



RESEARCH ARTICLE

10.1002/2015WR018253

Key Points:

- A bottom-up approach to climate change impact assessment using formal optimization techniques
- Identifies maximum operational adaptive capacity of water resource systems by adapting management policies
- Suitable for adapting system operation when the number of adaptation options is large

Correspondence to:

S. Culley,
sam.culley@adelaide.edu.au

Citation:

Culley, S., S. Noble, A. Yates, M. Timbs, S. Westra, H. R. Maier, M. Giuliani, and A. Castelletti (2016), A bottom-up approach to identifying the maximum operational adaptive capacity of water resource systems to a changing climate, *Water Resour. Res.*, 52, 6751–6768, doi:10.1002/2015WR018253.

Received 16 OCT 2015

Accepted 12 AUG 2016

Accepted article online 18 AUG 2016

Published online 9 SEP 2016

A bottom-up approach to identifying the maximum operational adaptive capacity of water resource systems to a changing climate

S. Culley¹, S. Noble¹, A. Yates¹, M. Timbs¹, S. Westra¹, H. R. Maier¹, M. Giuliani², and A. Castelletti^{2,3}

¹School of Civil, Environmental and Mining Engineering, The University of Adelaide, Adelaide, Australia, ²Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milano, Italy, ³Institute of Environmental Engineering, ETH Zurich, Zurich, Switzerland

Abstract Many water resource systems have been designed assuming that the statistical characteristics of future inflows are similar to those of the historical record. This assumption is no longer valid due to large-scale changes in the global climate, potentially causing declines in water resource system performance, or even complete system failure. Upgrading system infrastructure to cope with climate change can require substantial financial outlay, so it might be preferable to optimize existing system performance when possible. This paper builds on decision scaling theory by proposing a bottom-up approach to designing optimal feedback control policies for a water system exposed to a changing climate. This approach not only describes optimal operational policies for a range of potential climatic changes but also enables an assessment of a system's upper limit of its operational adaptive capacity, beyond which upgrades to infrastructure become unavoidable. The approach is illustrated using the Lake Como system in Northern Italy—a regulated system with a complex relationship between climate and system performance. By optimizing system operation under different hydrometeorological states, it is shown that the system can continue to meet its minimum performance requirements for more than three times as many states as it can under current operations. Importantly, a single management policy, no matter how robust, cannot fully utilize existing infrastructure as effectively as an ensemble of flexible management policies that are updated as the climate changes.

1. Introduction

Water resource infrastructure, including reservoirs, levees, and river regulators, are assets requiring a significant financial outlay, with annual global investments in such infrastructure exceeding US \$500 billion [Ashley, 2006]. Most of this infrastructure has been designed on the assumption that the statistical characteristics of future inflows are equivalent to those of the historical data. This assumption of a stationary climate is unlikely to be valid in the future, due to climatic changes that will affect most aspects of the hydrological cycle [Milly et al., 2008; IPCC, 2013]. As a result, many water resource systems are expected to become increasingly vulnerable and experience degraded performance, as the climate approaches the limit of what was originally accounted for by design safety factors [Risbey, 2011; Paton et al., 2013; IPCC, 2014]. However, the economic and environmental cost of upgrading infrastructure to counteract this is likely to be high [Paton et al., 2014; Beh et al., 2015a, 2015b]. This provides impetus to focus on identifying *operational* strategies that maximize system performance under a changing climate, thus enabling the best use to be made of existing infrastructure [Gleick, 2003; Giuliani et al., 2016].

Traditionally, top-down approaches have been used as the basis for developing adaptation strategies, by describing the performance of water resource systems under a discrete set of global climate projections. Projections are acquired using general circulation models (GCMs) [Arnell et al., 2004; Brekke et al., 2009; Vano et al., 2010; Wilby and Dessai, 2010; Anghileri et al., 2011; Giuliani and Castelletti, 2016], the outputs of which are fed into an integrated water resource system model to determine the system's performance with respect to each projection. The system's performance can be classified as "acceptable" or "unacceptable" for each projection, and the potential benefits of alternative adaptation strategies can be explored [Prudhomme et al., 2010]. Due to the discrete nature of these projections, such top-down approaches are generally not suitable for identifying

thresholds of performance with respect to changes in climate exposure, as it may be difficult to identify the exact degree of climate change at which system performance changes from acceptable to unacceptable.

Bottom-up approaches are an alternative to the top-down procedure described above, and have been designed to identify performance thresholds independently from climate models' projections. To implement a bottom-up approach, climate exposures are generated for a range of plausible changes in climate, including those beyond the bounds projected by GCMs, and system response is assessed against each climate exposure [Lempert *et al.*, 2004; Prudhomme *et al.*, 2010; Brown *et al.*, 2012; Brown and Wilby, 2012; Weaver *et al.*, 2013]. This enables a more thorough understanding of how a system responds to changes in climate variables, for example, by identifying the changes in climate exposure that can cause unsatisfactory degradation in system performance, or thresholds for system failure [Whateley *et al.*, 2014; Steinschneider *et al.*, 2015]. Thus, the purpose of the bottom-up approach is to identify the exposures under which a particular system performs satisfactorily, rather than to assess system performance under GCM-based climate change projections [Lempert and Collins, 2007].

For climate exposures that are associated with degraded performance or system failure, an alternative management strategy—one better suited to the new climate exposure—may improve performance or avoid failure. Approaches that focus on the identification of system performance thresholds at which a particular course of action becomes preferable to another have been successfully used in many areas of water resources management [Hyde and Maier, 2006; Ravalico *et al.*, 2010; Herman *et al.*, 2014; Guillaume *et al.*, 2016]. An example in the climate impact assessment arena is “decision scaling” [Brown *et al.*, 2012], where a scenario neutral climate space is divided into regions for which different discrete decisions would be preferable, thus allowing the articulation of preferred adaptation options in response to specific changes in climate. This is a convenient approach as it (i) provides an understanding of system vulnerability, (ii) identifies decision thresholds that can be compared easily with climate predictions, and (iii) demonstrates whether a particular decision can achieve acceptable performance under given climate conditions [Turner *et al.*, 2014; Poff *et al.*, 2015]. In the early work on decision scaling, the approach was used to determine regions of future climate, referred to as “climate states,” for which one infrastructure decision is preferable over another [Brown *et al.*, 2012; Herman *et al.*, 2015]. However, recently decision scaling has also been used for operational problems, identifying the changes in hydrometeorological variables required before an operational plan is no longer successful [Whateley *et al.*, 2014; Steinschneider *et al.*, 2015]. While Steinschneider *et al.* [2015] used this approach to test the robustness of different management alternatives, Whateley *et al.* [2014] examined the robustness of a physical system, by identifying management strategies as an optimal response to addressing the impact of future climates.

The approach presented in this paper builds on the above work by introducing a formal optimization formulation that identifies the operational strategy, formulated as a closed loop control policy, which performs best for a particular climate exposure. Here “climate exposure” is expressed as a combination of hydrometeorological variables scaled from current climate. Traditional reservoir operations are generally defined as rule curves, which are not able to adapt when the system deviates from the hydrological conditions used in the design of the rule [Loucks and Sigvaldason, 1981]. Instead of static rule curves, the formulation of the system operation as a feedback control policy enhances the robustness and adaptive capacity of the system, as the operational decisions are informed by the feedback loop depending on the current state of the system. As this process is repeated for all combinations of hydrometeorological variables considered, the optimal solutions (decisions) for all exposures that are of interest are identified. Consequently, the proposed approach is also able to find the theoretical upper limit for adaptation of a water resource system by identifying optimal management policies with respect to all climate exposures of interest. We call this limit the maximum operational adaptive capacity of the system. Beyond this limit, infrastructure upgrades would be required to adapt to further changes in climate. Therefore, knowledge of this failure boundary provides operators with the full extent of climate change that their system could withstand, thereby potentially avoiding unnecessary or premature infrastructure upgrades and reducing the vulnerability of the system [Mastrandrea *et al.*, 2010]. A better understanding of the maximum operational adaptive capacity and of the associated failure boundary can be embedded in existing methods of sequential decision making [Haasnoot *et al.*, 2013; Beh *et al.*, 2015b; Kwakkel *et al.*, 2016; Maier *et al.*, 2016] to determine when adaptation options through changes in management are no longer available.




Figure 1. Outline of proposed approach for identifying the maximum operational adaptive capacity of a water resource system, demonstrated for two attributes (i.e., $\Theta = [\theta_1, \theta_2]$).

The proposed approach is presented in the following section and then demonstrated in section 3, where it is applied to a case study on Lake Como, a regulated lake in Northern Italy where the lake operator has to balance water supply and flood control. The results of this case study are presented in section 4, followed by a discussion of the utility of identifying an upper limit to adaptation in section 5. Conclusions are presented in section 6.

2. Proposed Approach

The proposed approach first requires the identification of an exposure space (Θ), which is defined as the set of “hydrometeorological states” that a system may confront in the future. The performance of the water resource system is then assessed under both historical and alternative operating policies optimized for each climate exposure, with the latter used to identify a theoretical upper limit of the system’s adaptive capacity. Finally, climate change projections can be superimposed on the exposure space, to understand the plausibility of failure thresholds being reached given the current understanding of climate change. Each of these steps is summarized in Figure 1 and elaborated upon in the following subsections.

2.1. Generate an Exposure Space

The exposure space (Θ) is obtained by making incremental changes to attributes (e.g., mean, 7 day maximum, etc.) of hydrological variables such as temperature and precipitation. This leads to an n -dimensional exposure space $\Theta = [\theta_1 \dots \theta_n]$ made up of the selected attributes of all the variables considered, which should represent the variables and attributes to which a system’s performance is most sensitive [Mastrandrea et al., 2010].

Once the axes of the exposure space have been selected, it is necessary to sample across the exposure space by developing perturbed time series of the relevant climatic variables. This has generally been achieved using a simple scaling of historical data [Prudhomme et al., 2010], although more sophisticated techniques, such as the use of stochastic weather generators, can also be used [Steinschneider and Brown, 2013; Guo et al., 2016].

It should be noted that the accuracy of the maximum operational adaptive capacity boundary will depend on the resolution of the generated states. However, increased resolution must be balanced against the computational effort required in subsequent steps of the methodology, as separate model (Step 2) and optimization (Step 3) runs are required for each generated hydrometeorological state.

2.2. Assess Current System Performance

To identify the system’s maximum operational adaptive capacity, it is necessary to first determine the robustness of the system, assuming that the system’s operations remain unchanged. This is achieved by simulating system performance for each point of the exposure space using current operating policy. Considering all system objectives, the hydrometeorological states under which system performance is acceptable and the states under which this is not the case are distinguished on the exposure space. Subspaces of “failure” and “success” are thus formed, with the success subspace (denoted S_f) given as

$$S_f = \{(\theta_1 \cdots \theta_n) \in R^n : M(\theta_1 \cdots \theta_n, p) < \mathbf{f}\}, \tag{1}$$

where $\theta_1 \cdots \theta_n$ are the selected attributes of the exposure space, p is the operating policy, and $M(\theta_1 \cdots \theta_n, p)$ is the integrated water system model that maps each point on the exposure space into system performance for a given policy p . System performance is represented by a set of failure criteria for each objective, with \mathbf{f} representing the threshold of acceptable values. A simple approach for finding the full failure boundary when considering multiple objectives is a binary classification system, where “failure” is assigned at each state if any failure criterion value exceeds a threshold, and “success” is assigned if all failure criteria do not exceed the failure threshold. However, which measure of overall system failure is most appropriate is case study dependent.

2.3. Assess Optimized System Performance

This step involves using formal optimization techniques to determine the management strategy that maximizes system performance for each hydrometeorological state in the exposure space. By optimizing the operating policy, p , against each point in the exposure space, it becomes possible to delineate a boundary between “success” and “failure” given the optimized policy. This boundary is defined as the system’s maximum operational adaptive capacity. The optimization formulation for achieving this is given by

$$\arg \min_p J = M(\theta_1 \cdots \theta_n, p), \tag{2}$$

where the feedback control policy p that minimizes the one or more failure criteria J (with thresholds \mathbf{f}) for that hydrometeorological state is identified. Equation (1) can be used to map the postoptimization failure boundary, where p is now an optimized policy specific for each hydrometeorological state. Once the states under which system performance is satisfactory under optimized operational regimes have been identified, the maximum operational adaptive capacity can be measured as the ratio of the number of states under which system performance is satisfactory postoptimization to the number of states under which system performance is satisfactory under historical operation, as follows:

$$C = \frac{\int_{\Theta} S_{f_optimized}}{\int_{\Theta} S_{f_historical}}, \tag{3}$$

where C is the maximum operational adaptive capacity. In practice this operation is undertaken over a discrete exposure space, where each axis of the exposure space θ_i is uniformly sampled between a set of plausible bounds for that attribute.

The process of determining adaptation options by optimizing the operating policy for each point in the exposure space, which capitalizes on the flexibility of operational decisions compared to longer-term and potentially more expensive structural upgrades, is the differentiating characteristic of this approach. Furthermore, this approach provides a look-up table for optimal management strategies, given certain hydro-meteorological conditions. This potentially allows adaptation to climate change through adaptive management, up to the point where the physical limitations of the system infrastructure prevent successful performance, even under optimal management. While the above optimization problem can be solved using a range of approaches, the use of evolutionary algorithms (EAs) is attractive in this setting, as they can be easily linked with existing simulation models of the water resources system of interest [Kingston et al., 2008; Giuliani et al., 2014a; Maier et al., 2014; Giuliani et al., 2015; Maier et al., 2015; Salazar et al., 2016].

2.4. Superimpose Projections Onto the Exposure Space

The possibility of the failure boundaries being reached under both current and optimal system operations is assessed by considering climate projections for the region in which the system is located. Although climate model projections are highly uncertain, mapping the projections onto the exposure space may enable




Figure 2. Map of Lake Como system.

patterns, disparities or similarities to be identified and compared to the optimal failure thresholds. This can provide consensus or otherwise about the possibility of the climate reaching a state that causes system failure, time frames for reaching failure, the value of adapting management policies, and the potential timing of these policies.

3. Application

3.1. Case Study Description

This section describes the implementation of the proposed approach for a real-world case study based on Lake Como. It should be noted that while the case study provides a reasonably realistic representation of the Lake Como system, a number of simplifications have been made as the primary purpose is to illustrate the proposed approach, rather than to identify optimal management approaches for the actual Lake Como system.

Lake Como, a regulated lake located in Northern Italy (Figure 2) with an active storage capacity of 254 Mm³, was selected as the illustrative case study as it is a significant system that is situated in a region with complex hydrological relationships, and because it has a variety of stakeholders with conflicting objectives. Due to the subalpine location of the system, inflows into the lake are largely derived from snowmelt, with very short travel times associated with the steep terrain. Lake Como has been regulated since 1946 to provide a more reliable water supply to downstream users, particularly to support irrigation in five agricultural districts, which collectively represent one of the largest irrigated regions in Europe. Major crops in the irrigated regions downstream of the reservoir are cereals, especially maize, along with temporary grasslands for livestock. Irrigation is practiced with the border method or free-surface flooding.

The lake is regulated on an annual cycle, typically storing the large snowmelt inflows in spring, drawing down to provide irrigation supply during summer, and then receiving inflows as a result of heavy rain in autumn. However, preventing floods along the lake shores, particularly in Como city, limits the storage capacity of the reservoir and introduces a clear conflict between flood control and irrigation supply [Castelletti *et al.*, 2010].

The model of the Lake Como system is made up of three main components: the upstream catchment, the reservoir dynamics, and the reservoir discharge into the Adda River, which serves the irrigation districts in the downstream part of the system. In this study, the upstream catchment is represented by a lumped

Table 1. GCM-RCM Combinations of Climate Models

Model Reference #	GCM	RCM	RCP
1	CM5 (CNRM CERFACS)	CCLM4 (CLMcom)	4.5
2	CM5 (CNRM CERFACS)	CCLM4 (CLMcom)	8.5
3	CM5 (CNRM CERFACS)	RCA4	4.5
4	CM5 (CNRM CERFACS)	RCA4	8.5
5	EARTH (ICEC)	CCLM4 (CLMcom)	4.5
6	EARTH (ICEC)	CCLM4 (CLMcom)	8.5
7	EARTH (ICEC)	HIRHAM5 (DMI)	4.5
8	EARTH (ICEC)	HIRHAM5 (DMI)	8.5
9	EARTH (ICEC)	RACMO22E (KNMI)	4.5
10	EARTH (ICEC)	RACMO22E (KNMI)	8.5
11	EARTH (ICEC)	RCA4	2.6
12	EARTH (ICEC)	RCA4	4.5
13	EARTH (ICEC)	RCA4	8.5
14	ESM LR (MPI)	REMO 2009 (MPI)	4.5
15	CanESM2 (CCCma)	RCA4	4.5
16	CanESM2 (CCCma)	RCA4	8.5
17	MIROC	RCA4	4.5
18	MIROC	RCA4	8.5
19	NCC	RCA4	4.5
20	NCC	RCA4	8.5
21	NOAA	RCA4	4.5
22	NOAA	RCA4	8.5

HBV model [Bergström and Singh, 1995], which simulates the soil water balance and subsequent runoff produced by rainfall, snowmelt, and evapotranspiration. The reservoir dynamics is described by a mass balance equation assuming a modeling and decision-making time step of 24 h,

$$s_{t+1} = s_t + q_{t+1} - r_{t+1}, \quad (4)$$

where s_t is the lake storage [m^3] at time t , and q_{t+1} and r_{t+1} are the net inflow and release volumes in the time interval $[t, t + 1)$, respectively. In particular, the regulated release, which depends on the daily release decision u_t , is defined as $r_{t+1} = R(s_t, u_t, q_{t+1})$, where the function $R()$ takes into account the legal and physical constraints on the lake level and release, includ-

ing spills when the level exceeds the maximum capacity [Soncini Sessa, 2007]. The Adda River is described by a plug-flow model at a daily time step, which simulates the routing of the lake releases from the lake outlet to the intake of the irrigation canals.

The regulation of Lake Como is driven by two primary objectives: flood control and downstream irrigation supply, subject to a minimum environmental flow constraint on the lake releases to ensure adequate ecological conditions in the Adda River. The flood objective is formulated as the daily average flooded surface in Como [m^2], which is a function of the lake level averaged over the simulation horizon. The irrigation objective is formulated as the quadratic daily average water deficit [kL^2] with respect to the daily water demand of the downstream system. This quadratic formulation penalizes severe deficits in a single time step, while allowing for more frequent, small shortages [Hashimoto et al., 1982].

For this study, the system is considered to fail if its operations result in the exceedance of a daily average flooded area of $100 m^2$ and/or a daily average squared water deficit value of $400 kL^2$. These failure criteria are defined as slightly higher than current failure levels (which are $79 m^2$ and $359 kL^2$, respectively), based on the assumption that the system is currently efficiently operated and some small degradation in performance is allowed.

Time series of daily mean areal precipitation, daily maximum and minimum temperature, and lake inflows were available over the historical period 1965–1980. In addition, projected time series of the same variables were obtained by applying a statistical downscaling method based on quantile mapping [Boé et al., 2007; Déqué, 2007] over 22 different scenarios of climate change, comprising seven general circulation models (GCMs), six regional climate models (RCMs) and three representative concentration pathways (RCPs) (see Table 1). These projections come from the EURO CORDEX project [see Jacob et al., 2014].

3.2. Generating the Climate Exposure Space

The exposure space was generated by systematically perturbing the time series for the chosen attributes of the two hydrometeorological input variables of the hydrological model. These two attributes are the annual average precipitation depth and annual average temperature, and they form the axes of the exposure space, with each grid point in the space representing a unique hydrometeorological state. Perturbations were made uniformly across a year, as opposed to changing extremes or seasonality, because the failure criteria are also defined in terms of annual averages of system performance.

Precipitation was perturbed as a percentage change to daily wet days (>1 mm rainfall), from 90 to 130% of current values. This leads to an equivalent change in total annual precipitation volume. Average

temperature was perturbed additively for each daily time step, from -5 to 15°C of current values. These perturbation ranges were selected to ensure climate projections for 2025 would be contained within the bounds of the exposure space. Furthermore, although the upper bound for the temperature change is substantially higher than would be expected based on the current generation of climate change projections, this value was selected through an iterative process to show a limit of the system’s operational adaptive capacity with respect to increases in temperature (discussed further in section 4.2 below). With a step size of 1°C for average annual temperature and 1% for precipitation volume change, this corresponds to 861 unique hydrometeorological states. These perturbations were used by the HBV hydrological model to generate the reservoir inflow time series, enabling these inflows to be input into the reservoir dynamic model.

3.3. Assessing Current System Performance

The reservoir model for Lake Como, described in section 3.1, was used to estimate the system’s performance for each of the 861 inflow states that correspond to the generated hydrometeorological states. Each simulation results in the generation of the performance values for the two criteria—flood and irrigation—under a modeled operating policy representing the historical lake regulation. Failure was considered to occur when at least one failure criterion was exceeded. The resulting performance was mapped onto the exposure space by assigning a success or failure outcome to each state.

Modeling the historical regulation of Lake Como requires the mathematical formulation of the operating policy adopted by the lake operator (p in equation (2)), which provides the release decision u_t , at each time step t , as a function of the lake level h_t . In general, we can assume that the policy p is a periodic sequence, with period one year, of operating rules of the form $u_t = m_{\gamma}(t, h_t, \gamma)$, where γ is a vector of unknown parameters [Castelletti et al., 2008]. For a preassigned family of functions $m_{\gamma}(t, h_t, \gamma)$, the values of γ can be determined via regression as those that minimize the distance metric between historical and modeled releases [Guariso et al., 1986; Corani et al., 2009; ICSC, 2009].

In the case of the Lake Como case study, historical releases are known, but there are no official operating rules, as the dam is operated within a legal regulation range based on operator experience. To model the lake’s historical operation, we assume that the lake operator is a rational agent who, according to our model formulation, is balancing flood control and irrigation supply following a parameterized operating policy depending on time and lake level. Optimizing the policy parameters, γ , with respect to this objective function yields a management policy that implicitly captures the actual decisions of the lake operator [Giuliani et al., 2014a]. The optimization tool used is as described in section 3.4.

To estimate the historical policy, we parameterize the operating rules $m_{\gamma}(t, h_t, \gamma)$ by Gaussian radial basis functions (RBFs) to approximate the multi-input, single-output nature of the Lake Como operating policy. RBF-based policy parameterization has been demonstrated to outperform other nonlinear approximating networks (e.g., traditional artificial neural networks) in their representation of operating policies [Giuliani et al., 2014b]. The parameterized RBF policy requires the estimation of 28 parameters (the vector γ) representing the centers and radii of the Gaussian functions and the weights used in the convex combinations of their values [see Busoniu et al., 2011]. Using the evolutionary multiobjective direct policy search [Giuliani et al., 2015], a set of Pareto-optimal RBF policies is designed. Among this set, the solution characterized by the closest performance to the historical operating policies in the two considered objectives is selected. The results over the 16 year modeled period closely approximate those of the historical period (Table 2). Only two objective functions were used to select the solution, as the Lake Como case study considered in this paper for illustration purposes is a simplification of the real system, where other objectives (i.e., environmental interests) are considered in the historical regulation, and so the trajectories of the policy used in this study and the historical one are not directly comparable. However, for the purposes of illustrating the proposed approach, a solution that provides similar performance for the two main competing objectives (shown in Table 2) is suitable.

Table 2. Summary of Modeled System Behavior

System Objectives		Magnitude
Floods	Observed	Flooded area = 79.0 m ²
	Modeled	Flooded area = 77.2 m ²
Irrigation	Observed	Irrigation deficit = 359 kL ²
	Modeled	Irrigation deficit = 352 kL ²

3.4. Assessing Optimized System Performance

The performance of any reservoir system is dictated by the day-to-day release decisions of the reservoir operator. The same RBF policy

structure used for reproducing the historical regulation is used for assessing the optimized system performance obtained by solving the problem formulated in equation (2), where the 28 parameters of the RBF policy are the decision variables to be optimized. These decision variables are subjected to the following constraints: the radii of the Gaussian functions range from 0 to 1, the center coordinate of each function has the range -1 to 1 , and the weighted combinations are non-negative and sum to 1 [Busoniu et al., 2011]. A genetic algorithm (GA) was linked with the system model in order to optimize system performance, as evolutionary algorithms have been shown to perform well for the optimization of reservoir operation [Giuliani et al., 2015; Li et al., 2015; Tsoukalas and Makropoulos, 2015; Yang et al., 2015].

The objective function, OF , minimized by the GA, was

$$OF = \sqrt{(1-r_f)^2 + (1-r_s)^2} * PF, \tag{5}$$

where r_f is the flood reliability function (the ratio of the days without a flood event to the length of the simulation horizon), r_s is the irrigation reliability function (the ratio of available supply to the demand averaged across all time steps, with ratio values capped at 1), and PF is the penalty function used. Reliability functions were found to be more effective than simply considering the magnitudes of flood and irrigation deficit that make up the failure criteria, as they provide additional emphasis on the timing of releases during optimization. The multiobjective problem was reformulated into a single objective function, to allow automation of the optimization of the 861 hydrometeorological states considered, as no decision about optimal trade-off selection needed to be made.

The failure criteria were used in a penalty function, PF in equation (5), to control the magnitude of flood or drought events in the final solution. If the average flood inundation for a day was more than 100 m^2 , then the solution was penalized regardless of the corresponding storage reliability. Similarly, if the average squared irrigation deficit exceeded 400 kL^2 , the solution was penalized regardless of volumetric reliability performance. Each violation of the constraints resulted in the fitness of that evaluation increasing by a factor of 5. This ensures that the optimal solutions found by the GA maximize system reliability for each hydrometeorological state while achieving “successful” overall performance.

System performance resulting from the optimal operating policies for each of the 861 states was mapped onto the exposure space using the process outlined in section 3.3, distinguishing regions of success under current management, hydrometeorological states that can experience successful performance under optimal management, and regions of the exposure space that result in failure even under optimal management.

3.5. Compare Failure Boundaries to Climate Projections

To understand the plausibility of the scenario-neutral failure boundaries in the exposure space being reached, climate projections for the system were mapped onto the exposure space. This was done for four time slices: 2025, 2050, 2075, and 2100. The hydrometeorological states were estimated using 11 year windows centered on each time slice (i.e., 5 years either side of a year selected for consideration). The values were calculated for each time slice for the 22 GCM-RCM combinations.

4. Results

4.1. Current System Performance Thresholds

Figure 3 shows the performance of the current system operations for each hydrometeorological state with respect to two criteria: the flood objective and the irrigation objective. The failure boundary of the Lake Como case study under current management indicates that the lake can continue to maintain adequate performance for significant increases in total precipitation and temperature—but only if the increases occur for both variables simultaneously (blue points in Figure 3).

Focusing initially on changes in precipitation, the irrigation supply failure (orange points) is primarily driven by decreases in precipitation, as would be expected due to the reduced inflow that results under such exposures. In contrast, flood failure (red points) mostly occurs with an increase in precipitation volume (without a simultaneous increase in temperature), due to the augmented peak flow that results from larger rainfall events.




Figure 3. Multiobjective failure boundary of current operations mapped onto the exposure space, relative to the 1965–1980 baseline.

ical rainfall alters only the magnitude, not the seasonality), but can still cause a reduction in the magnitude of peak inflow due to increased evaporation and reduced snowmelt.

The historical operation can accommodate temperature increases up to 4°C, without failure in either objective (assuming no change in rainfall). The flood objective is met under these conditions because of the offset of peak snowmelt inflow from peak rainfall inflow, and the reduction in snowmelt magnitude, both of which reduce the overall peak from what the historical policy is designed for. The irrigation objective is also satisfied because the earlier snowmelt inflow occurs at a time in the year when the historical policy is releasing less each day, so the reduction in peak inflow volume is offset by the extra inflow stored. For temperature increases beyond about 4°C, evaporation in the catchment further reduces inflows, and for both competing objectives to be satisfied, this loss in inflow needs to be offset with increases in precipitation.

Interestingly, the system is far more sensitive to decreases in annual average temperature. Decreases of about 3°C from historical values cause an increase in peak inflow, because the colder temperatures result in more precipitation stored as snow during winter and thus more snowmelt in spring/summer. This also changes the timing of the peak inflow, with snowmelt occurring later in the year. Decreases in temperature greater than 3°C result in a reduction in inflow, because temperature reaches the point where snowmelt is induced less often, thus causing the precipitation to be stored in the upper catchment snowpack.

This initial increase and subsequent decline of inflow causes the changing failure region of the exposure space that can be seen in Figure 3, just to the left of the current conditions. With no change in precipitation, a decrease in temperature first causes failure in the flood objective (red dots), due to inflow from increased snow storage. Failure in both objectives occurs next (purple dots), as the additional inflows not only causes additional flooding, but due to the later timing of the inflows do not provide sufficient storage for irrigation. Decreases in temperature greater than 3°C only result in irrigation failure, as inflow is reduced significantly.

4.2. Optimal Adaptive Responses to Each Exposure

By optimizing the operating policies to each exposure, the system would continue performing at a level equal to or above its current level for a broader range of hydrometeorological states (Figure 4). The measure of maximum operational adaptive capacity for the system (defined by a discrete form of equation (3)) is $C = 3.23$, where the historical operation resulted in 172 successful hydrometeorological states, compared to 555 states after optimization. When considering the multi objective description of performance in Figure 3 and the adaptive capacity in Figure 4, it becomes apparent that the region of climate exposure that the

The relationship between performance and temperature is more complicated, as inflow dynamics are largely driven by snowmelt, which is triggered by temperature in the HBV model. As temperature increases up to 4°C, less precipitation is held as snow, reducing the magnitude of the snowmelt inflow. Additionally, the inflow from snowmelt will occur earlier in the year, no longer coinciding with the peak rainfall (historically occurring in late spring/early summer). At the point of about a 4°C increase, the largest daily inflows are now due to rainfall, rather than snowmelt. Temperature increases beyond this point have no further impact on the timing of peak inflows (as the perturbation method for historical




Figure 4. Optimal adaptive capacity of the Lake Como system.

system is most capable of adapting its performance for its irrigation supply. In these regions of the exposure space, warmer temperatures and lower precipitation reduce inflow and hence reduce flood concerns, allowing the lake to be operated at higher levels. The reoptimized operating policies focus mainly on storing water in larger volumes and for longer durations.

However, when considering the adaptation to regions of the exposure space that cause flooding failure, there are significantly fewer states that can be adapted to, suggesting that the current infrastructure is sensitive to changes in hydrometeorological states that cause increased inflows. Even under optimized operating policies, the system

fails for a temperature decrease of more than 3°C. This implies that the existing infrastructure’s ability to respond to changes in flood occurrence and snow storage is limited.

To illustrate the effects of a reoptimized operating policy, Figure 5 compares the lake storage throughout the first year of simulation. One storage trajectory is based on the daily releases of the current lake regulation (blue line) and the other from the daily releases due to optimal operating policy (green line). The difference in releases due to the two alternative operating policies translates to the different lake storages through time. Two points in the exposure space were examined, one that caused failure of the historical system’s operations in the irrigation objective (Figure 5, top) and another that caused failure in the flood objective (Figure 5, bottom). Both states were on the boundary for adaptation, meaning that they correspond to the most extreme climatic conditions that could be adapted to. These two states were selected to demonstrate how, at the adaptation boundary, the RBF used to simulate different operating policies responded to each of the system objectives almost in isolation, as boundary conditions that cause one objective to be close to failure are actually more favorable for the other objective.

The hydrometeorological state with poor irrigation performance (Figure 5, top) has a lower volume of total annual precipitation than the historical climate, increasing the need to store excess water. When applying the historical operating policy to this exposure, the lake releases are too aggressive in spring and summer, shown by the steep slope and resulting decrease in storage level for the blue line. In addition to this change in magnitude, there is a large shift of approximately 20 days in the timing of a major release, demonstrating the ability of the reoptimized policies to satisfy system objectives. Historical releases are also greater than the optimal release in the first 90 days of the simulation, due to more consistent availability of water replenishing supply under the current climate, but such a strategy drains the lake under the drier climate condition.

The operating policy for a flood sensitive hydrometeorological state (Figure 5, bottom) did not greatly alter the timing of the major releases. This is because the storage is near empty prior to the major inflow, and lowering the level of the reservoir sooner was not an option. Thus, the resulting effect on the reoptimized operating policy is that the timing of releases did not have to be managed, but the magnitude of release did. The quantity of water released after the first major inflow was increased compared to the current operation, providing the irrigator objective with enough storage and preventing the second peak inflow from causing as much flooding.




Figure 5. Comparison of alternative management options resulting from the use of current and re-optimized operating policies in the first year of the evaluation horizon. The flood sensitive climate corresponds to a 15% increase in precipitation and a 3° increase in temperature, and the irrigation sensitive climate corresponds to a 10% decrease in precipitation and a 13° increase in temperature.

4.3. Superposition of Projections Onto the Exposure Space

To understand the potential system failure timeframe, the model projections for four future time slices were overlaid onto the exposure space (Figure 6). This allows the proportion of GCM projections indicating failure of the system (under both the current and optimized management policies) to be examined, which is useful in providing recommendations on the future management of the system, and whether infrastructure upgrades are likely to be an unavoidable conclusion. Note that the fraction of scenarios under which the system performs successfully should not be taken as a proxy for the likelihood of that event occurring, as the climate projections are not sampled in equal fashion, and Table 1 makes it clear that these projections are not independent.

The first time slice (2025) shows there is disagreement between model combinations (i.e., GCM-RCM) as to the potential future hydrometeorological state the system will be exposed to. Slightly more than half (13 out of 22) of the projections suggest system “success” under current operation,

whereas four indicate complete system failure even with system reoptimization. The remaining five projections would lead to system failure under the current operation policies, but potentially could be adapted to.

By the 2050 time slice, an increase in both average temperature and precipitation can be observed relative to 2025. This shift has moved the projections towards a critical region of the exposure space, with more projections indicating adaptation would be required. There are in fact four projections that have moved beyond the bounds of the axes. It should be noted that it is not necessary to have all projections shown on the exposure space for this study, as they are being used solely as a tool to provide information on system behavior around failure boundaries shown. Overall, only six projections lie within the ‘success’ region of the current management policy, with 13 falling into the region where optimized management may be possible, and three predicting a failure to uphold the chosen level of acceptable performance.

Projections for the next time slice (2075) show the increase in temperature and precipitation has continued, and the spread of projections has increased dramatically. With projections becoming more varied, it makes it difficult to use them to provide information on system behavior. The last time slice (2100) has continued the trend of increasing temperature and precipitation, coupled with an even greater spread of model projections compared with the previous time interval. Once again, the spread further reduces the ability to understand how the system will behave in the future, highlighting the need to account for uncertainty when making decisions this far into the future.




Figure 6. Snapshots of climate projections for the next 100 years. Numbering corresponds to GCM-RCM combinations described in Table 1.

Figure 7 summarizes the performance of each model at every projection snapshot. Eight models predict the system will be under sufficient stress to result in failure in relation to both objectives at one or more of the future time slices, even under optimized operating conditions. It is important to note that in some cases, such as for model 4, “failure” at an earlier time period is followed by a return to the adaptation region at a later date. This highlights the importance of checking system failure at regular time intervals, as only examining system performance at a later time interval could be misleading, as the system might have already failed at an earlier point in time. The results in Figure 7 also show that 12 models predict that operational adaptation will be required at some point in the future in order for the system to remain operational. This highlights the utility of the framework, as it enables an assessment of the extent to which modification of operation policies can delay potentially costly upgrades to infrastructure. For this case study, only two models predict a climate that the system will be able to cope with into the future under its current operating policy, suggesting that a combination of operation policy changes and/or infrastructure upgrades will be necessary in the future.

Model	1	2	3	4	5	6	7	8	9	10	11
2025	●	●	●	●	●	●	●	●	●	●	●
2050	●	●	●	●	●	●	●	●	●	●	●
2075	●	●	●	●	●	●	●	●	●	●	●
2100	●	●	●	●	●	●	●	●	●	●	●
Model	12	13	14	15	16	17	18	19	20	21	22
2025	●	●	●	●	●	●	●	●	●	●	●
2050	●	●	●	●	●	●	●	●	●	●	●
2075	●	●	●	●	●	●	●	●	●	●	●
2100	●	●	●	●	●	●	●	●	●	●	●

Legend	
●	In region of current management policy
●	In region of re-optimized management policy
●	In failure region

Figure 7. Summary of 22 climate projections and their performance over the next 100 years.

mentation would require knowledge of the hydrometeorological state the system is, and will be, exposed to. Such information is difficult to obtain due to both the high uncertainty in future climate projections and the difficulty in separating interannual and interdecadal variability from longer-term systematic climate changes in observed data. This means that, in many cases, it will be challenging to adapt the operating policies to reach the discovered theoretical upper limit.

However, it is possible to use the proposed approach to identify operating policies that are locally robust to a large range of possible climate exposures. In the case of Lake Como, operating policies that are optimal for a specific point in the exposure space will still yield successful performance in other parts of the exposure space. This phenomenon is shown in Figure 8, where an operating policy is optimized for temperature remaining unchanged and precipitation decreasing by 10%. The points enclosed by black circles represent the extent of hydrometeorological states that can be accommodated by this single management policy. By exploiting the feedback loop, these policies are adapting operational decisions to the current system conditions they




Figure 8. Local robustness of operating policies. An operation strategy is optimal for the yellow state, but still results in successful performance for neighboring states enclosed by black circles.

5. Discussion

5.1. Benefit of Identifying the Upper Limit of Adaptive Capacity

This paper presents a novel optimization formulation for identifying the theoretical upper limit to a system's adaptive capacity based on optimal modification of its operating policies. However, practical implementation of the operating policies found through optimization, especially those close to the failure boundary, needs to be addressed. Deciding which of the many potential optimal policies to select for imple-

mentation would require knowledge of the hydrometeorological state the system is, and will be, exposed to. Such information is difficult to obtain due to both the high uncertainty in future climate projections and the difficulty in separating interannual and interdecadal variability from longer-term systematic climate changes in observed data. This means that, in many cases, it will be challenging to adapt the operating policies to reach the discovered theoretical upper limit. However, it is possible to use the proposed approach to identify operating policies that are locally robust to a large range of possible climate exposures. In the case of Lake Como, operating policies that are optimal for a specific point in the exposure space will still yield successful performance in other parts of the exposure space. This phenomenon is shown in Figure 8, where an operating policy is optimized for temperature remaining unchanged and precipitation decreasing by 10%. The points enclosed by black circles represent the extent of hydrometeorological states that can be accommodated by this single management policy. By exploiting the feedback loop, these policies are adapting operational decisions to the current system conditions they are exposed to and, consequently, their performance is acceptable not only for the single climate exposure used in the policy design but is robust for a range of possible climate exposures. This suggests that precise knowledge of which climate trajectory is occurring is not necessary to prevent failure. Although optimized over a single hydrometeorological state, the feedback control policies can attain adequate system performance for neighboring states.

This local robustness discussed in the previous paragraph is different to the "climate states" defined by decision scaling in Brown et al. [2012]. While for the decision scaling approach climate states are the regions of the exposure space for which a particular decision is preferable




Figure 9. Suggested operation for climate projections from now until 2050. The black outlined states show successful performance under operation identified for the yellow state.

compared with another decision, here, the locally robust region of the exposure space corresponds to the hydrometeorological states under which a particular operational strategy does not result in system failure. For the case study, this local robustness is large for regions of the exposure space where the timing and magnitude of inflows is similar, which is shown by both the success subspace of the current operation (Figure 3) and the operation strategy presented in Figure 8. However, the local robustness of a solution is significantly less near failure boundaries, as the operation strategies that are found specifically for these extreme climates are less applicable to nearby climate states. This is a key difference between

other bottom-up approaches that focus on finding solutions that are robust to multiple climate states, like multiobjective robust decision making [Kasprzyk *et al.*, 2013]. Such approaches could find operating strategies that are successful for a wider range of hydrometeorological states, but doing so would sacrifice the optimality necessary to identify the theoretical upper limit to adaptive capacity [Dittrich *et al.*, 2016].

In addition, the local robustness of a solution found using the proposed approach is enough to identify a set of operating policies that can accommodate both the current climate and the most severe set of climate model projections. For the Lake Como case study, by optimizing the system to an 18% precipitation increase and a 4°C temperature increase, it is possible to cater to a large majority of the projected climate conditions for the next 40 years (Figure 9). Note that while this is referred to as a suggested operation, it is used only to demonstrate the robustness of solutions, and due to simplifications in the case study would not be implemented in the actual Lake Como system. There is a large amount of uncertainty in the climate projections shown below, and given many are close to the failure boundary, there is a good possibility that the system would experience failure. However, the selected operating policy does result in “successful” performance for hydrometeorological states changing from current conditions and encompasses a large number of climate model projections. This means that for such an operating policy, the full benefit from the existing system could be obtained.

The inherent local robustness of each optimal solution allows decision makers to accommodate some level of climate uncertainty when determining which management policy they should implement in the future. However, what this robustness also shows is that, in the event where climate conditions were to change to those of the bottom right portion of the exposure space, one management decision would not be sufficient to maintain satisfactory system performance. This suggests that focus on adaptive operating rules, with changes in management based on updated climate information, is the only way to ensure maximum benefit is derived from existing infrastructure in the face of uncertain climate change.

Finally, implementing the proposed framework provides information that informs existing methods of sequential decision making under deep uncertainty [Haasnoot *et al.*, 2013; Beh *et al.*, 2015b] as an upper bound on operational adaptive capacity can be used to determine when adaption options through changes in management are no longer available. Additionally, the superposition of GCM projections (Figure 6) offers insights about the time evolution of the system, including the “sell-by-dates” of given operating policies that must then be adapted to (or combined with other infrastructural actions) in order to maintain

acceptable system performance [Haasnoot *et al.*, 2013; Guillaume *et al.*, 2016]. For example, a possible adaptive policy pathway for the Lake Como case study may consist of implementing the operation shown in Figure 9, until around 2050 when infrastructure upgrades are implemented. It should be noted that the use of optimal operational strategies in sequential decision making depends on the local robustness of the solutions. For the simplified Lake Como case study these regions were quite large, but this will change depending on the physical system and operation mechanisms in place. Also a factor is the method used to generate the exposure space, and as this study used perturbation techniques it is possible that stochastic generation of exposure points will see different behavior for the same system.

5.2. Ability of the Proposed Framework to Identify an Upper Limit

Finding the absolute limit of the adaptive capacity of a system relies on (i) the accurate modeling of stakeholders' demands to identify system failure criteria; (ii) appropriate representation of system exposure; and (iii) the structure and constraints of the operating policy used in the optimization process.

The usefulness of the exposure space depends greatly on the performance of the hydrological model as it converts scaled hydrometeorological variables to runoff. Any error and uncertainty in the simulation should be considered when examining the boundaries that represent the difference between success and failure. Additionally, in this case study, both the physical system model and the stakeholders' demands have been held cyclo-stationary throughout the modelling of different climate exposures. While this is a reasonable assumption for the physical model of the lake as the study is for the existing unchanged infrastructure, it is not a reasonable assumption for the downstream stakeholders' demands, which are likely to change under different climate exposures. In the case of Lake Como, where farmers are a major water user, the water demand may change substantially if climate change was to alter crop types in the area. This assumption is not a limitation of the proposed framework; however, it is a limitation on this case study, as it affects the practical implementation of any results. If a relationship between demand and hydrometeorological states could be described, a nonstationary system model could be applied within this framework to obtain more realistic information about how the system behaves in the vicinity of the threshold for adaptation. One of the benefits of generating an exposure space is that it readily allows for dynamic failure criteria to be incorporated, as the performance at each hydrometeorological state is recorded, so that thresholds can be re-adjusted if the acceptable performance levels change with exposure.

In this study, the failure of the Lake Como system was defined with respect to two hydrometeorological variables, mean annual temperature and annual precipitation volume. To accurately capture adaptive capacity, the failure boundary identified must be the first instance of failure the system experiences. In other words, the variables used to generate the exposure space must have the greatest impact on performance. A sensitivity analysis of the inputs to the hydrological model is recommended to identify what these variables might be. As a result, it may be necessary to represent exposure with three or more variables in order to identify adaptive capacity. While the proposed approach and optimization formulation can be implemented in higher dimensions, the ability to identify the failure boundary accurately also depends on the resolution used to create the exposure space, as mentioned previously. Increasing the number of variables considered will only compound an already large computational effort required to optimize for each hydrometeorological state. As a result, significant reductions in resolution of the exposure space may be required, although advanced sampling methods could be used to ensure good coverage of the higher dimensional exposure space.

The location of the failure boundary of the optimized system is also dependent on the structure and flexibility of the operating policy model used (e.g., the RBF policy parameterization for the Lake Como case study). The decision model must not be constricted in terms of flexibility in release, and must be capable of changing releases in response to extreme climates. As demonstrated in Figure 8, the inherent robustness of an optimal solution found by the RBF model in this case allows for a measure of uncertainty in the knowledge of the future climate exposure when determining which management alternative should be implemented. This supports the evidence that the policies identified using the RBF model represent reservoir decisions well [Giuliani *et al.*, 2014b, 2015], but for other water resource systems, alternative models may be more appropriate. Where possible, the model chosen should not overcalibrate to a particular hydrometeorological state, and provide robustness across neighboring states in the same manner as the RBFs, to reduce the reliance on exact knowledge of the climate state should a decision be implemented.

6. Conclusion

The approach presented in this paper identifies the maximum operational adaptive capacity of water resource systems with respect to future hydrometeorological states. The approach uses formal optimization techniques to identify the optimal response to a future exposure, and allows for dynamic management that updates optimal operating policies as the climate changes. When applied to the Lake Como case study it first defined a failure boundary of the system under current system operations, and then, by using policies optimized for future hydrometeorological states, identified the upper limit of system performance.

It was found that by modifying the operating policies, the Lake Como system can adapt to more than three times as many hydrometeorological states than it would under the current operating policies. While the solutions were generated for a simplified case study and are not designed for the actual Lake Como system, they illustrate the utility of the proposed approach. A key outcome of this study is that the generation of an exposure space provides an informative context for further system performance and climate impact analysis. It demonstrates that when searching for optimal management alternatives, a single operating policy, no matter how robust, does not provide a system with an adaptive capacity as large as an ensemble of operating policies that update as the climate changes. By modifying the current system operations to account for a changing climate, it is hoped that the life of existing water resources systems can be extended, thereby reducing the need for expensive and disruptive physical modifications to water infrastructure.

Acknowledgments

The data used in this study are from Agenzia Regionale per la Protezione dell'Ambiente (<http://ita.arpalombardia.it/ita/inde>) and Consorzio dell'Adda (<http://www.addaconsorzio.it/>). The authors would like to thank ARPA and Eng. Bertoli from Consorzio dell'Adda for providing these data. S. Westra was supported by Discovery project DP120100338.

References

- Anghileri, D., F. Pianosi, and R. Soncini-Sessa (2011), A framework for the quantitative assessment of climate change impacts on water-related activities at the basin scale, *Hydrol. Earth Syst. Sci.*, *15*(6), 2025–2038.
- Arnell, N. W., M. J. L. Livermore, S. Kovats, P. E. Levy, R. Nicholls, M. L. Parry, and S. R. Gaffin (2004), Climate and socio-economic scenarios for global-scale climate change impacts assessments: Characterising the SRES storylines, *Global Environ. Change*, *14*(1), 3–20.
- Ashley, R. C. A. (2006), *Infrastructure to 2030*, OECD Publ., University of Sheffield, U. K.
- Beh, E. H. Y., H. R. Maier, and G. C. Dandy (2015a), Adaptive, multiobjective optimal sequencing approach for urban water supply augmentation under deep uncertainty, *Water Resour. Res.*, *51*, 1529–1551, doi:10.1002/2014WR016254.
- Beh, E. H. Y., H. R. Maier, and G. C. Dandy (2015b), Scenario driven optimal sequencing under deep uncertainty, *Environ. Modell. Software*, *68*, 181–195.
- Bergström, S., and V. Singh (1995), The HBV model, *Computer Models of Watershed Hydrology*, edited by V. P. Singh, Chapter 13: The HBV mode, pp. 1–22, *Water Resour. Publ.*, Highlands Ranch, Colo.
- Boé, J., L. Terray, F. Habets, and E. Martin (2007), Statistical and dynamical downscaling of the Seine basin climate for hydro-meteorological studies, *Int. J. Climatol.*, *27*(12), 1643–1655.
- Brekke, L. D., E. P. Maurer, J. D. Anderson, M. D. Dettinger, E. S. Townsley, A. Harrison, and T. Pruitt (2009), Assessing reservoir operations risk under climate change, *Water Resour. Res.*, *45*, W04411, doi:10.1029/2008WR006941.
- Brown, C., and R. L. Wilby (2012), An alternate approach to assessing climate risks, *Eos Trans. AGU*, *93*(41), 401–402.
- Brown, C., Y. Ghile, M. Laverty, and K. Li (2012), Decision scaling: Linking bottom-up vulnerability analysis with climate projections in the water sector, *Water Resour. Res.*, *48*, W09537, doi:10.1029/2011WR011212.
- Busoni, L., D. Ernst, B. D. Schutter, and R. Babuska (2011), Cross-entropy optimization of control policies with adaptive basis functions, *Trans. Sys. Man Cyber. Part B*, *41*(1), 196–209.
- Castelletti, A., F. Pianosi, and R. Soncini-Sessa (2008), Water reservoir control under economic, social and environmental constraints, *Automatica*, *44*(6), 1595–1607.
- Castelletti, A., S. Galelli, M. Restelli, and R. Soncini-Sessa (2010), Tree-based reinforcement learning for optimal water reservoir operation, *Water Resour. Res.*, *46*, W09507, doi:10.1029/2009WR008898.
- Corani, G., A. E. Rizzoli, A. Salvetti, and M. Zaffalon (2009), Reproducing human decisions in reservoir management: The case of lake Lugano, in *Information Technologies in Environmental Engineering*, edited by I. Athanasiadis, A. Rizzoli, P. Mitkas, and J. Gómez, pp. 252–263, Springer, Berlin.
- Déqué, M. (2007), Frequency of precipitation and temperature extremes over France in an anthropogenic scenario: Model results and statistical correction according to observed values, *Global Planet. Change*, *57*(1–2), 16–26.
- Dittrich, R., A. Wreford, and D. Moran (2016), A survey of decision-making approaches for climate change adaptation: Are robust methods the way forward?, *Ecol. Econ.*, *122*, 79–89.
- Giuliani, M., and A. Castelletti (2016), Is robustness really robust? How different definitions of robustness impact decision-making under climate change, *Clim. Change*, *135*, 409–424.
- Giuliani, M., J. D. Herman, A. Castelletti, and P. Reed (2014a), Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management, *Water Resour. Res.*, *50*, 3355–3377, doi:10.1002/2013WR014700.
- Giuliani, M., E. Mason, A. Castelletti, F. Pianosi, and R. Soncini-Sessa (2014b), Universal approximators for direct policy search in multi-purpose water reservoir management: A comparative analysis, in *Proceedings of the 19th IFAC World Congress*, *47*, 6234–6239, [Available at <http://dx.doi.org/10.3182/20140824-6-ZA-1003.01962>].
- Giuliani, M., A. Castelletti, F. Pianosi, E. Mason, and P. M. Reed (2015), Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations, *J. Water Resour. Plann. Manage.*, *421*(2), 04015050.
- Giuliani, M., D. Anghileri, A. Castelletti, P. N. Vu, and R. Soncini-Sessa (2016), Large storage operations under climate change: Expanding uncertainties and evolving tradeoffs, *Environ. Res. Lett.*, *11*(3), 035009.
- Gleick, P. H. (2003), Global freshwater resources: Soft-path solutions for the 21st century, *Science*, *302*(5650), 1524–1528.
- Guariso, G., S. Rinaldi, and R. Soncini-Sessa (1986), The management of Lake Como: A multiobjective analysis, *Water Resour. Res.*, *22*(2), 109–120.

- Guillaume, J. H. A., M. Arshad, A. J. Jakeman, M. Jalava, and M. Kummu (2016), Robust discrimination between uncertain management alternatives by iterative reflection on crossover point scenarios: Principles, design and implementations, edited by M. Giuliani et al., universal approximators for direct policy search in multi-purpose water reservoir management: A comparative analysis, *IFAC Proceedings Vol-umes*, 47(3), 6234–6239, *Environ. Modell. Software*, 83, 326–343.
- Guo, D., S. Westra, and H. R. Maier (2016), An inverse approach to perturb historical rainfall data for scenario-neutral climate impact studies, *J. Hydrol.*, in press.
- Haasnoot, M., J. H. Kwakkel, W. E. Walker, and J. ter Maat (2013), Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world, *Global Environ. Change*, 23(2), 485–498.
- Hashimoto, T., J. R. Stedinger, and D. P. Loucks (1982), Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation, *Water Resour. Res.*, 18(1), 14–20.
- Herman, J. D., H. B. Zeff, P. M. Reed, and G. W. Characklis (2014), Beyond optimality: Multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty, *Water Resour. Res.*, 50(10), 7692–7713.
- Herman, J. D., P. M. Reed, H. B. Zeff, and G. W. Characklis (2015), How should robustness be defined for water systems planning under change?, *J. Water Resour. Plann. Manage.*, 141(10), 04015012.
- Hyde, K. M., and H. R. Maier (2006), Distance-based and stochastic uncertainty analysis for multi-criteria decision analysis in Excel using Visual Basic for Applications, *Environ. Modell. Software*, 21(12), 1695–1710.
- ICSC (2009), *Proceedings of the 4th International ICSC Symposium Thessaloniki, Greece*, Springer, Berlin.
- IPCC (2013), Summary for policymakers, in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by T. F. Stocker et al., Cambridge University Press, Cambridge, N. Y., U. K.
- IPCC (2014), Summary for policymakers, in *Climate Change 2014: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by C. B. Field et al., pp. 1–32, Cambridge University Press, Cambridge, N. Y., U. K.
- Jacob, D., J. Petersen, B. Eggert, A. Alias, O. B. Christensen, L. M. Bouwer, A. Braun, A. Colette, M. Déqué, and G. Georgievski (2014), EURO-CORDEX: New high-resolution climate change projections for European impact research, *Reg. Environ. Change*, 14(2), 563–578.
- Kasprzyk, J. R., S. Nataraj, P. M. Reed, and R. J. Lempert (2013), Many objective robust decision making for complex environmental systems undergoing change, *Environ. Modell. Software*, 42, 55–71.
- Kingston, G. B., G. C. Dandy, and H. R. Maier (2008), *AI Techniques for Hydrological Modeling and Management*, pp. 67–69, Nova, Hauppauge, N. Y.
- Kwakkel, J. H., W. E. Walker, and M. Haasnoot (2016), Coping with the wickedness of public policy problems: Approaches for decision making under deep uncertainty, *J. Water Resour. Plann. Manage.*, 142, 01816001.
- Lempert, R., and M. T. Collins (2007), Managing the risk of uncertain threshold responses: Comparison of robust, optimum, and precautionary approaches, *Risk Anal.*, 27(4), 1009–1026.
- Lempert, R., N. Nakicenovic, D. Sarewitz, and M. Schlesinger (2004), Characterizing climate-change uncertainties for decision-makers. An editorial essay, *Clim. Change*, 65(1), 1–9.
- Li, F.-F., C. A. Shoemaker, J. Qiu, and J.-H. Wei (2015), Hierarchical multi-reservoir optimization modeling for real-world complexity with application to the Three Gorges system, *Environ. Modell. Software*, 69, 319–329.
- Loucks, D. P., and O. T. Sigvaldason (1981), Multiple reservoir operation in North America, paper presented at Surface Water Impoundments, ASCE, Laxenbourg, Austria.
- Maier, H. R., et al. (2014), Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions, *Environ. Modell. Software*, 62, 271–299.
- Maier, H. R., Z. Kapelan, J. Kasprzyk, and L. S. Matott (2015), Thematic issue on evolutionary algorithms in water resources, *Environ. Modell. Software*, 69, 222–225.
- Maier, H. R., J. H. A. Guillaume, H. van Delden, G. A. Riddell, M. Haasnoot, and J. H. Kwakkel (2016), An uncertain future, deep uncertainty, scenarios, robustness and adaptation: How do they fit together?, *Environ. Modell. Software*, 81, 154–164.
- Mastrandrea, M., N. Heller, T. Root, and S. Schneider (2010), Bridging the gap: Linking climate-impacts research with adaptation planning and management, *Clim. Change*, 100(1), 87–101.
- Milly, P. C. D., J. Betancourt, M. Falkenmark, R. M. Hirsch, Z. W. Kundzewicz, D. P. Lettenmaier, and R. J. Stouffer (2008), Climate change—Stationarity is dead: Whither water management?, *Science*, 319(5863), 573–574.
- Paton, F. L., H. R. Maier, and G. C. Dandy (2013), Relative magnitudes of sources of uncertainty in assessing climate change impacts on water supply security for the southern Adelaide water supply system, *Water Resour. Res.*, 49, 1643–1667, doi:10.1002/wrcr.20153.
- Paton, F. L., H. R. Maier, and G. C. Dandy (2014), Including adaptation and mitigation responses to climate change in a multiobjective evolutionary algorithm framework for urban water supply systems incorporating GHG emissions, *Water Resour. Res.*, 50, 6285–6304, doi: 10.1002/2013WR015195.
- Poff, N. L., C. M. Brown, T. E. Grantham, J. H. Matthews, M. A. Palmer, C. M. Spence, R. L. Wilby, M. Haasnoot, G. F. Mendoza, and K. C. Dominique (2015), Sustainable water management under future uncertainty with eco-engineering decision scaling, *Nat. Clim. Change*, in press.
- Prudhomme, C., R. L. Wilby, S. Crooks, A. L. Kay, and N. S. Reynard (2010), Scenario-neutral approach to climate change impact studies: Application to flood risk, *J. Hydrol.*, 390(3–4), 198–209.
- Ravalico, J. K., G. C. Dandy, and H. R. Maier (2010), Management option rank equivalence (MORE)—A new method of sensitivity analysis for decision-making, *Environ. Modell. Software*, 25(2), 171–181.
- Risbey, J. S. (2011), Dangerous climate change and water resources in Australia, *Reg. Environ. Change*, 11, S197–S203.
- Salazar, J. Z., P. M. Reed, J. D. Herman, M. Giuliani, and A. Castelletti (2016), A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control, *Adv. Water Resour.*, 92, 172–185.
- Soncini Sessa, R., Castelletti, A. and Weber, E. (2007), *Integrated and Participatory Water Resources Management Theory*, Elsevier Science, Amsterdam.
- Steinschneider, S., and C. Brown (2013), A semiparametric multivariate, multisite weather generator with low-frequency variability for use in climate risk assessments, *Water Resour. Res.*, 49(11), 7205–7220.
- Steinschneider, S., R. McCrary, S. Wi, K. Mulligan, L. Mearns, and C. Brown (2015), Expanded decision-scaling framework to select robust long-term water-system plans under hydroclimatic uncertainties, *J. Water Resour. Plann. Manage.*, 141, 04015023.
- Tsoukalas, I., and C. Makropoulos (2015), Multiobjective optimisation on a budget: Exploring surrogate modelling for robust multi-reservoir rules generation under hydrological uncertainty, *Environ. Modell. Software*, 69, 396–413.

- Turner, S. W. D., D. Marlow, M. Ekström, B. G. Rhodes, U. Kularathna, and P. J. Jeffrey (2014), Linking climate projections to performance: A yield-based decision scaling assessment of a large urban water resources system, *Water Resour. Res.*, *50*, 3553–3567, doi:10.1002/2013WR015156.
- Vano, J. A., M. J. Scott, N. Voisin, C. O. Stockle, A. F. Hamlet, K. E. B. Mickelson, M. M. Elsner, and D. P. Lettenmaier (2010), Climate change impacts on water management and irrigated agriculture in the Yakima River Basin, Washington, USA, *Clim. Change*, *102*(1–2), 287–317.
- Weaver, C. P., R. J. Lempert, C. Brown, J. A. Hall, D. Revell, and D. Sarewitz (2013), Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks, *WIREs Clim. Change*, *4*(1), 39–60.
- Whateley, S., S. Steinschneider, and C. Brown (2014), A climate change range-based method for estimating robustness for water resources supply, *Water Resour. Res.*, *50*(11), 8944–8961.
- Wilby, R. L., and S. Dessai (2010), Robust adaptation to climate change, *Weather*, *65*(7), 180–185.
- Yang, T., X. Gao, S. L. Sellars, and S. Sorooshian (2015), Improving the multi-objective evolutionary optimization algorithm for hydropower reservoir operations in the California Oroville–Thermalito complex, *Environ. Modell. Software*, *69*, 262–279.