance response simulated by the complex model (i.e., approximate objective values) (Kleijnen, 2008). Conversely, dynamic emulators preserve the dynamic nature of the original model when providing computationally fast estimates of state dynamics (Castelletti, Galelli, Ratto, et al., 2012). Despite being largely adopted in a number of water systems analysis publications especially for groundwater applications (e.g., Asher et al., 2015; Bizzi et al., 2015; Müller & Shoemaker, 2014; Regis & Shoemaker, 2007; Siade et al., 2010; Ushijima & Yeh, 2013), their use in designing optimal reservoir operations is still in its infancy with few applications using either a response surface for simplifying the controlled downstream system dynamics (Galelli & Soncini-Sessa, 2010; Giuliani, Anghileri, et al., 2016) or a dynamic emulator for approximating the reservoir's dynamics including water quality processes (Galelli et al., 2015). Emulation modeling is then expected to significantly contribute in isolating key decision-relevant processes across the four interactive dynamic systems illustrated in Figure 4 for accommodating realistic representations of the original system complexity in computationally efficient models.

Advancing the representation of optimal reservoir control will also contribute to the design of diverse portfolios of water management actions that combine reservoir operations with groundwater sources (e.g., Macian-Sorribes et al., 2017; Nayak et al., 2018; Pulido-Velazquez et al., 2016), water markets (e.g., Erfani et al., 2015; Kasprzyk et al., 2012), demand management (e.g., Escriva-Bou et al., 2018; Gonzales & Ajami, 2017), and financial tools (e.g., Denaro et al., 2018; Foster et al., 2015) by providing a better understanding of the potential role of reservoir operations within a candidate portfolio. Moreover, the portfolio design should include these recent advances in reservoir operations to avoid misrepresenting important systemic complexities that could bias the estimation of their impacts. Lastly, the ubiquitous presence of dams (Mulligan et al., 2020) makes water reservoirs one of the most important infrastructures impacting hydrological processes, which should be properly represented by including advanced reservoir operation models and closed-loop policies in large-scale hydrological models in order to enable the development of reliable and credible projections about the future coevolution of coupled human-natural systems and better inform the planning of adaptation options (e.g., Dang et al., 2020; Giuliani, Li, et al., 2016; Thompson et al., 2013).

6. Conclusions and Challenges Ahead

Advancing existing approaches for the design of optimal water reservoir operations is an ongoing relevant research problem aimed at supporting both the reoperation of existing infrastructure and the increasing rate of dam construction in fast developing countries. Despite being extensively studied, this topic is still timely due to the presence of new challenges, such as changes in societal perception of natural resources and increasing uncertainties altering decision makers' preferences, and emerging opportunities including better information and more computing power. For this paper, we reviewed 336 recent publications focusing on how the operation design problem is formulated, rather than solved. Moreover, we provide a critical analysis of 114 studies, which attempt to overcome some or all of the curses of Stochastic Dynamic Programming.

Our analysis identifies Approximation in Value Space methods as the approach that better copes with the curse of dimensionality. It is worth mentioning that AVS applications mostly scale in terms of number of reservoirs but generally neglect the state of hydrologic and climate systems. Implicit AVS (i.e., Model Predictive Control) and Approximation in Policy Space methods emerge as the best options to overcome the curse of modeling, even though the exogenous information used for conditioning the operating policy is generally limited to observed or predicted inflow to the reservoir. Lastly, DPS applications are the most effective in dealing with the curse of

multiple objectives; studies adopting DPS methods consider, on average, almost three objectives and some works scale beyond 6, also allowing nontime separable as well as mixed objective function formulations. Our analysis also suggests two promising options for addressing the challenges posed by conflicting human pressure and a changing world: enlarging the set of information especially from the hydrologic and climate state contributes in better informing reservoir operations under increasingly variable hydroclimatic conditions, and combining control policy design methods with emulation modeling allows isolating key decision-relevant processes and accommodates more realistic system representations especially in terms of climate and hydrologic dynamics.

Finally, an unresolved challenge affecting most methods reviewed in this paper is their limited uptake by practitioners. Although research efforts on more and more advanced approaches for optimizing reservoir operations combined with the growing tendency of developing open-source tools (see Table 1), repeated surveys (e.g., Pianosi et al., 2020; Rogers & Fiering, 1986; Simonovic, 1992; Teegavarapu & Simonovic, 2001; Whateley et al., 2014) find low adoption of these methods in practice (for a recent review focused on this point, see Dobson et al., 2019). Operators generally reject optimization models to directly inform actual real-time operations and prefer simpler tools, such as rule curves (Loucks & Van Beek, 2017). Among the factors that hinder the uptake of research results is the black box perception of system models, which make their recommendations difficult to understand and explain (Castelvecchi, 2016). This issue is particularly relevant for APS approaches relying on policy parameterized as nonlinear approximating networks, where ongoing research efforts are exploring the possibility of approximating the optimized policy with if/then/else-based rules fit with machine-learning algorithms, such as Classification and Regression Trees.

Acknowledging the resistance of practitioners to use sophisticated optimization methods, the assumption that simplified rules are valid alternatives to optimization-based policies in complex multireservoir systems operated for meeting multiple objectives over diverse time scales is still to be demonstrated via proper numerical benchmarking experiments. Besides, traditional rule curves that specify a target storage trajectory are also somehow “opaque” policies as they do not completely specify how the operator should track such reference. As a consequence, opening the black box of optimization models may be the best option to facilitate the transition from static rule curves to dynamic closed-loop operating policies. Indeed, recent studies suggest that human decisions are often outperformed by optimization tools (e.g., Fraternali et al., 2012; Parameswaran et al., 2011; Quinn & Bederson, 2011), but few people enjoy being replaced by automatic devices or algorithms (Yates et al., 2003). We therefore recommend a hybrid situation as suggested in psychological studies (Kahneman & Klein, 2009), where optimized operating policies remain under the reservoir operator supervision to ensure continuous monitoring of their performance and of relevant changes in the environment. This transition is not going to reduce the key role of water operators as the supervision of the optimized policies will remain a challenging task due to the automation bias, where human operators risk becoming passive when automatic tools are adopted (Skitka et al., 1999).

Data Availability Statement

No data, models, or code were generated or used during this study.

Acknowledgments

The authors would like to thank Prof. Demetris Koutsoyiannis for his inspiring keynote presentation at the first workshop on control methods for water resource systems. Co-author P.M. Reed was supported by the National Science Foundation INFEWS/T2 program, award #1639268. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding entities. Open Access Funding provided by Politecnico di Milano within the CRUI-CARE Agreement.

References

- Adams, L., Lund, J., Moyle, P., Quiñones, R., Herman, J., & O'Rear, T. (2017). Environmental hedging: A theory and method for reconciling reservoir operations for downstream ecology and water supply. *Water Resources Research*, 53(9), 7816–7831. <https://doi.org/10.1002/2016wr020128>
- AghaKouchak, A. (2015). Recognize anthropogenic drought. *Nature*, 524(7566), 409–411. <https://doi.org/10.1038/524409a>
- Ahmadi, M., Bozorg-Haddad, O., & Mariño, M. (2014). Extraction of flexible multi-objective real-time reservoir operation rules. *Water Resources Management*, 28(1), 131–147. <https://doi.org/10.1007/s11269-013-0476-z>
- Ahn, J., & Lyu, S. (2017). Analysis of the coordinated operation of reservoirs and weirs during the management of Nakdong River water resources. *Journal of Water Resources Planning and Management*, 143(8), 04017030. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000780](https://doi.org/10.1061/(asce)wr.1943-5452.0000780)
- Anghileri, D., Castelletti, A., Pianosi, F., Soncini-Sessa, R., & Weber, E. (2013). Optimizing watershed management by coordinated operation of storing facilities. *Journal of Water Resources Planning and Management*, 139(5), 492–500. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000313](https://doi.org/10.1061/(asce)wr.1943-5452.0000313)
- Anghileri, D., Voisin, N., Castelletti, A., Pianosi, F., Nijssen, B., & Lettenmaier, D. (2016). Value of long-term streamflow forecasts to reservoir operations for water supply in snow-dominated river catchments. *Water Resources Research*, 52(6), 4209–4225. <https://doi.org/10.1002/2015wr017864>
- Ansar, A., Flyvbjerg, B., Budzier, A., & Lunn, D. (2014). Should we build more large dams? The actual costs of hydropower megaproject development. *Energy Policy*, 69, 43–56. <https://doi.org/10.1016/j.enpol.2013.10.069>
- Asher, M., Croke, B., Jakeman, A., & Peeters, L. (2015). A review of surrogate models and their application to groundwater modeling. *Water Resources Research*, 51(8), 5957–5973. <https://doi.org/10.1002/2015wr016967>

- Ashofteh, P.-S., Bozorg-Haddad, O., & Loáiciga, H. (2015). Evaluation of climatic-change impacts on multiobjective reservoir operation with multiobjective genetic programming. *Journal of Water Resources Planning and Management*, 141(11), 04015030. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000540](https://doi.org/10.1061/(asce)wr.1943-5452.0000540)
- Ashrafi, S., & Dariane, A. (2017). Coupled operating rules for optimal operation of multi-reservoir systems. *Water Resources Management*, 31(14), 4505–4520. <https://doi.org/10.1007/s11269-017-1762-y>
- Bahrami, M., Bozorg-Haddad, O., & Chu, X. (2017). Application of cat swarm optimization algorithm for optimal reservoir operation. *Journal of Irrigation and Drainage Engineering*, 144(1), 04017057.
- Bauer, P., Thorpe, A., & Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525(7567), 47–55. <https://doi.org/10.1038/nature14956>
- Bellman, R. (1957). *Dynamic programming*. Princeton University Press.
- Bennett, S. (1996). A brief history of automatic control. *IEEE Control Systems Magazine*, 16(3), 17–25.
- Benson, R. (2016). Reviewing reservoir operations: Can federal water projects adapt to change. *Columbia Journal of Environmental Law*, 42, 353.
- Bertoni, F., Castelletti, A., Giuliani, M., & Reed, P. M. (2019). Discovering dependencies, trade-offs, and robustness in joint dam design and operation: An ex-post assessment of the Kariba Dam. *Earth's Future*, 7(12), 1367–1390. <https://doi.org/10.1029/2019EF001235>
- Bertsekas, D. (1976). *Dynamic programming and stochastic control*. Academic Press.
- Bertsekas, D. (2005). Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4–5). <https://doi.org/10.3166/ejc.11.310-334>
- Bertsekas, D. (2019). *Reinforcement learning and optimal control*. Athena Scientifi.
- Beyer, H.-G., & Sendhoff, B. (2007). Robust optimization—a comprehensive survey. *Computer Methods in Applied Mechanics and Engineering*, 196(33–34), 3190–3218. <https://doi.org/10.1016/j.cma.2007.03.003>
- Biglarbeigi, P., Giuliani, M., & Castelletti, A. (2018). Partitioning the impacts of streamflow and evaporation uncertainty on the operations of multipurpose reservoirs in arid regions. *Journal of Water Resources Planning and Management*, 144(7), 05018008. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000945](https://doi.org/10.1061/(asce)wr.1943-5452.0000945)
- Billington, D., & Jackson, D. (2017). *Big dams of the new deal era: A confluence of engineering and politics*. University of Oklahoma Press.
- Bizzi, S., Dinh, Q., Bernardi, D., Denaro, S., Schippa, L., & Soncini-Sessa, R. (2015). On the control of riverbed incision induced by run-of-river power plant. *Water Resources Research*, 51(7), 5023–5040. <https://doi.org/10.1002/2014wr016237>
- Bozorg-Haddad, O., Afshar, A., & Mariño, M. (2008). Honey-bee mating optimization (HBMO) algorithm in deriving optimal operation rules for reservoirs. *Journal of Hydroinformatics*, 10(3), 257–264. <https://doi.org/10.2166/hydro.2008.018>
- Bozorg-Haddad, O., Hosseini-Moghari, S.-M., & Loáiciga, H. (2015). Biogeography-based optimization algorithm for optimal operation of reservoir systems. *Journal of Water Resources Planning and Management*, 142(1), 04015034.
- Brodeur, Z. P., Herman, J. D., & Steinschneider, S. (2020). Bootstrap aggregation and cross-validation methods to reduce overfitting in reservoir control policy search. *Water Resources Research*, 56(8), e2020WR027184. <https://doi.org/10.1029/2020wr027184>
- Brunet, G., Jones, S., & Ruti, P. (2015). *Seamless prediction of the Earth system: From minutes to months*. World Meteorological Organization.
- Cash, D., Clark, W., Alcock, F., Dickson, N., Eckley, N., Guston, D., et al. (2003). Knowledge systems for sustainable development. *Proceedings of the National Academy of Sciences*, 100(14), 8086–8091. <https://doi.org/10.1073/pnas.1231332100>
- Castelletti, A., Corani, G., & Weber, E. (2002). Reinforcement learning in the operational management of a water system. In *Proceedings of the ifac workshop on modeling and control in environmental issues* (pp. 325–330).
- Castelletti, A., de Rigo, D., Rizzoli, A., Soncini-Sessa, R., & Weber, E. (2007). Neuro-dynamic programming for designing water reservoir network management policies. *Control Engineering Practice*, 15(8), 1031–1038. <https://doi.org/10.1016/j.conengprac.2006.02.011>
- Castelletti, A., Galelli, S., Ratto, M., Soncini-Sessa, R., & Young, P. (2012). A general framework for dynamic emulation modelling in environmental problems. *Environmental Modelling & Software*, 34, 5–18. <https://doi.org/10.1016/j.envsoft.2012.01.002>
- Castelletti, A., Galelli, S., Restelli, M., & Soncini-Sessa, R. (2010). Tree-based reinforcement learning for optimal water reservoir operation. *Water Resources Research*, 46(W09507). <https://doi.org/10.1029/2009wr008898>
- Castelletti, A., Galelli, S., Restelli, M., & Soncini-Sessa, R. (2012). Data-driven dynamic emulation modelling for the optimal management of environmental systems. *Environmental Modelling & Software*, 34, 30–43. <https://doi.org/10.1016/j.envsoft.2011.09.003>
- Castelletti, A., Pianosi, F., Quach, X., & Soncini-Sessa, R. (2012). Assessing water reservoirs management and development in Northern Vietnam. *Hydrology and Earth System Sciences*, 16(1), 189–199. <https://doi.org/10.5194/hess-16-189-2012>
- Castelletti, A., Pianosi, F., & Restelli, M. (2013). A multiobjective reinforcement learning approach to water resources systems operation: Pareto Frontier approximation in a single run. *Water Resources Research*, 49, 3476–3486. <https://doi.org/10.1002/wrcr.20295>
- Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2008). Receding horizon control for water resources management. *Applied Mathematics and Computation*, 204(2), 621–631. <https://doi.org/10.1016/j.amc.2008.05.044>
- Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2012). Stochastic and robust control of water resource systems: Concepts, methods and applications. In *System identification, environmental modelling, and control system design* (pp. 383–401). Springer. https://doi.org/10.1007/978-0-85729-974-1_19
- Castelletti, A., Yajima, H., Giuliani, M., Soncini-Sessa, R., & Weber, E. (2014). Planning the optimal operation of a multi-outlet water reservoir with water quality and quantity targets. *Journal of Water Resources Planning and Management*, 140, 496–510. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000348](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000348)
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nature News*, 538(7623), 20. <https://doi.org/10.1038/538020a>
- Celeste, A., & Billib, M. (2009). Evaluation of stochastic reservoir operation optimization models. *Advances in Water Resources*, 32(9), 1429–1443. <https://doi.org/10.1016/j.advwatres.2009.06.008>
- Celeste, A., & El-Shafie, A. (2018). Assessment of stochastic operation optimization for reservoirs of contrasting scales. *Water Resources Management*, 32(11), 3751–3763. <https://doi.org/10.1007/s11269-018-2017-2>
- Chaves, P., & Kojiri, T. (2007). Deriving reservoir operational strategies considering water quantity and quality objectives by stochastic fuzzy neural networks. *Advances in Water Resources*, 30(5), 1329–1341. <https://doi.org/10.1016/j.advwatres.2006.11.011>
- Chen, V. C., Ruppert, D., & Shoemaker, C. A. (1999). Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Operations Research*, 47(1), 38–53. <https://doi.org/10.1287/opre.47.1.38>
- Cheng, C., Wu, H., Wu, X., Shen, J., & Wang, J. (2017). Power generation scheduling for integrated large and small hydropower plant systems in Southwest China. *Journal of Water Resources Planning and Management*, 143(8), 04017027. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000771](https://doi.org/10.1061/(asce)wr.1943-5452.0000771)
- Cheng, X., Tang, Y.-L., Liu, J., Zhong, H., Li, W.-W., & Wu, Y.-L. (2020). Ultrashort-term scheduling of interbasin cascaded hydropower plants to rapidly balance the load demand. *IEEE Access*, 8, 32737–32756. <https://doi.org/10.1109/ACCESS.2020.2973680>
- Clark, E. (1950). New York control curves. *Journal of the American Water Works Association*, 42(9), 823–827.

- Clark, E. (1956). Impounding reservoirs. *Journal of the American Water Works Association*, 48(4), 349–354.
- Cloke, H., Pappenberger, F., Smith, P., & Wetterhall, F. (2017). How do i know if i've improved my continental scale flood early warning system. *Environmental Research Letters*, 12(4), 044006. <https://doi.org/10.1088/1748-9326/aa625a>
- Cohen, J., & Herman, J. (2021). Dynamic adaptation of water resources systems under uncertainty by learning policy structure and indicators. *Water Resources Research*, 57, e2021WR030433. <https://doi.org/10.1029/2021WR030433>
- Cohon, J. L., & Marks, D. (1975). A review and evaluation of multiobjective programming techniques. *Water Resources Research*, 11(2), 208–220. <https://doi.org/10.1029/wr011i002p00208>
- Côté, P., & Leconte, R. (2015). Comparison of stochastic optimization algorithms for hydropower reservoir operation with ensemble streamflow prediction. *Journal of Water Resources Planning and Management*, 142(2), 04015046.
- Crochemore, L., Ramos, M., & Pechlivanidis, I. (2020). Can continental models convey useful seasonal hydrologic information at the catchment scale? *Water Resources Research*, 56. <https://doi.org/10.1029/2019wr025700>
- Dai, A. (2011). Drought under global warming: A review. *Wiley Interdisciplinary Reviews: Climate Change*, 2(1), 45–65. <https://doi.org/10.1002/wcc.81>
- Dang, T., Chowdhury, A., & Galelli, S. (2020). On the representation of water reservoir storage and operations in large-scale hydrological models: Implications on model parameterization and climate change impact assessments. *Hydrology and Earth System Sciences*, 24, 397–416. <https://doi.org/10.5194/hess-24-397-2020>
- Davidsen, C., Pereira-Cardenal, S., Liu, S., Mo, X., Rosbjerg, D., & Bauer-Gottwein, P. (2014). Using stochastic dynamic programming to support water resources management in the ziya river basin, china. *Journal of Water Resources Planning and Management*, 141(7), 04014086.
- Deisenroth, M., Neumann, G., & Peters, J. (2011). A survey on policy search for robotics. In *Foundations and trends in robotics* (Vol. 2, pp. 1–142). <https://doi.org/10.1561/23000000021>
- Delipetrev, B., Jonoski, A., & Solomatine, D. (2015). A novel nested dynamic programming (ndp) algorithm for multipurpose reservoir optimization. *Journal of Hydroinformatics*, 17(4). <https://doi.org/10.2166/hydro.2015.066>
- Denaro, S., Anghileri, D., Giuliani, M., & 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. <https://doi.org/10.1016/j.advwatres.2017.02.012>
- Denaro, S., Castelletti, A., Giuliani, M., & Characklis, G. W. (2018). Fostering cooperation in power asymmetrical water systems by the use of direct release rules and index-based insurance schemes. *Advances in Water Resources*, 115, 301–314. <https://doi.org/10.1016/j.advwatres.2017.09.021>
- Desreumaux, Q., Côté, P., & Leconte, R. (2014). Role of hydrologic information in stochastic dynamic programming: A case study of the Kemano hydropower system in British Columbia. *Canadian Journal of Civil Engineering*, 41(9), 839–844. <https://doi.org/10.1139/cjce-2013-0370>
- Desreumaux, Q., Côté, P., & Leconte, R. (2018). Comparing model-based and model-free streamflow simulation approaches to improve hydropower reservoir operations. *Journal of Water Resources Planning and Management*, 144(3), 05018002. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000860](https://doi.org/10.1061/(asce)wr.1943-5452.0000860)
- Dias, N., Pereira, M., & Kelman, J. (1985). Optimization of flood control and power generation requirements in a multi-purpose reservoir. *IFAC Proceedings Volumes*, 18(7), 101–104. (IFAC Symposium on Planning and Operation of Electric Energy Systems., Rio de Janeiro, Brazil, 22-25 July). [https://doi.org/10.1016/S1474-6670\(17\)60421-X](https://doi.org/10.1016/S1474-6670(17)60421-X)
- Ding, W., Zhang, C., Cai, X., Li, Y., & Zhou, H. (2017). Multiobjective hedging rules for flood water conservation. *Water Resources Research*, 53(3), 1963–1981. <https://doi.org/10.1002/2016wr019452>
- Dobson, B., Wagener, T., & Pianosi, F. (2019). An argument-driven classification and comparison of reservoir operation optimization methods. *Advances in Water Resources*, 128, 74–86. <https://doi.org/10.1016/j.advwatres.2019.04.012>
- Draper, A., & Lund, J. (2004). Optimal Hedging and Carryover storage value. *Journal of Water Resources Planning and Management*, 130(1), 83–87. [https://doi.org/10.1061/\(ASCE\)0733-9496\(2004\)130:1\(83\)](https://doi.org/10.1061/(ASCE)0733-9496(2004)130:1(83))
- Erfani, T., Binions, O., & Harou, J. (2015). Protecting environmental flows through enhanced water licensing and water markets. *Hydrology and Earth System Sciences*, 19(2), 675–689. <https://doi.org/10.5194/hess-19-675-2015>
- Ernst, D., Geurts, P., & Wehenkel, L. (2005). Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6, 503–556.
- Escriva-Bou, A., Lund, J., Pulido-Velazquez, M., Hui, R., & Medellín-Azuara, J. (2018). Developing a water-energy-GHG emissions modeling framework: Insights from an application to California's water system. *Environmental Modelling & Software*, 109, 54–65. <https://doi.org/10.1016/j.envsoft.2018.07.011>
- Esogbue, A. (1989). Dynamic programming and water resources: Origins and interconnections. In *Dynamic programming for optimal water resources systems analysis*. Prentice-Hall.
- Faber, B., & Stedinger, J. (2001). Reservoir optimization using sampling SDP with ensemble streamflow prediction (ESP) forecasts. *Journal of Hydrology*, 249(1), 113–133. [https://doi.org/10.1016/S0022-1694\(01\)00419-X](https://doi.org/10.1016/S0022-1694(01)00419-X)
- Fallah-Mehdipour, E., Bozorg-Haddad, O., Alimohammadi, S., & Loáiciga, H. (2015). Development of real-time conjunctive use operation rules for aquifer-reservoir systems. *Water Resources Management*, 29(6), 1887–1906. <https://doi.org/10.1007/s11269-015-0917-y>
- Feng, M., Liu, P., Guo, S., Shi, L., Deng, C., & Ming, B. (2017). Deriving adaptive operating rules of hydropower reservoirs using time-varying parameters generated by the EnKF. *Water Resources Research*, 53(8), 6885–6907. <https://doi.org/10.1002/2016wr020180>
- Feng, Z.-K., Niu, W.-J., Zhang, R., Wang, S., & Cheng, C.-T. (2019). Operation rule derivation of hydropower reservoir by k-means clustering method and extreme learning machine based on particle swarm optimization. *Journal of hydrology*, 576, 229–238. <https://doi.org/10.1016/j.jhydrol.2019.06.045>
- Ficchi, A., Raso, L., Dorchie, D., Pianosi, F., Malaterre, P., Van Overloop, P., & 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. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000571](https://doi.org/10.1061/(asce)wr.1943-5452.0000571)
- Fleming, P., Purshouse, R., & Lygoe, R. (2005). Many-objective optimization: An engineering design perspective. In *Proceedings of the third international conference on evolutionary multi-criterion optimization* (pp. 14–32). Guanajuato. https://doi.org/10.1007/978-3-540-31880-4_2
- Formentin, S., Karimi, A., & Savaresi, S. (2013). Optimal input design for direct data-driven tuning of model-reference controllers. *Automatica*, 49(6), 1874–1882. <https://doi.org/10.1016/j.automatica.2013.02.054>
- Foster, B., Kern, J., & Characklis, G. (2015). Mitigating hydrologic financial risk in hydropower generation using index-based financial instruments. *Water Resources and Economics*, 10, 45–67. <https://doi.org/10.1016/j.wre.2015.04.001>
- Foufoula-Georgiou, E., & Kitanidis, P. K. (1988). Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems. *Water resources research*, 24(8), 1345–1359. <https://doi.org/10.1029/wr024i008p01345>
- Fraternali, P., Castelletti, A., Soncini-Sessa, R., Ruiz, C., & Rizzoli, A. (2012). Putting humans in the loop: Social computing for Water Resources Management. *Environmental Modelling & Software*, 37, 68–77. <https://doi.org/10.1016/j.envsoft.2012.03.002>

- Galelli, S., Castelletti, A., & Goedbloed, A. (2015). High-performance integrated control of water quality and quantity in urban water reservoirs. *Water Resources Research*, 51(11), 9053–9072. <https://doi.org/10.1002/2015wr017595>
- Galelli, S., Goedbloed, A., Schwanenberg, D., & van Overloop, P. (2014). Optimal real-time operation of multi-purpose urban reservoirs: A case study in Singapore. *Journal of Water Resources Planning and Management*, 140(4), 511–523. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000342](https://doi.org/10.1061/(asce)wr.1943-5452.0000342)
- Galelli, S., & Soncini-Sessa, R. (2010). Combining metamodeling and stochastic dynamic programming for the design of reservoir release policies. *Environmental Modeling & Software*, 25(2), 209–222. <https://doi.org/10.1016/j.envsoft.2009.08.001>
- Gass, S., & Saaty, T. (1955). Parametric objective function - Part II. *Operations Research*, 3, 316–319. <https://doi.org/10.1287/opre.3.4.395>
- Georgakakos, A. (1989). Extended linear quadratic Gaussian control: Further extensions. *Water Resources Research*, 25(2), 191–201. <https://doi.org/10.1029/wr025i002p00191>
- Georgakakos, A., Yao, H., Mullusky, M., & Georgakakos, K. (1998). Impacts of climate variability on the operational forecast and management of the upper Des Moines River basin. *Water Resources Research*, 34(4), 799–821. <https://doi.org/10.1029/97wr03135>
- Geressu, R. T., & Harou, J. J. (2015). Screening reservoir systems by considering the efficient trade-offs? Informing infrastructure investment decisions on the blue Nile. *Environmental Research Letters*, 10(12), 125008. <https://doi.org/10.1088/1748-9326/10/12/125008>
- Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, 63(1), 3–42. <https://doi.org/10.1007/s10994-006-6226-1>
- Ghorbanidehno, H., Kokkinaki, A., Kitanidis, P., & Darve, E. (2017). Optimal estimation and scheduling in aquifer management using the rapid feedback control method. *Advances in Water Resources*, 110, 310–318. <https://doi.org/10.1016/j.advwatres.2017.10.011>
- Giuliani, M., Anghileri, D., Vu, P., Castelletti, A., & Soncini-Sessa, R. (2016). Large storage operations under climate change: Expanding uncertainties and evolving tradeoffs. *Environmental Research Letters*, 11(3), 035009. <https://doi.org/10.1088/1748-9326/11/3/035009>
- Giuliani, M., & Castelletti, A. (2013). Assessing the value of cooperation and information exchange in large water resources systems by agent-based optimization. *Water Resources Research*, 49, 3912–3926. <https://doi.org/10.1002/wrcr.20287>
- Giuliani, M., & Castelletti, A. (2016). Is robustness really robust? How different definitions of robustness impact decision-making under climate change. *Climatic Change*, 135, 409–424. <https://doi.org/10.1007/s10584-015-1586-9>
- Giuliani, M., Castelletti, A., Fedorov, R., & Fraternali, P. (2016). Using crowdsourced web content for informing water systems operations in snow-dominated catchments. *Hydrology and Earth System Sciences*, 20, 5049–5062. <https://doi.org/10.5194/hess-20-5049-2016>
- Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P. (2016). 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), 04015050. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000570](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000570)
- Giuliani, M., Crochemore, L., Pechlivanidis, I., & Castelletti, A. (2020). From skill to value: Isolating the influence of end user behaviour on seasonal forecast assessment. *Hydrology and Earth System Sciences*, 24, 5891–5902. <https://doi.org/10.5194/hess-24-5891-2020>
- Giuliani, M., Galelli, S., & 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. <https://doi.org/10.1016/j.envsoft.2014.02.011>
- Giuliani, M., Herman, J., Castelletti, A., & Reed, P. (2014). Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management. *Water Resources Research*, 50, 3355–3377. <https://doi.org/10.1002/2013WR014700>
- Giuliani, M., Li, Y., Castelletti, A., & Gandolfi, C. (2016). A coupled human-natural systems analysis of irrigated agriculture under changing climate. *Water Resources Research*, 52(9), 6928–6947. <https://doi.org/10.1002/2016wr019363>
- Giuliani, M., Pianosi, F., & 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. <https://doi.org/10.1002/2015WR017044>
- Giuliani, M., Quinn, J. D., Herman, J. D., Castelletti, A., & 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. <https://doi.org/10.1109/TCST.2017.2705162>
- Giuliani, M., Zaniolo, M., Castelletti, A., Davoli, G., & 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. <https://doi.org/10.1029/2019WR025035>
- Gneiting, T., & Raftery, A. (2005). Weather forecasting with ensemble methods. *Science*, 310(5746), 248–249. <https://doi.org/10.1126/science.1115255>
- Gonzales, P., & Ajami, N. K. (2017). The changing water cycle: Impacts of an evolving supply and demand landscape on urban water reliability in the bay area. *Wiley Interdisciplinary Reviews: Water*, 4(6), e1240. <https://doi.org/10.1002/wat2.1240>
- Goor, Q., Halleux, C., Mohamed, Y., & Tilmant, A. (2010). Optimal operation of a multipurpose multireservoir system in the Eastern Nile River Basin. *Hydrology and Earth System Sciences*, 14(10), 1895–1908. <https://doi.org/10.5194/hess-14-1895-2010>
- Goor, Q., Kelman, R., & Tilmant, A. (2011). Optimal multipurpose-multireservoir operation model with variable productivity of hydropower plants. *Journal of Water Resources Planning and Management*, 137(3), 258–267. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000117](https://doi.org/10.1061/(asce)wr.1943-5452.0000117)
- Grill, G., Lehner, B., Thieme, M., Geenen, B., Tickner, D., Antonelli, F., et al. (2019). Mapping the world's free-flowing rivers. *Nature*, 569(7755), 215–221. <https://doi.org/10.1038/s41586-019-1111-9>
- Guariso, G., Rinaldi, S., & Zielinski, P. (1984). The value of information in reservoir management. *Applied Mathematics and Computation*, 15(2), 165–184. [https://doi.org/10.1016/0096-3003\(84\)90015-8](https://doi.org/10.1016/0096-3003(84)90015-8)
- Haimes, Y., Lasdon, L., & Wismer, D. (1971). On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE Transactions on Systems, Man and Cybernetics*, 1, 296–297.
- Hall, J., Grey, D., Garrick, D., Fung, F., Brown, C., Dadson, S., & Sadoff, C. (2014). Coping with the curse of freshwater variability. *Science*, 346(6208), 429–430. <https://doi.org/10.1126/science.1257890>
- Hall, W., & Buras, N. (1961). The dynamic programming approach to water-resources development. *Journal of Geophysical Research*, 66(2), 517–520. <https://doi.org/10.1029/jz066i002p00517>
- Hanasaki, N., Kanae, S., & Oki, T. (2006). A reservoir operation scheme for global river routing models. *Journal of Hydrology*, 327(1), 22–41. <https://doi.org/10.1016/j.jhydrol.2005.11.011>
- Hapuarachchi, H., Wang, Q., & Pagano, T. (2011). A review of advances in flash flood forecasting. *Hydrological Processes*, 25(18), 2771–2784. <https://doi.org/10.1002/hyp.8040>
- Hart, J. K., & Martinez, K. (2006). Environmental sensor networks: A revolution in the earth system science? *Earth-Science Reviews*, 78(3), 177–191. <https://doi.org/10.1016/j.earscirev.2006.05.001>
- Hayes, C. F., Rădulescu, R., Bargiacchi, E., Källström, J., Macfarlane, M., Reymond, M., et al. (2021). A practical guide to multi-objective reinforcement learning and planning. arXiv preprint arXiv:2103.09568.
- Hazen, A. (1914). Storage to be provided in impounding reservoirs for municipal water supply. *Transactions of the American Society of Civil Engineers*, 77(1539). [https://doi.org/10.1061/\(taceat.0002563](https://doi.org/10.1061/(taceat.0002563)

- Herman, J., & Giuliani, M. (2018). Policy tree optimization for threshold-based water resources management over multiple timescales. *Environmental Modelling & Software*, 99, 39–51. <https://doi.org/10.1016/j.envsoft.2017.09.016>
- Hou, Z.-S., & Wang, Z. (2013). From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, 235, 3–35. <https://doi.org/10.1016/j.ins.2012.07.014>
- Howard, C. (1992). Experience with probabilistic forecasts for cumulative stochastic optimization. In *Proc. American water resources association symposium*.
- Howard, C. (1999). Death to rule curves. In *Wrpm'd'99: Preparing for the 21st century* (pp. 1–5). [https://doi.org/10.1061/40430\(1999\)232](https://doi.org/10.1061/40430(1999)232)
- Howitt, R., MacEwan, D., Medellin-Azuara, J., Lund, J., & Sumner, D. (2017). *Economic analysis of the 2015 drought for California agriculture (Tech. Rep.)*.
- Hu, T., Zhang, X.-Z., Zeng, X., & Wang, J. (2016). A two-step approach for analytical optimal hedging with two triggers. *Water*, 8(2), 52. <https://doi.org/10.3390/w8020052>
- Hui, R., & Lund, J. R. (2015). Flood storage allocation rules for parallel reservoirs. *Journal of Water Resources Planning and Management*, 141(5). [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000469](https://doi.org/10.1061/(asce)wr.1943-5452.0000469)
- Hui, R., Lund, J. R., Zhao, J., & Zhao, T. (2016). Optimal flood pre-releases—Flood hedging for a single reservoir. *Water Resources Management*, 30(5113). <https://doi.org/10.1007/s11269-016-1472-x>
- Hurford, A., & Harou, J. (2014). Balancing ecosystem services with energy and food security—assessing trade-offs for reservoir operation and irrigation investment in Kenya's Tana basin. *Hydrology and Earth System Sciences*, 11(1), 1343–1388.
- Johnson, S., Stedinger, J., Shoemaker, C., Li, Y., & Tejada-Guibert, J. (1993). Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Operations Research*, 41(3), 484–500. <https://doi.org/10.1287/opre.41.3.484>
- Johnson, S., Stedinger, J., & Staschus, K. (1991). Heuristic operating policies for reservoir system simulation. *Water Resources Research*, 27(5), 673–685. <https://doi.org/10.1029/91wr00320>
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6), 515–526. <https://doi.org/10.1037/a0016755>
- Kalyanakrishnan, S., & Stone, P. (2007). Batch reinforcement learning in a complex domain. In *Proceedings of the 6th international joint conference on autonomous agents and multiagent systems* (p. 94). <https://doi.org/10.1145/1329125.1329241>
- Karami, F., & Dariane, A. (2018). Many-objective multi-scenario algorithm for optimal reservoir operation under future uncertainties. *Water Resources Management*, 32(12), 3887–3902. <https://doi.org/10.1007/s11269-018-2025-2>
- Karamouz, M., Ahmadi, A., & Moridi, A. (2009). Probabilistic reservoir operation using Bayesian stochastic model and support vector machine. *Advances in Water Resources*, 32(11), 1588–1600. <https://doi.org/10.1016/j.advwatres.2009.08.003>
- Karamouz, M., & Houck, M. H. (1987). Comparison of stochastic and deterministic dynamic programming for reservoir operating rule generation. *JAWRA Journal of the American Water Resources Association*, 23(1), 1–9. <https://doi.org/10.1111/j.1752-1688.1987.tb00778.x>
- Karimanzira, D., Schwabenberg, D., Allen, C., & Barton, S. (2016). Short-term hydropower optimization and assessment of operational flexibility. *Journal of Water Resources Planning and Management*, 142(2), 04015048. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000577](https://doi.org/10.1061/(asce)wr.1943-5452.0000577)
- Kasprzyk, J., Reed, P., Characklis, G., & Kirsch, B. (2012). Many-objective de Novo water supply portfolio planning under deep uncertainty. *Environmental Modelling & Software*, 34(0), 87–104. <https://doi.org/10.1016/j.envsoft.2011.04.003>
- Kelman, J., Cooper, L. A., Hsu, E., & Yuan, S.-Q. (1988). Use of probabilistic constraints in reservoir operation policies with sampling stochastic dynamic programming. In *Proceedings of 3rd water resources operations management workshop: Computerized decision report systems for water managers* (pp. 266–279).
- Kelman, J., Stedinger, J., Cooper, L., Hsu, E., & Yuan, S. (1990). Sampling stochastic dynamic programming applied to reservoir operation. *Water Resources Research*, 26(3), 447–454. <https://doi.org/10.1029/wr026i003p00447>
- Kerachian, R., & Karamouz, M. (2006). Optimal reservoir operation considering the water quality issues: A stochastic conflict resolution approach. *Water Resources Research*, 42(W12401). <https://doi.org/10.1029/2005wr004575>
- Kim, Y.-O., Eum, H.-I., Lee, E.-G., & Ko, I. H. (2007). Optimizing operational policies of a Korean multireservoir system using sampling stochastic dynamic programming with ensemble streamflow prediction. *Journal of Water Resources Planning and Management*, 133(1), 4–14. [https://doi.org/10.1061/\(asce\)0733-9496\(2007\)133:1\(4\)](https://doi.org/10.1061/(asce)0733-9496(2007)133:1(4))
- Kistenmacher, M., & Georgakakos, A. (2015). Assessment of reservoir system variable forecasts. *Water Resources Research*, 51(5), 3437–3458. <https://doi.org/10.1002/2014wr016564>
- Kleijnen, J. (2008). Response surface methodology for constrained simulation optimization: An overview. *Simulation Modelling Practice and Theory*, 16(1), 50–64. <https://doi.org/10.1016/j.simpat.2007.10.001>
- Klemeš, V. (1987). One hundred years of applied storage reservoir theory. *Water Resources Management*, 1(3), 159–175.
- Kollet, S., Maxwell, R., Woodward, C., Smith, S., Vanderborght, J., Vereecken, H., & Simmer, C. (2010). Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources. *Water Resources Research*, 46(4). <https://doi.org/10.1029/2009wr008730>
- Koutsyiannis, D., & Economou, A. (2003a). Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems. *Water Resources Research*, 39(6), 1170–1187. <https://doi.org/10.1029/2003wr002148>
- Koutsyiannis, D., & Economou, A. (2003b). *Mathematical study of the concept of equivalent reservoir of a reservoir system (Tech. Rep.)*. National Technical University of Athens.
- Kwakkel, J., Eker, S., & Pruyt, E. (2016). How robust is a robust policy? Comparing alternative robustness metrics for robust decision-making. In *Robustness analysis in decision aiding, optimization, and analytics* (pp. 221–237). Springer. https://doi.org/10.1007/978-3-319-33121-8_10
- Labadie, J. (2004). Optimal operation of multireservoir systems: State-of-the-art review. *Journal of Water Resources Planning and Management*, 130(2), 93–111. [https://doi.org/10.1061/\(asce\)0733-9496\(2004\)130:2\(93\)](https://doi.org/10.1061/(asce)0733-9496(2004)130:2(93))
- Labadie, J. (2005). Closure to optimal operation of multireservoir systems: State-of-the-art review? by John W. Labadie. *Journal of Water Resources Planning and Management*, 131(5), 407–408. [https://doi.org/10.1061/\(asce\)0733-9496\(2005\)131:5\(407\)](https://doi.org/10.1061/(asce)0733-9496(2005)131:5(407))
- Labadie, J. W., & Wan, Y. (2010). Fuzzy optimal control of reservoir-assisted stormwater treatment areas for aquatic ecosystem restoration. *Environmental Modelling & Software*, 25(12), 1692–1701. <https://doi.org/10.1016/j.envsoft.2010.05.001>
- Lamontagne, J. (2015). *Representation of uncertainty and corridor DP for hydropower optimization (Unpublished doctoral dissertation)*. Cornell University.
- Laplace, P. (1951). *A philosophical essays on probabilities*. Dover.
- Lee, J., & Labadie, J. (2007). Stochastic optimization of multireservoir systems via reinforcement learning. *Water resources research*, 43(11). <https://doi.org/10.1029/2006wr005627>
- Lempert, R. J. (2002). A new decision sciences for complex systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3), 7309–7313. <https://doi.org/10.1073/pnas.082081699>

- Lettenmaier, D. P., Alsdorf, D., Dozier, J., Huffman, G. J., Pan, M., & Wood, E. F. (2015). Inroads of remote sensing into hydrologic science during the WRR era. *Water Resources Research*, *51*(9), 7309–7342. <https://doi.org/10.1002/2015wr017616>
- Li, F.-F., Shoemaker, C., Qiu, J., & Wei, J.-H. (2015). Hierarchical multi-reservoir optimization modeling for real-world complexity with application to the Three Gorges system. *Environmental Modelling & Software*, *69*, 319–329. <https://doi.org/10.1016/j.envsoft.2014.11.030>
- Libisch-Lehner, C. P., Nguyen, H. T. T., Taormina, R., Nachtnebel, H. P., & Galelli, S. (2019). On the value of ENSO state for urban water supply system operators: Opportunities, trade-offs, and challenges. *Water Resources Research*, *55*(4), 2856–2875. <https://doi.org/10.1029/2018WR023622>
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., et al. (2015). *Continuous control with deep reinforcement learning*. arXiv preprint arXiv:1509.02971.
- Liu, P., Li, L., Guo, S., Xiong, L., Zhang, W., Zhang, J., & Xu, C.-Y. (2015). Optimal design of seasonal flood limited water levels and its application for the three gorges reservoir. *Journal of Hydrology*, *527*, 1045–1053. <https://doi.org/10.1016/j.jhydrol.2015.05.055>
- Loucks, D., & Sigvaldason, O. (1982). Multiple-reservoir operation in North America. In Z. Kaczmarck, & J. Kindler (Eds.), *The operation of multiple reservoir systems* (pp. 1–103).
- Loucks, D., Stedinger, J., & Shamir, U. (1985). Modelling water resource systems: Issues and experiences. *Civil Engineering Systems*, *2*(4), 223–231. <https://doi.org/10.1080/02630258508970410>
- Loucks, D., van Beek, E., Stedinger, J., Dijkman, J., & Villars, M. (2005). *Water resources systems planning and management: An introduction to methods, models and applications*. UNESCO.
- Loucks, D. P., & Van Beek, E. (2017). *Water resource systems planning and management: An introduction to methods, models, and applications*. Springer.
- Lu, J. Z. (2015). Closing the gap between planning and control: A multiscale mpc cascade approach. *Annual Reviews in Control*, *40*, 3–13. <https://doi.org/10.1016/j.arcontrol.2015.09.016>
- Lund, J. (2000). Derived power production and energy drawdown rules for reservoirs. *Journal of Water Resources Planning and Management*, *126*(2), 108–111. [https://doi.org/10.1061/\(asce\)0733-9496\(2000\)126:2\(108\)](https://doi.org/10.1061/(asce)0733-9496(2000)126:2(108))
- Lund, J., & Guzman, J. (1999). Derived operating rules for reservoirs in series or in parallel. *Journal of Water Resources Planning and Management*, *125*(3), 143–153. [https://doi.org/10.1061/\(asce\)0733-9496\(1999\)125:3\(143\)](https://doi.org/10.1061/(asce)0733-9496(1999)125:3(143))
- Maass, A., Hufschmidt, M., Dorfman, R., Thomas, H., Jr, Marglin, S., & Fair, G. (1962). *Design of water–resource systems*. Harvard University Press Cambridge.
- Mabaya, G., Unami, K., & Fujihara, M. (2017). Stochastic optimal control of agrochemical pollutant loads in reservoirs for irrigation. *Journal of Cleaner Production*, *146*, 37–46. <https://doi.org/10.1016/j.jclepro.2016.05.108>
- Macian-Sorribes, H., & Pulido-Velazquez, M. (2020). Inferring efficient operating rules in multireservoir water resource systems: A review. *Wiley Interdisciplinary Reviews: Water*, *7*(1), e1400.
- Macian-Sorribes, H., Tilmant, A., & Pulido-Velazquez, M. (2017). Improving operating policies of large-scale surface-groundwater systems through stochastic programming. *Water Resources Research*, *53*(2), 1407–1423.
- Maciejowski, J. M. (2002). *Predictive control: With constraints*. Pearson education.
- Madani, K., & Hooshyar, M. (2014). A game theory–reinforcement learning (gt–rl) method to develop optimal operation policies for multi-operator reservoir systems. *Journal of Hydrology*, *519*, 732–742. <https://doi.org/10.1016/j.jhydrol.2014.07.061>
- Mallakpour, I., AghaKouchak, A., & Sadegh, M. (2019). Climate-induced changes in the risk of hydrological failure of major dams in California. *Geophysical Research Letters*, *46*(4), 2130–2139. <https://doi.org/10.1029/2018gl081888>
- Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Sokaert, P. O. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, *36*(6), 789–814. [https://doi.org/10.1016/s0005-1098\(99\)00214-9](https://doi.org/10.1016/s0005-1098(99)00214-9)
- McPhail, C., Maier, H., Kwakkel, J., Giuliani, M., Castelletti, A., & Westra, S. (2018). Robustness metrics: How are they calculated, when should they be used and why do they give different results? *Earth's Future*, *6*(2), 169–191. <https://doi.org/10.1002/2017EF000649>
- Milanese, M., Norton, J., Piet-Lahanier, H., & Walter, É. (1996). *Bounding approaches to system identification*. Springer. <https://doi.org/10.1007/978-1-4757-9545-5>
- Mora, C., Spirandelli, D., Franklin, E., Lynham, J., Kantar, M., Miles, W., et al. (2018). Broad threat to humanity from cumulative climate hazards intensified by greenhouse gas emissions. *Nature Climate Change*, *8*(12), 1062–1071. <https://doi.org/10.1038/s41558-018-0315-6>
- Mortazavi-Naeini, M., Kuczera, G., & Cui, L. (2014). Application of multiobjective optimization to scheduling capacity expansion of urban water resource systems. *Water Resources Research*, *50*(6), 4624–4642. <https://doi.org/10.1002/2013wr014569>
- Müller, J., & Shoemaker, C. (2014). Influence of ensemble surrogate models and sampling strategy on the solution quality of algorithms for computationally expensive black-box global optimization problems. *Journal of Global Optimization*, *60*(2), 123–144. <https://doi.org/10.1007/s10898-014-0184-0>
- Mulligan, M., van Soesbergen, A., & Sáenz, L. (2020). GOODD, a global dataset of more than 38,000 georeferenced dams. *Scientific Data*, *7*(1), 1–8. <https://doi.org/10.1038/s41597-020-0362-5>
- Nalbantis, I., & Koutsoyiannis, D. (1997). A parametric rule for planning and management of multiple-reservoir systems. *Water Resources Research*, *33*(9), 2165–2177. <https://doi.org/10.1029/97wr01034>
- Nandalal, K., & Bogardi, J. (2007). *Dynamic programming based operation of reservoirs: Applicability and limits*. Cambridge University Press.
- National Research Council. (2005). *Getting up to speed: The future of supercomputing*. National Academies Press. <https://doi.org/10.17226/11148>
- Nayak, M., Herman, J., & Steinschneider, S. (2018). Balancing flood risk and water supply in California: Policy search integrating short-term forecast ensembles with conjunctive use. *Water Resources Research*, *54*(10), 7557–7576. <https://doi.org/10.1029/2018WR023177>
- Orlob, G., & Simonovic, S. (1982). Reservoir operation for water quality control. In T. Unny, & E. McBean (Eds.), *Experience in operation of hydrosystems* (pp. 263–285). Water Resources Publications.
- Orlovski, S., Rinaldi, S., & Soncini-Sessa, R. (1983). A min-max approach to storage control problems. *Applied Mathematics and Computation*, *12*(2–3), 237–254. [https://doi.org/10.1016/0096-3003\(83\)90010-3](https://doi.org/10.1016/0096-3003(83)90010-3)
- Ostfeld, A., & Salomons, S. (2005). A hybrid genetic-instance based learning algorithm for CE-QUAL-W2 calibration. *Journal of Hydrology*, *310*, 122–142. <https://doi.org/10.1016/j.jhydrol.2004.12.004>
- Padovani, J., Gorelick, S., Thompson, B., Rozelle, S., & Fendorf, S. (2015). Assessment of human–natural system characteristics influencing global freshwater supply vulnerability. *Environmental Research Letters*, *10*(10), 104014. <https://doi.org/10.1088/1748-9326/10/10/104014>
- Pan, L., Housh, M., Liu, P., Cai, X., & Chen, X. (2015). Robust stochastic optimization for reservoir operation. *Water Resources Research*, *51*(1), 409–429. <https://doi.org/10.1002/2014wr015380>
- Parameswaran, A., Sarma, A., Garcia-Molina, H., Polyzotis, N., & Widom, J. (2011). Human-assisted graph search: It's okay to ask questions. *Proceedings of the VLDB Endowment*, *4*(5), 267–278. <https://doi.org/10.14778/1952376.1952377>

- Pereira, M., & Pinto, L. (1985). Stochastic optimization of a multireservoir hydroelectric system: A decomposition approach. *Water resources research*, 21(6), 779–792. <https://doi.org/10.1029/wr021i006p00779>
- Pereira, M., & Pinto, L. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1–3), 359–375. <https://doi.org/10.1007/bf01582895>
- Pereira-Cardenal, S., Mo, B., Gjelsvik, A., Riegels, N., Arnbjerg-Nielsen, K., & Bauer-Gottwein, P. (2016). Joint optimization of regional water-power systems. *Advances in Water Resources*, 92, 200–207. <https://doi.org/10.1016/j.advwatres.2016.04.004>
- Pianosi, F., Castelletti, A., & Restelli, M. (2013). Tree-based fitted Q-iteration for multi-objective Markov decision processes in water resource management. *Journal of Hydroinformatics*, 15(2), 258–270. <https://doi.org/10.2166/hydro.2013.169>
- Pianosi, F., Dobson, B., & Wagener, T. (2020). Use of reservoir operation optimization methods in practice: Insights from a survey of water resource managers. *Journal of Water Resources Planning and Management*, 146(12), 02520005. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0001301](https://doi.org/10.1061/(asce)wr.1943-5452.0001301)
- Pianosi, F., Quach, X., & Soncini-Sessa, R. (2011). Artificial neural networks and multi objective genetic algorithms for water resources management: An application to the Hoabinh reservoir in Vietnam. In *Proceedings of the 18th ifac world congress*. <https://doi.org/10.3182/20110828-6-it-1002.02208>
- Pianosi, F., & Soncini-Sessa, R. (2009). Real-time management of a multipurpose water reservoir with a heteroscedastic inflow model. *Water Resources Research*, 45(10), W10430. <https://doi.org/10.1029/2008wr007335>
- Pina, J., Tilmant, A., & Côté, P. (2017). Optimizing multireservoir system operating policies using exogenous hydrologic variables. *Water Resources Research*, 53(11), 9845–9859. <https://doi.org/10.1002/2017WR021701>
- Poff, N. L., Brown, C. M., Grantham, T. E., Matthews, J. H., Palmer, M. A., Spence, C. M., et al. (2015). Sustainable water management under future uncertainty with eco-engineering decision scaling. *Nature Climate Change*. <https://doi.org/10.1038/nclimate2765>
- Powell, W. (2007). *Approximate dynamic programming: Solving the curses of dimensionality*. Wiley.
- Powell, W. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, 275(3), 795–821. <https://doi.org/10.1016/j.ejor.2018.07.014>
- Pulido-Velazquez, M., Marques, G., Harou, J., & Lund, J. (2016). Hydroeconomic models as decision support tools for conjunctive management of surface and groundwater. In A. Jakeman (Ed.), *Integrated groundwater management* (pp. 693–710). Springer Nature. https://doi.org/10.1007/978-3-319-23576-9_27
- Quinn, A., & Bederson, B. (2011). Human computation: A survey and taxonomy of a growing field. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1403–1412).
- Quinn, J., Reed, P., Giuliani, M., & Castelletti, A. (2017). Rival framings: A framework for discovering how problem formulation uncertainties shape risk management trade-offs in water resources systems. *Water Resources Research*, 53(8), 7208–7233. <https://doi.org/10.1002/2017WR020524>
- Quinn, J., Reed, P., Giuliani, M., & Castelletti, A. (2019). What is controlling our control rules? Opening the black box of multi-reservoir operating policies using time-varying sensitivity analysis. *Water Resources Research*, 55, 5962–5984. <https://doi.org/10.1029/2018wr024177>
- Quinn, J., Reed, P., & Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environmental Modelling & Software*, 92, 125–141. <https://doi.org/10.1016/j.envsoft.2017.02.017>
- Raso, L., & Malaterre, P. (2016). Combining short-term and long-term reservoir operation using infinite horizon model predictive control. *Journal of Irrigation and Drainage Engineering*, 143(3), B4016002.
- Raso, L., Malaterre, P., & Bader, J. (2017). Effective streamflow process modeling for optimal reservoir operation using stochastic dual dynamic programming. *Journal of Water Resources Planning and Management*, 143(4), 04017003. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000746](https://doi.org/10.1061/(asce)wr.1943-5452.0000746)
- Raso, L., Schwabenberg, D., van de Giesen, N., & van Overloop, P. (2014). Short-term optimal operation of water systems using ensemble forecasts. *Advances in Water Resources*, 71, 200–208. <https://doi.org/10.1016/j.advwatres.2014.06.009>
- Razavi, S., Tolson, B., & Burn, D. (2012). Review of surrogate modeling in water resources. *Water Resources Research*, 48(7). <https://doi.org/10.1029/2011wr011527>
- Reed, P., & Hadka, D. (2014). Evolving many-objective water management to exploit exascale computing. *Water Resources Research*, 50(10), 8367–8373. <https://doi.org/10.1002/2014WR015976>
- Reed, P. M., & Kollat, J. B. (2013). Visual analytics clarify the scalability and effectiveness of massively parallel many-objective optimization: A groundwater monitoring design example. *Advances in Water Resources*, 56, 1–13. <https://doi.org/10.1016/j.advwatres.2013.01.011>
- Regis, R., & Shoemaker, C. (2007). A stochastic radial basis function method for the global optimization of expensive functions. *INFORMS Journal on Computing*, 19(4), 497–509. <https://doi.org/10.1287/ijoc.1060.0182>
- Revelle, C., Loucks, D., & Lynn, W. (1968). Linear programming applied to water quality management. *Water Resources Research*, 4(1), 1–9. <https://doi.org/10.1029/wr004i001p00001>
- Rheinheimer, D., Null, S., & Lund, J. (2014). Optimizing selective withdrawal from reservoirs to manage downstream temperatures with climate warming. *Journal of Water Resources Planning and Management*, 141(4), 04014063.
- Richter, B., Baumgartner, J., Powell, J., & Braun, D. (1996). A method for assessing hydrologic alteration within ecosystems. *Conservation Biology*, 10(4), 1163–1174. <https://doi.org/10.1046/j.1523-1739.1996.10041163.x>
- Rieker, J., & Labadie, J. (2012). An intelligent agent for optimal river-reservoir system management. *Water Resources Research*, 48(9). <https://doi.org/10.1029/2012wr011958>
- Rinaldi, S., Soncini-Sessa, R., Stehfest, H., & Tamura, H. (1979). *Modeling and control of river quality* (Vol. 283). McGraw-Hill.
- Rippl, W. (1883). The capacity of storage reservoirs for water supply. In *Minutes of the proceedings* (Vol. 71, pp. 270–278). Institution of Civil Engineers. <https://doi.org/10.1680/imotp.1883.21797>
- Rogers, P., & Fiering, M. (1986). Use of systems analysis in water management. *Water Resources Research*, 22(9S). <https://doi.org/10.1029/wr022i09sp0146s>
- Rosenstein, M., & Barto, A. (2001). Robot weightlifting by direct policy search. In *International joint conference on artificial intelligence* (Vol. 17, pp. 839–846).
- Rückstieß, T., Sehnke, F., Schaul, T., Wierstra, D., Sun, Y., & Schmidhuber, J. (2010). Exploring parameter space in reinforcement learning. *Paladyn. Journal of Behavioral Robotics*, 1(1), 14–24.
- Rupp, K. (2020). 48 years of microprocessor trend data. *Zenodo*. <https://doi.org/10.5281/zenodo.3947824>
- Saad, M., Turgeon, A., Bigras, P., & Duquette, R. (1994). Learning disaggregation technique for the operation of long-term hydroelectric power systems. *Water Resources Research*, 30(11), 3195–3202. <https://doi.org/10.1029/94wr01731>
- Saad, M., Turgeon, A., & Stedinger, J. R. (1992). Censored-data correlation and principal component dynamic programming. *Water Resources Research*, 28(8), 2135–2140. <https://doi.org/10.1029/92WR00896>
- Salas, J., Tabios, G., & Bartolini, P. (1985). Approaches to multivariate modeling of water resources time series. *JAWRA Journal of the American Water Resources Association*, 21(4), 683–708. <https://doi.org/10.1111/j.1752-1688.1985.tb05383.x>

- Sangiorgio, M., & Guariso, G. (2018). NN-based implicit stochastic optimization of multi-reservoir systems management. *Water*, *10*(3), 303. <https://doi.org/10.3390/w10030303>
- Savage, L. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, *46*(253), 55–67. <https://doi.org/10.1080/01621459.1951.10500768>
- Scattolini, R. (2009). Architectures for distributed and hierarchical Model Predictive Control - A review. *Journal of Process Control*, *19*(5), 723–731. <https://doi.org/10.1016/j.jprocont.2009.02.003>
- Schardong, A., & Simonovic, S. (2015). Coupled self-adaptive multiobjective differential evolution and network flow algorithm approach for optimal reservoir operation. *Journal of Water Resources Planning and Management*, *141*(10), 04015015. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000525](https://doi.org/10.1061/(asce)wr.1943-5452.0000525)
- Schmidhuber, J. (2001). Sequential decision making based on direct search. In R. Sun, & C. Giles (Eds.), *Sequence learning* (Vol. 1828, pp. 213–240). Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-44565-x_10
- Schmitt, R., Bizzi, S., Castelletti, A., & Kondolf, G. (2018). Improved trade-offs of hydropower and sand connectivity by strategic dam planning in the Mekong. *Nature Sustainability*, *1*(2), 96–104. <https://doi.org/10.1038/s41893-018-0022-3>
- Séguin, S., Côté, P., & Audet, C. (2016). Self-scheduling short-term unit commitment and loading problem. *IEEE Transactions on Power Systems*, *31*(1), 133–142. <https://doi.org/10.1109/TPWRS.2014.2383911>
- Shannon, C. (1949). Communication in the presence of noise. *Proceedings of the Institute of Radio Engineering*, *37*(1), 10–21. <https://doi.org/10.1109/jrproc.1949.232969>
- Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, *209*(1), 63–72. <https://doi.org/10.1016/j.ejor.2010.08.007>
- Siade, A., Putti, M., & Yeh, W. (2010). Snapshot selection for groundwater model reduction using proper orthogonal decomposition. *Water Resources Research*, *46*(8). <https://doi.org/10.1029/2009wr008792>
- Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. <https://doi.org/10.1037/h0042769>
- Simonovic, S. (1992). Reservoir systems analysis: Closing gap between theory and practice. *Journal of Water Resources Planning and Management*, *118*(3), 262–280. [https://doi.org/10.1061/\(asce\)0733-9496\(1992\)118:3\(262\)](https://doi.org/10.1061/(asce)0733-9496(1992)118:3(262))
- Skitka, L., Mosier, K., & Burdick, M. (1999). Does automation bias decision-making? *International Journal of Human-Computer Studies*, *51*(5), 991–1006. <https://doi.org/10.1006/ijhc.1999.0252>
- Soncini-Sessa, R., Castelletti, A., & Weber, E. (2007). *Integrated and participatory water resources management: Theory*. Elsevier.
- Soncini-Sessa, R., Cellina, F., Pianosi, F., & Weber, E. (2007). *Integrated and participatory water resources management: Practice*. Elsevier.
- Sovacool, B., Gilbert, A., & Nugent, D. (2014). Risk, innovation, electricity infrastructure and construction cost overruns: Testing six hypotheses. *Energy*, *74*, 906–917. <https://doi.org/10.1016/j.energy.2014.07.070>
- Stedinger, J. R., Faber, B. A., & Lamontagne, J. R. (2013). Developments in stochastic dynamic programming for reservoir operation optimization. In *World environmental and water resources congress 2013: Showcasing the future* (pp. 1266–1278). <https://doi.org/10.1061/9780784412947.125>
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. MIT Press.
- Tan, Q.-F., Wang, X., Wang, H., Wang, C., Lei, X.-H., Xiong, Y.-S., & Zhang, W. (2017). Derivation of optimal joint operating rules for multi-purpose multi-reservoir water-supply system. *Journal of Hydrology*, *551*, 253–264. <https://doi.org/10.1016/j.jhydrol.2017.06.009>
- Teegavarapu, R., & Simonovic, S. (2001). Optimal operation of water resource systems: Trade-offs between modeling and practical solutions. In S. S. Marino (Ed.), *Integrated water resources management* (Vol. 272, pp. 257–263).
- Tejada-Guibert, J., Johnson, S., & Stedinger, J. (1995). The value of hydrologic information in stochastic dynamic programming models of a multireservoir system. *Water Resources Research*, *31*(10), 2571–2579. <https://doi.org/10.1029/95wr02172>
- Thompson, S., Sivapalan, M., Harman, C., Srinivasan, V., Hipsey, M., Reed, P., et al. (2013). Developing predictive insight into changing water systems: Use-inspired hydrologic science for the anthropocene. *Hydrology and Earth System Sciences*, *17*(12). <https://doi.org/10.5194/hess-17-5013-2013>
- Tikk, D., Kóczy, L., & Gedeon, T. (2003). A survey on universal approximation and its limits in soft computing techniques. *International Journal of Approximate Reasoning*, *33*(2), 185–202. [https://doi.org/10.1016/s0888-613x\(03\)00021-5](https://doi.org/10.1016/s0888-613x(03)00021-5)
- Tilmant, A., Beevers, L., & Muyunda, B. (2010). Restoring a flow regime through the coordinated operation of a multireservoir system: The case of the Zambezi River basin. *Water Resources Research*, *46*(7). <https://doi.org/10.1029/2009wr008897>
- Tilmant, A., Lettany, J., & Kelman, R. (2007). Hydrological risk assessment in the euphrates-tigris river basin: A stochastic dual dynamic programming approach. *Water International*, *32*(2), 294–309. <https://doi.org/10.1080/02508060708692208>
- Tilmant, A., Pinte, D., & Goor, Q. (2008). Assessing marginal water values in multipurpose multireservoir systems via stochastic programming. *Water Resources Research*, *44*(12). <https://doi.org/10.1029/2008wr007024>
- Trenberth, K. (2011). Changes in precipitation with climate change. *Climate Research*, *47*(1–2), 123–138. <https://doi.org/10.3354/cr00953>
- Tsitsiklis, J., & Van Roy, B. (1996). Feature-based methods for large scale dynamic programming. *Machine Learning*, *22*(1–3), 59–94. https://doi.org/10.1007/978-0-585-33656-5_5
- Turner, S. W. D., & Galelli, S. (2016). Regime-shifting streamflow processes: Implications for water supply reservoir operations. *Water Resources Research*, *52*(5), 3984–4002. <https://doi.org/10.1002/2015WR017913>
- Ushijima, T., & Yeh, W. (2013). Experimental design for estimating unknown groundwater pumping using genetic algorithm and reduced order model. *Water Resources Research*, *49*(10), 6688–6699. <https://doi.org/10.1002/wrcr.20513>
- Uysal, G., Schwanenber, D., Alvarado-Montero, R., & Şensoy, A. (2018). Short term optimal operation of water supply reservoir under flood control stress using model predictive control. *Water Resources Management*, *32*(2), 583–597. <https://doi.org/10.1007/s11269-017-1828-x>
- Varlet, H. (1923). Graphical analysis of operating conditions of a flow-control reservoir (in French). *Ann. Pans Chausstes Mem. Doc. Partie Tech*, *93*, 61–79.
- Vicuna, S., Dracup, J., Lund, J., Dale, L., & Maurer, E. (2010). Basin-scale water system operations with uncertain future climate conditions: Methodology and case studies. *Water Resources Research*, *46*(4). <https://doi.org/10.1029/2009wr007838>
- Wada, Y., Bierkens, M., de Roo, A., Dirmeyer, P., Famiglietti, J., Hanasaki, N., et al. (2017). Human-water interface in hydrological modeling: Current status and future directions. *Hydrology and Earth System Sciences Discussions*, 1–39.
- Wald, A. (1950). *Statistical decision functions*. John Wiley.
- Wang, L., Koike, T., Ikeda, M., Tinh, D., Nyunt, C., Saavedra, O., et al. (2014). Optimizing multireservoir releases in large river basins by combining distributed hydrological inflow predictions with rolling-horizon decision making. *Journal of Water Resources Planning and Management*, *140*(10), 05014006. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000452](https://doi.org/10.1061/(asce)wr.1943-5452.0000452)
- Wang, W.-C., Chau, K.-W., Cheng, C.-T., & Qiu, L. (2009). A comparison of performance of several artificial intelligence methods for forecasting monthly discharge time series. *Journal of hydrology*, *374*(3–4), 294–306. <https://doi.org/10.1016/j.jhydrol.2009.06.019>

- Ward, V., Singh, R., Reed, P., & Keller, K. (2015). Confronting tipping points: Can multi-objective evolutionary algorithms discover pollution control tradeoffs given environmental thresholds? *Environmental Modelling & Software*, 73, 27–43. <https://doi.org/10.1016/j.envsoft.2015.07.020>
- Watkins, C., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4), 279–292. <https://doi.org/10.1023/a:1022676722315>
- Watt, J., Borhani, R., & Katsaggelos, A. (2016). *Machine learning refined: Foundations, algorithms, and applications*. Cambridge University Press.
- Whateley, S., Steinschneider, S., & Brown, C. (2014). A climate change range-based method for estimating robustness for water resources supply. *Water Resources Research*, 50(11), 8944–8961. <https://doi.org/10.1002/2014wr015956>
- Wild, T., Reed, P., Loucks, D., Mallen-Cooper, M., & Jensen, E. (2019). Balancing Hydropower Development and Ecological Impacts in the Mekong: Tradeoffs for Sambor Mega Dam. *Journal of Water Resources Planning and Management*, 145(2), 05018019. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0001036](https://doi.org/10.1061/(asce)wr.1943-5452.0001036)
- Winemiller, K., McIntyre, P., Castello, L., Fluet-Chouinard, E., Giarrizzo, T., Nam, S., et al. (2016). Balancing hydropower and biodiversity in the Amazon, Congo, and Mekong. *Science*, 351(6269), 128–129. <https://doi.org/10.1126/science.aac7082>
- World Bank. (2009). *Directions in hydropower* (Tech. Rep.).
- Wurbs, R. (1993). Reservoir-system simulation and optimization models. *Journal of Water Resources Planning and Management*, 119(4), 455–472. [https://doi.org/10.1061/\(asce\)0733-9496\(1993\)119:4\(455\)](https://doi.org/10.1061/(asce)0733-9496(1993)119:4(455))
- Xu, B., Zhong, P.-A., Huang, Q., Wang, J., Yu, Z., & Zhang, J. (2017). Optimal hedging rules for water supply reservoir operations under forecast uncertainty and conditional value-at-risk criterion. *Water*, 9(8), 568. <https://doi.org/10.3390/w9080568>
- Xu, W., Meng, F., Guo, W., Li, X., & Fu, G. (2021). Deep reinforcement learning for optimal hydropower reservoir operation. *Journal of Water Resources Planning and Management*, 147(8), 04021045. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0001409](https://doi.org/10.1061/(asce)wr.1943-5452.0001409)
- Xu, W., Zhang, X., Peng, A., & Liang, Y. (2020). Deep reinforcement learning for cascaded hydropower reservoirs considering inflow forecasts. *Water Resources Management*, 34(9), 3003–3018. <https://doi.org/10.1007/s11269-020-02600-w>
- Xu, W., Zhao, J., Zhao, T., & Wang, Z. (2015). Adaptive reservoir operation model incorporating nonstationary inflow prediction. *Journal of Water Resources Planning and Management*, 141(8), 04014099. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000502](https://doi.org/10.1061/(asce)wr.1943-5452.0000502)
- Yakowitz, S. (1982). Dynamic programming applications in water resources. *Water Resources Research*, 18(4), 673–696. <https://doi.org/10.1029/wr018i004p00673>
- Yang, T., Gao, X., Sellars, S., & Sorooshian, S. (2015). Improving the multi-objective evolutionary optimization algorithm for hydropower reservoir operations in the California Oroville–Thermalito complex. *Environmental Modelling & Software*, 69, 262–279. <https://doi.org/10.1016/j.envsoft.2014.11.016>
- Yates, J., Veinott, E., & Patalano, A. (2003). Hard decisions, bad decisions: On decision quality and decision aiding. *Emerging Perspectives on Judgment and Decision Research*, 13–63. <https://doi.org/10.1017/cbo9780511609978.003>
- Yeh, W. (1985). Reservoir management and operations models: A state of the art review. *Water Resources Research*, 21(12), 1797–1818. <https://doi.org/10.1029/wr021i012p01797>
- Young, G. K. (1967). Finding reservoir operating rules. *Journal of the Hydraulics Division*, 93(6), 297–322. <https://doi.org/10.1061/jycej.0001714>
- Zaniolo, M., Giuliani, M., & Castelletti, A. (2021). Neuro-evolutionary direct policy search for multiobjective optimal control. *IEEE Transactions on Control Systems Technology*, 1–13. <https://doi.org/10.1109/TNNLS.2021.3071960>
- Zarfl, C., Lumsdon, A., Berlekamp, J., Tydecks, L., & Tockner, K. (2015). A global boom in hydropower dam construction. *Aquatic Sciences*, 77(1), 161–170. <https://doi.org/10.1007/s00027-014-0377-0>
- Zhang, C., Li, Y., Chu, J., Fu, G., Tang, R., & Qi, W. (2017). Use of many-objective visual analytics to analyze water supply objective trade-offs with water transfer. *Journal of Water Resources Planning and Management*, 143(8), 05017006. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000800](https://doi.org/10.1061/(asce)wr.1943-5452.0000800)
- Zhang, W., Liu, P., Chen, X., Wang, L., Ai, X., Feng, M., et al. (2016). Optimal operation of multi-reservoir systems considering time-lags of flood routing. *Water Resources Management*, 30(2), 523–540. <https://doi.org/10.1007/s11269-015-1175-8>
- Zhang, W., Liu, P., Wang, H., Chen, J., Lei, X., & Feng, M. (2017). Reservoir adaptive operating rules based on both of historical streamflow and future projections. *Journal of Hydrology*, 553, 691–707. <https://doi.org/10.1016/j.jhydrol.2017.08.031>
- Zhang, X., Liu, P., Xu, C.-Y., Ming, B., Xie, A., & Feng, M. (2018). Conditional value-at-risk for nonstationary streamflow and its application for derivation of the adaptive reservoir flood limited water level. *Journal of Water Resources Planning and Management*, 144(3), 04018005. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000906](https://doi.org/10.1061/(asce)wr.1943-5452.0000906)
- Zhao, Q., Cai, X., & Li, Y. (2019). Determining inflow forecast horizon for reservoir operation. *Water Resources Research*, 55(5), 4066–4081. <https://doi.org/10.1029/2019wr025226>
- Zhou, J., Jia, B., Chen, X., Qin, H., He, Z., & Liu, G. (2019). Identifying efficient operating rules for hydropower reservoirs using system dynamics approach—A case study of three gorges reservoir, China. *Water*, 11(12), 2448. <https://doi.org/10.3390/w11122448>
- Zmijewski, N., Bottacin-Busolin, A., & Wörman, A. (2016). Incorporating hydrologic routing into reservoir operation models: Implications for hydropower production planning. *Water Resources Management*, 30(2), 623–640. <https://doi.org/10.1007/s11269-015-1181-x>