Environmental Modelling and Software 141 (2021) 105047

Contents lists available at ScienceDirect

Environmental Modelling and Software

journal homepage: http://www.elsevier.com/locate/envsoft

How do the properties of training scenarios influence the robustness of reservoir operating policies to climate uncertainty?

Jonathan S. Cohen [a,*], Harrison B. Zeff [b], Jonathan D. Herman [a]

[a] Department of Civil and Environmental Engineering, University of California, Davis, CA, USA
[b] Department of Environmental Sciences and Engineering, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

ARTICLE INFO

Keywords:
Policy search
Reservoir operations
Climate adaptation
Robustness
Scenario selection

ABSTRACT

Reservoir control policies provide a flexible option to adapt to the uncertain hydrologic impacts of climate change. This challenge requires robust policies capable of navigating scenarios that are wetter, drier, or more variable than anticipated. While a number of prior studies have trained robust policies using large scenario ensembles, there remains a need to understand how the properties of training scenarios impact policy robustness. Specifically, this study investigates scenario properties including annual runoff, snowpack, and baseline regret—the difference between baseline policy and perfect foresight performance in an individual scenario. Results indicate that policies trained to scenario subsets with high baseline regret outperform those generated with other training sets in both wetter and drier futures, largely by adopting an intra-annual hedging strategy. The approach highlights the potential to improve the efficiency and robustness of policy training by considering both the hydrologic properties and baseline regret of the training ensemble.

1. Introduction

Adaptation to the multi-scale impacts of climate change in water resources systems is challenged by substantial uncertainty in future hydrologic projections, particularly in flood and drought risks (Wilby and Dessai, 2010; Asadieh and Krakauer, 2017; Dottori et al., 2018). This hinders the ability to use traditional prediction-based planning methods and has resulted in the recent consensus toward robust planning (Dessai and Hulme, 2004; Wilby and Dessai, 2010). Robust and adaptive planning have been widely considered for both expansion of water resources infrastructure (e.g. Haasnoot et al., 2013; Beh et al., 2015; Zeff et al., 2016; Kwakkel et al., 2016; Maier et al., 2016; Trindade et al., 2017) as well as changes to reservoir control policies (e.g. Giuliani et al., 2014; Quinn et al., 2018; Herman and Giuliani, 2018). Of the two alternatives, control policies provide a more flexible "soft path" approach, as they can be reversed if the future unfolds differently than predicted (Gleick, 2002; Fletcher et al., 2017). In this case, the physical constraints of the existing system establish the range of uncertain scenarios that can be adapted to before new infrastructure is needed (e.g. Culley et al., 2016).

Robust planning of reservoir control policies generally consists of two phases that have been studied using a variety of different approaches: policy design and robustness analysis. A number of studies have focused on the robustness of current system operations to a range of future climate changes represented either by downscaled Global Circulation Model (GCM) scenarios (e.g., Brekke et al., 2009; Karamouz et al., 2013; Knowles et al., 2018) or synthetically generated scenarios based on perturbed statistics of hydrologic timeseries (e.g., Prudhomme et al., 2010; Brown et al., 2012; Weaver et al., 2013; Turner et al., 2014). Both approaches often serve as precursors to adaptation studies in which a discrete set of proposed management alternatives are tested to mitigate vulnerabilities in future scenarios (e.g., Groves et al., 2013; Steinschneider et al., 2015a; Mateus and Tullos, 2017). In this case, the policies are not trained or optimized to a particular set of scenarios, but instead arise from stakeholder expertise and negotiations.

An alternative approach is to generate candidate alternatives via optimization approaches (Kasprzyk et al., 2013). In this case, policy design and robustness analysis are analogous to the train-test terminology often used in machine learning (e.g., Russell and Norvig, 2002), and recently in the water resources field (e.g., Brodeur et al., 2020). Policy design (training) involves optimizing learned policy parameters to a specific set of input data (the training set). In robustness analysis (testing), a test set consisting of input data separate from the training set are used to assess the performance of an optimized policy. In the context

* Corresponding author.
E-mail address: joncohen@ucdavis.edu (J.S. Cohen).

https://doi.org/10.1016/j.envsoft.2021.105047
Accepted 25 March 2021
Available online 2 April 2021

of reservoir control under climate change, the most relevant optimization approach is policy search, in which parameterized operating rules are optimized for system performance objectives under a set of training scenarios (Koutsoyiannis and Economou, 2003; Giuliani et al., 2015a, 2017). This heuristic approach functions as both a simulation-optimization problem (Salazar et al., 2016) and an information selection problem for the policy inputs (Giuliani et al., 2015b; Nayak et al., 2018). While the training performance of a policy on an individual scenario represents a best-case outcome with perfect foresight, the key challenge is whether the policy can generalize to a different test set, which is also the case for any optimization method applied in the context of climate adaptation. Test sets often include additional stochastic realizations of the same uncertainties used in training, i.e., to obtain a thorough representation of sampling uncertainty (e.g., Quinn et al., 2017; Trindade et al., 2017). They may also include scenarios representing a different characterization of uncertainty altogether (Watson and Kasprzyk, 2017; Eker and Kwakkel, 2018).

The robustness of alternatives generated by policy search therefore strongly depends on the choice of scenarios used for training and testing, both in terms of coarse-scale statistics (wet vs. dry) and the realizations of natural variability that lead to extreme events (Herman et al., 2020). This includes the case where the training data represent a baseline or historical scenario (e.g., Kasprzyk et al., 2013; Giuliani and Castelletti, 2016; Quinn et al., 2018). These studies evaluate resulting alternatives over test scenarios spanning a wide range of potential future hydrology, but without analyzing the influence of the training scenarios on robustness. Robust optimization studies overcome this by optimizing robustness metrics over many samples of uncertain scenarios (e.g., Hamarat et al., 2014; Kwakkel et al., 2015). While they optimize over a large range of training scenarios, studies using robust optimization generally have not considered how well solutions can generalize out-of-sample. Robustness measures will typically include either regret, which quantifies the cost of choosing an incorrect solution, or satisficing, which calculates the fraction of scenarios in which a policy meets a set of performance criteria (Lempert and Collins, 2007; Herman et al., 2015). While both of these methods are effective in evaluating the performance of individual solutions, there is also the consideration of the robustness of the Pareto set as a whole to quantify deviations in multi-objective performance.

While the properties of the test scenarios have been the focus of many prior studies using scenario discovery and related methods, the properties of the training scenarios and their influence on policy robustness have received relatively less attention. However, several recent studies have begun to analyze the effect of the choice of training scenarios in optimization problems. For example, Watson and Kasprzyk (2017) extend many-objective robust decision making by optimizing to several different sets of scenarios with varying properties. They then re-evaluate solutions in out-of-sample scenarios, quantifying robustness for individual solutions using the satisficing metric. Eker and Kwakkel (2018) optimize to scenarios with maximum diversity and policy relevance and re-evaluate solutions under the same uncertainty characterization used in training. Giudici et al. (2020) propose an algorithm to select the smallest subset of training scenarios which can be used to generate robust solutions when re-evaluated against the full set to minimize computational cost. These studies all effectively aim to find training scenarios (scenario selection) that lead to robust out of sample performance. However, studies to date have not attributed multi-objective policy robustness to the hydrologic properties of the training scenarios, which holds significant implications for the design of robust policies under climate uncertainty.

This study proposes an experimental design to determine how the properties of forcing scenarios influence the robustness of multi-objective policy alternatives across several combinations of test scenarios. This framework is generalizable to any environmental planning problem that includes a no action case, an optimization component, and an ensemble of forcing scenarios exhibiting uncertainty. For the initial application, we focus specifically on how the hydrologic properties of climate scenarios influence reservoir policy alternatives. One additional scenario property is the baseline regret, which quantifies the extent to which policy search can improve upon the status quo based on a perfect foresight optimization for an individual scenario. Scenarios are clustered into groups with similar hydrology via unsupervised learning, and split into different combinations of training and test sets. The robustness of the resulting policies is quantified relative to the perfect foresight and baseline solutions to ensure that, at a minimum, all solutions outperform the baseline no-action policy. This is done with a normalized hypervolume metric to represent the robustness of the Pareto-set as a whole, which simultaneously quantifies the changes in the performance of the solutions as well as their diversity when re-evaluated on a given test set. Finally, we perform hypothesis tests on several iterations of the train-test split to identify the properties of training scenarios that lead to the most robust results for each test set, with a particular focus on training policies to scenarios which have high baseline regret. We demonstrate this approach using a simulation model of the northern California reservoir system coupled with an ensemble of transient downscaled climate scenarios.

2. Case study

2.1. Northern California reservoir system

To support urban and agricultural growth amid intense intra- and inter-annual variability in hydrology, California has built a complex system of water resources infrastructure. Reservoirs in the foothills of the Sierra Nevada range capture winter and spring flood season flows to be delivered for agriculture and municipal supply, particularly during summer months. The State Water Project (SWP) and federal Central Valley Project (CVP) consist of a number of reservoirs and aqueducts throughout the Sacramento-San Joaquin river basin. The terminal delta of this system is the site of pumped water exports from north to south, which support agriculture and municipal supply in the southern portion of the state. Critical environmental requirements related to the salinity of outflows from the Delta are a major constraint on these exports. Delta exports are a key metric for water supply reliability in the state and have been found vulnerable to climate change, due to both changes in precipitation levels and seasonal runoff timing (Anderson et al., 2008; Ray et al., 2020).

In the Sacramento River Basin, three of the largest Sierra foothill reservoirs by volume (Shasta, Oroville, and Folsom) combine to a total of 9 million acre-feet (11.1 km^3) of storage in parallel. These reservoirs play a major role in balancing the state's human and environmental water needs. Carryover storage in these reservoirs, measured at the end of the water year (September 30), is a strong indicator of overall system performance and potential economic vulnerabilities (Draper and Lund, 2004). Uncertainty in changing inter-annual precipitation patterns and reduced snowpack levels has the potential to be detrimental to carryover storage levels and their economic benefits (Medellín-Azuara et al., 2008). Under a variety of projected changes to the hydrologic regime, operational adaptations are needed to maintain carryover storage levels to support multiple environmental and water supply related objectives while continuing to provide adequate flood control functions (Cohen et al., 2020).

2.2. Simulation model (ORCA)

We use the open source model Operation of Reservoirs in California (ORCA) to simulate the northern California reservoir system (https://github.com/jscohen4/orca/tree/cohen-2021-properties-training-scenarios). ORCA is a simulation model that runs on a daily timestep and accurately reproduces historical operations (Cohen et al., 2020). The operating rules that drive ORCA are used as the baseline policy in the

current study. The model simulates the major components of the California system north of the Delta, including the Shasta, Oroville, and Folsom Reservoirs, and Delta water supply exports via the Harvey O. Banks (SWP) and Tracy (CVP) pumping plants (Fig. 1a and b). While not as spatially comprehensive as several other statewide models, ORCA is a pure simulation model, which allows for flexible adjustments to operating rules and straightforward evaluation of alternative hydrologic scenarios, as required by the proposed set of policy search experiments.

ORCA is driven by a basic mass balance update for each reservoir. Based on timestep t, storage S_t^r in reservoir r is updated based on inflows Q_r^t, evaporative losses L_t^r, and a release u_t^r:

$$S_t^r = S_{t-1}^r + Q_t^r - u_t^r - L_t^r \tag{1}$$

A target release RT_t^r is determined by the greatest of three minimum operating requirements that must be satisfied for each reservoir:

$$RT_t^r = \max\left(u_{t,\text{environment}}^r, u_{t,\text{flood}}^r, u_{t,\text{demand}}^r\right) \times c_t^r \tag{2}$$

The first is a minimum environmental flow requirement $u_{t,\text{environment}}^r$ that varies based on the time of year and water year type. The second is a flood control release target $u_{t,\text{flood}}^r$. The flood control release depends on a dynamic flood control rule curve, which is determined by a flood control index based on the previous day's precipitation, and current reservoir storage. Finally, the minimum demand release $u_{t,\text{demand}}^r$ consists of water supply demands north of the Delta, south of Delta demands to be delivered by Banks and Tracy pumping plants, and a Delta outflow demand for environmental benefits and salinity control. These water demands are also partially controlled by current and projected reservoir storage, creating a feedback between reservoir operations and downstream Delta management.

A snowpack-to-streamflow forecast enables projections of reservoir inflows throughout the irrigation season. This forecast also determines the water year type classification, which influences both environmental flow requirements and water supply demands along with other operational parameters. The snowpack-to-streamflow forecast is altered by an exceedance level $\overline{Z}_{WYI}$ for the water year type prediction and $\overline{Z}_{wyt}^r$ for individual reservoir inflows. The exceedance represents the confidence in the forecast. A lower exceedance level indicates a more conservative forecast, resulting in lower inflow forecasts and drier water year type classifications, and likely hedged reservoir releases. The forecast is updated each day in the simulation via Equation (3):

$$Qf_t^r = \beta_{dw_t}^r SWE_t^r + \alpha_{dw_t}^r + \overline{Z}_{wyt}^r \sigma_{dw_t}^r \tag{3}$$

Where Qf_t is the forecasted inflow for the remainder of the water year at reservoir r, SWE_t is the snow water equivalent at day t, $\beta_{dw_t}^k$ and $\alpha_{dw_t}^k$ are regression coefficients for day of water year dw_t based on historical streamflow records, and $\sigma_{dw_t}^k$ is the standard deviation of remaining streamflow on dw_t, also based on the historical record.

A curtailment multiplier c_t^r can hedge releases in cases where the system is not projected to meet a carryover storage target C_{wyt}^r at the end of the water year. The forecasted flow and current reservoir storage are used to determine what curtailment multiplier would be necessary to meet the carryover target (Equation (4)). The curtailment multiplier is also constrained by a maximum curtailment allowance $c_{\max,wyt}^r$. A higher maximum curtailment will allow for lower releases to occur in the irrigation season to maintain the cold-pool carryover storage. The daily curtailment multiplier during between May and September ($5 \le M \le 9$) is determined at each timestep via Equation (4):

$$c_{t,5\le M\le 9}^r = \min\left(1, \max\left\{\frac{Qf_t^r + S_{t-1}^r - \sum_{dw_t}^{365} RT_t}{C_{wyt}^r}, c_{\max,wyt}^r\right\}\right) \tag{4}$$

The release for each reservoir is then equal to the target release RT_t^k times the curtailment factor c_t^k:

$$u_t^r = RT_t^r \times c_t^r \tag{5}$$

Further details concerning operations modeled in ORCA are described in Cohen et al. (2020).

2.3. Data sources

Several hydroclimatic time series are used as inputs for the simulation model. These include daily streamflows, spatially gridded and site-specific precipitation, and air temperature, along with spatially averaged and site specific monthly spatial snow water equivalent (SWE). For simulating historical operations, these data are obtained from the California Data Exchange Center (CDEC, 2018). Downscaled CMIP5 Climate and Hydrology Projections are obtained from the United States Bureau of Reclamation (USBR) (Reclamation, 2013; Brekke et al., 2014). These

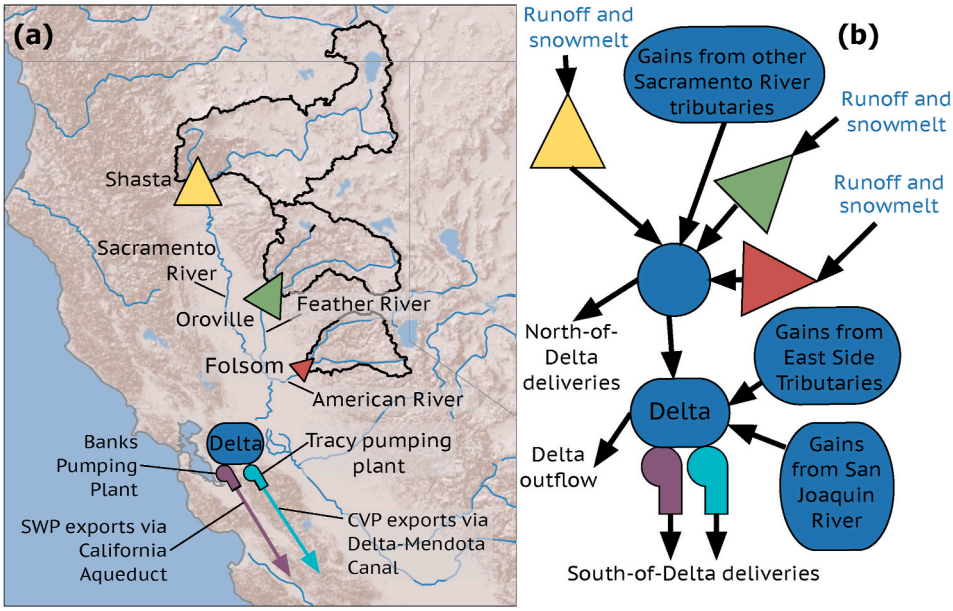

Fig. 1. (a) Map of northern California reservoir system modeled in ORCA. (b) Model schematic showing primary storage and pumping infrastructure.

consist of 31 GCMs simulated for various emissions scenarios to generate 97 scenarios of precipitation and temperature on a daily timestep through 2100 (see Section 1 in the Supplementary Material for a list of institutions providing GCM projections). In the USBR study, outputs from these GCM simulations were routed through the Variable Infiltration Capacity (VIC) model (Liang et al., 1994) calibrated for each basin, yielding additional streamflow and SWE projections to serve as model inputs. The choice to use a GCM ensemble in this study reflects several considerations. First, it provides the best available representation of physically-based transient changes to hydrology, including extreme events, despite the known limitations of GCM projections (Herman et al., 2020). Second, it provides an accurate link between hydrologic variables across space and time, linking precipitation, streamflow, temperature, and snowpack in multiple basins. This is difficult to achieve with synthetic generators, though these are rapidly improving for this purpose (e.g., Steinschneider et al., 2019).

This ensemble exhibits the high degree of uncertainty associated with future precipitation and, to a lesser extent, temperature. In Fig. 2a, the trajectories of annual streamflow show an end-of-century average annual flow ranging from −/+ 50% of historical values. This creates a challenge for how to best adapt operations to balance the tradeoff between flood control and water supply (Herman and Giuliani, 2018; Nayak et al., 2018). All scenarios in the ensemble show a decline in snowpack, ranging from 20 to 90% of the historical average. This is one of the best predicted aspects of climate change, although uncertainties exist in the extent and severity of this thermodynamic-hydrologic change (Cayan et al., 2001; Klos et al., 2014). This can be particularly impactful in mountainous regions where snowpack has historically functioned as a natural reservoir (Rhoades et al., 2018). As a result, the primary impact of rising temperatures in the region is earlier spring snowmelt timing (Knowles et al., 2006; Kapnick and Hall, 2010). These intra-annual streamflow shifts are predicted throughout the CMIP5 ensemble based on the water year centroid, a representation of the center of mass of the annual hydrograph (Fig. 2c). The ensemble also shows uncertainty in the extent of flood risk changes (Fig. 2d) based on both uncertain dynamic climate changes as well as the potential increases given a shift from snow to rain and more rain-on-snow events (McCabe et al., 2007; Surfleet and Tullos, 2013; Huang et al., 2018). Lastly, the ensemble shows severity in uncertainty related to changes in drought patterns (Fig. 1e). Neither changes in drought nor flood statistics show a reliable relationship to the overall changes in total annual streamflow. Overall, the downscaled projection ensemble exhibits the significant uncertainty typical of climate adaptation studies, requiring careful attention to the choice of scenarios under which reservoir control policies are trained and tested.

3. Methods

The proposed experiments aim to analyze reservoir policy performance on held-out climate projections by selecting different subsets of training scenarios based on their hydrologic properties and a baseline regret property. The experiments require several components (Fig. 3): (1) policy search, which is used at several steps throughout the experiment; (2) multi-objective baseline regret, which uses perfect foresight optimization to determine the upper bound of system performance in each scenario; (3) unsupervised clustering of the scenario ensemble, taking into account baseline regret as well as hydrologic properties to determine train-test splits; and (4) the robustness of policies trained to one set of scenarios when evaluated on another set. Finally, we analyze the decision variables and dynamics of several robust policies identified

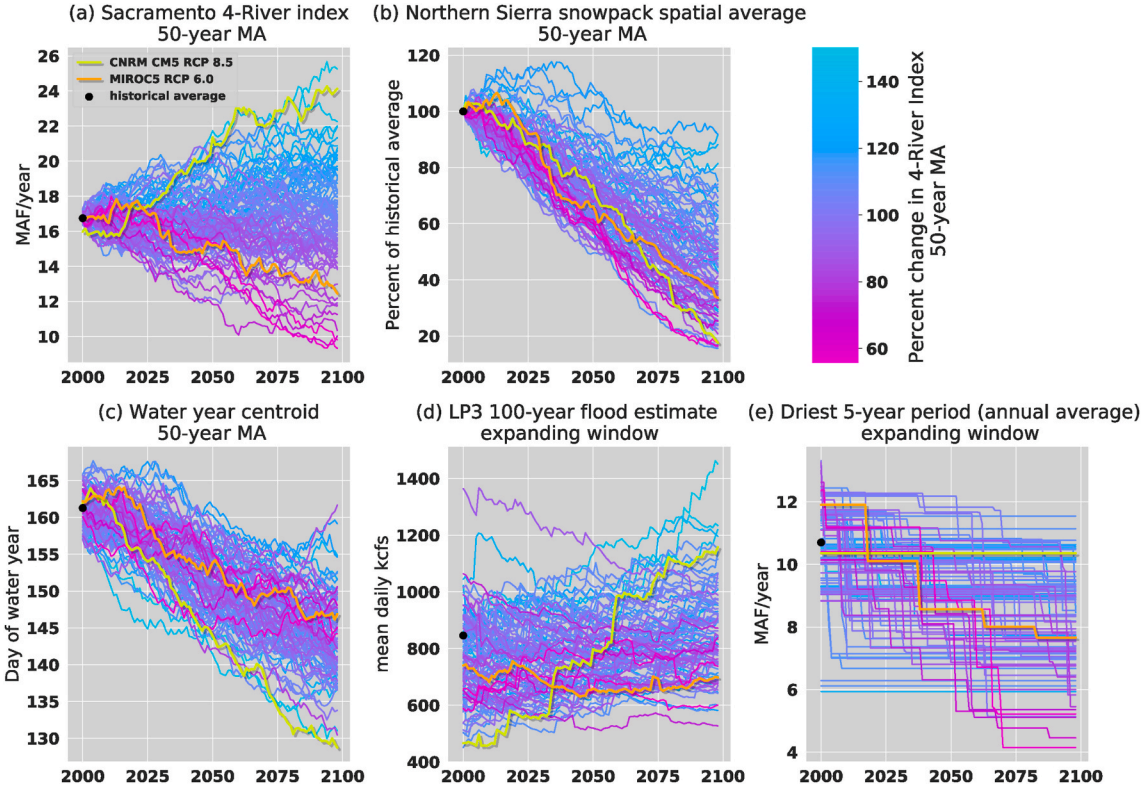

Fig. 2. 50-year moving averages of CMIP5 projections showing a wide range of uncertainty in future flood and drought risk. (a) Streamflow in the Sacramento River downstream of its three largest tributaries (the Feather, Yuba, and American Rivers), denoting the 4-river index. (b) snowpack in the northern Sierra Nevada. (c) water year centroid, defined as the day of the water year at which half of the total annual streamflow has been observed. (d) Log-Pearson Type III distribution (LP3) 100-year flood estimate for the Sacramento River flow below the American River. (e) The driest 5-year period in the 4-river index. The scenarios shown in yellow (RCP 8.5) and orange (RCP 6.0) are examined later in the analysis (Section 4.3). See Section 1 in Supplementary Material for institutions providing climate models and model abbreviations.

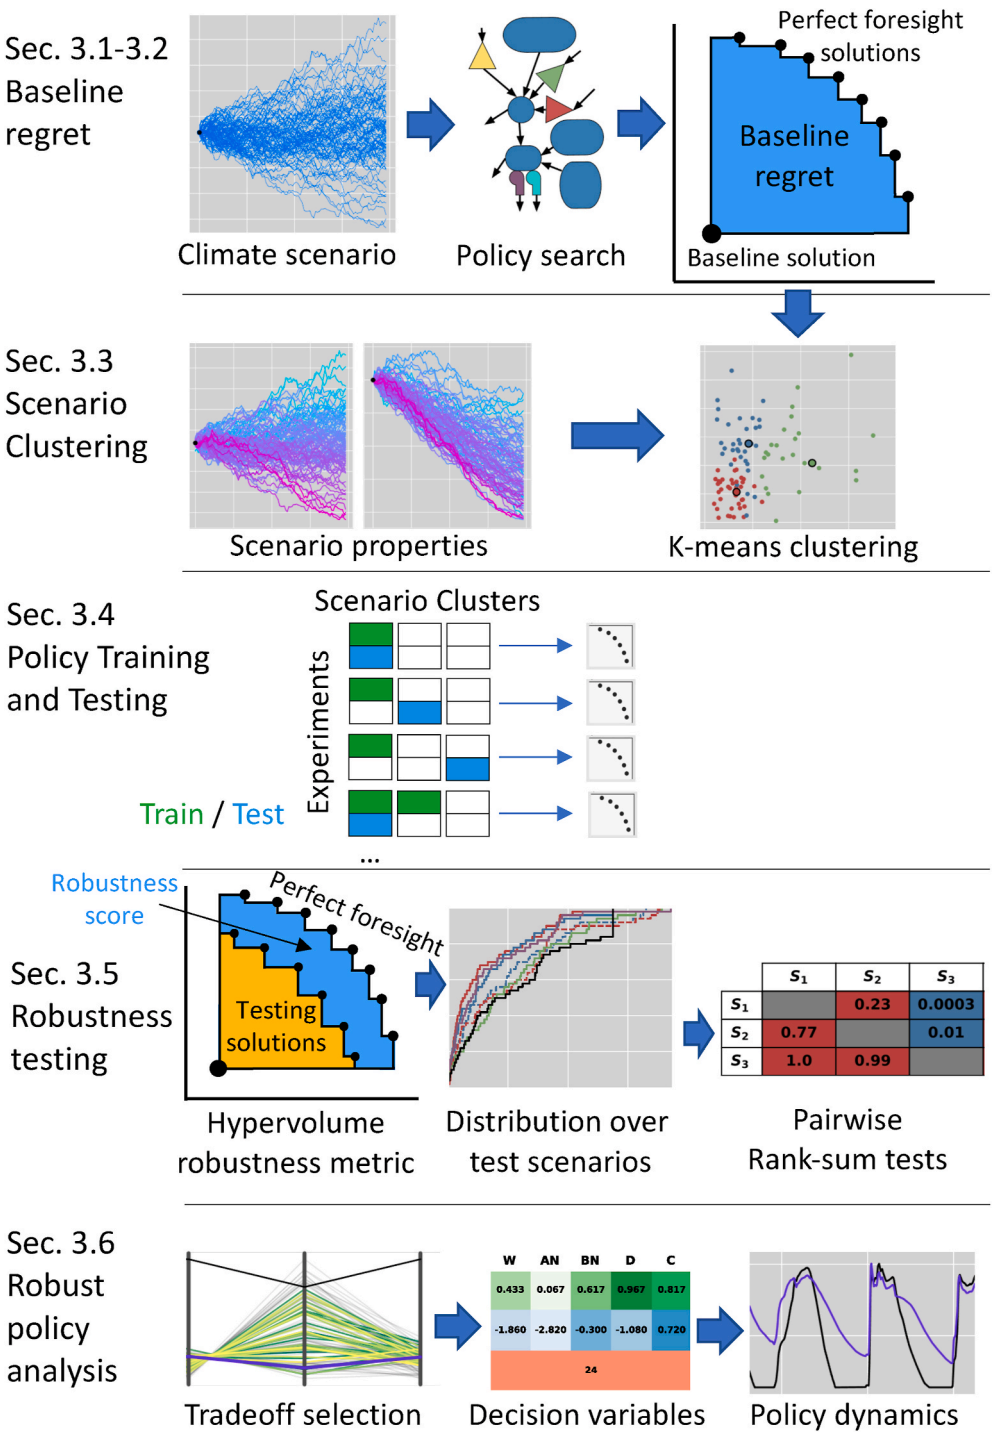

Fig. 3. Flow diagram of methods.

using different training sets in the policy search.

3.1. Policy search

We employ multi-objective policy search throughout this study to identify operational adaptations by parameterizing the structure of existing rules. In this study, policy search aims to solve the optimization problem:

$$\boldsymbol{\theta}^* = \text{argmax}_\theta \left(\frac{1}{n} \sum_s \boldsymbol{J}(\theta, s) \right) \quad (6)$$

Where $s \in \boldsymbol{S}$ are the training scenarios over which the particular optimization occurs, and n is the number of scenarios in set S. θ is the vector of decision variables representing the parameters of the operating policy, and $\boldsymbol{J}$ is the vector of objective functions. $\boldsymbol{\theta}^*$ is the set of policies which correspond to the Pareto-optimal solutions. Thus, the policy search attempts to maximize the expected value of objectives $\boldsymbol{J}$ in scenarios s across set $\boldsymbol{S}$. In the specific case of a perfect foresight optimization, $\boldsymbol{S}$ includes only one training scenario.

The alternatives represented by the decision variables include: a revised snowpack-to-streamflow forecasting method, updated release curtailment rules, and changes in the timing of a dynamic flood control

curve. Specifically:

$$\theta = \begin{bmatrix} \overline{Z}_{WYI} & \\ \overline{Z}^{r}_{wyt} & \forall r, wyt \\ c^{r}_{\max,wyt} & \forall r, wyt \\ FS^{r} & \forall r \end{bmatrix} \tag{7}$$

Where $\overline{Z}_{WYI}$ is the forecast exceedance level to determine water year indices and $\overline{Z}^{r}_{wyt}$ is the rest-of-year inflow forecast exceedance level for each reservoir r in each water year type *wyt*. $c^{r}_{\max,wyt}$ is the maximum curtailment ratio, and FS^{r} represents a shift of the reservoir flood-control refill period earlier in the water year. Given the $n_r = 3$ reservoirs and $n_{wyt} = 5$ water year types, this leads to a total of $1 + n_r(1 + 2n_{wyt}) = 34$ decision variables for each optimization. In prior work these individual actions have been shown to improve system performance by enumeration as snowpack decline continues later in the century (Cohen et al., 2020), but their effect in tandem has not yet been analyzed as a policy search problem. The decision variables are optimized using a normalized set, with 60 discrete values, in order for consistency with alternatives defined in Cohen et al. (2020). Values from the normalized set are transformed in the model to reflect the actual bounds for each variable from the previous study: $\overline{Z}_{WYI}, \overline{Z}^{r}_{wyt} \in [-3.6, 3.6]$, $c^{r}_{\max,wyt} \in [0, 1]$, and $FS^{r} \in [0, 60]$. The choice of decision variables reflecting system parameters rather than a universal approximator function, such as a neural network (e.g., Salazar et al., 2016; Giuliani et al., 2017), is intended to support the interpretability of the resulting policies, as well as their compatibility with already in-place system operations. However, there are currently efforts to formulate methods that improve the interpretability of neural network-based policies, for example via sensitivity analysis (Quinn et al., 2019).

The objectives $\boldsymbol{J}$ contain five performance metrics, including flood control, reservoir carryover storage at the end of the water year, Delta outflow representing salinity control and environmental benefits, Delta exports for water supply, and hydropower generation. The expectation of each objective across a scenario set is to be maximized (Equation (6)). The objective values over each scenario are calculated according to:

$$J_{\text{Flood}}(\theta) = -\sum_{t=1}^{T}\sum_{r=1}^{3} \max\left(u^{r}_{t} - DQ^{r}, 0\right)^{2} \tag{8}$$

$$J_{\text{Carryover}}(\theta) = \sum_{y=1}^{N}\sum_{r=1}^{3} Cr^{r}_{y} \tag{9}$$

$$J_{\text{Outflow}}(\theta) = \sum_{t=1}^{T}\left[Q_{in,t} - (TRP_t + HRO_t)\right] \tag{10}$$

$$J_{\text{Exports}}(\theta) = \sum_{t=1}^{T}\left[TRP_t + HRO_t\right] \tag{11}$$

$$J_{\text{Hydro}}(\theta) = \sum_{t=1}^{T}\sum_{r=1}^{3} HP^{r}_{t} \tag{12}$$

Where T is the number of days t in the simulation period, while N is the number of water years y. In the flooding objective, DQ^{r} is the downstream levee capacity of reservoir r. Note in Equation (8) we maximize negative flooding for consistency with maximization of the other objectives. Cr^{r}_{y} is the carryover storage in reservoir r at the end of water year y. $Q_{in,t}$ is the Delta inflow on day t, while TRP_t (Tracy pumping plant) and HRP_t (Harvey O. Banks pumping plant) are exports from the Delta to the Central Valley Project and State Water Project, respectively. Lastly, HP^{r}_{t} represents the hydropower production from reservoir r on day t. These objective functions are intended to capture the necessary balance between key aspects of system performance.

The optimization is performed using the Non-Dominated Genetic Sorting Algorithm (NSGAIII) (Deb and Jain, 2013) via the open source Platypus library (Hadka, 2015). To support this choice, algorithm performance was tested over 70 scenarios with three random trials of 50,000 maximum number of function evaluations (NFE) each, obtaining similar results for each trial. The results were compared with alternative MOEAs, including ε-MOEA, NSGAII, and SPEA2, which showed no significant improvement over NSGAIII for this problem. In further instances where this problem is solved, 10,000 NFE are used when optimizing to a single scenario (perfect foresight), while all other optimizations with various train-test splits use 50,000 NFE. As the number of scenarios increases, the optimization is slower to converge visually based on hypervolume, influencing the choice of 10,000 vs 50,000 NFE. All optimization runs were performed on the HPC1 cluster at UC Davis, which includes 60 nodes with 16 cores each running at 2.4 GHz.

3.2. Baseline regret

Multi-objective baseline regret quantifies the maximum level to which system performance can be improved for a particular scenario, within the constraints imposed by the policy function and existing infrastructure. This property is scenario-specific: each time it is calculated the baseline policy is held constant, while the scenario differs. This concept draws from the Expected Value of Perfect Information (EVPI) metric proposed by Giuliani et al. (2015b) for multi-objective problems. By incorporating the results of a perfect foresight optimization, it couples the EVPI approach with a regret metric (Savage, 1951), which describes the performance of a policy based on its distance from the best possible alternative. Traditional decision making under uncertainty problems often use the minimax regret approach, in which the goal is to choose the alternative that minimizes the maximum regret across all scenarios (e.g., Giuliani and Castelletti, 2016). The baseline regret metric differs from minimax regret because it applies only to the no action case of an individual scenario, rather than policy alternatives, reflecting the EVPI approach. It is based on the performance of both a baseline solution and perfect foresight solution set for each scenario, as the performance of any other effective policy solution is expected to be bounded by these two. As a result, the baseline regret partially depends on the suitability of the baseline policy for each climate scenario. However, this still reflects the ability of the system to adapt to future change, even if it is starting from a poor baseline.

3.2.1. Baseline policy performance and perfect foresight optimization

The baseline policy simulation uses parameters θ_B to best represent the dynamics of the system shown in historical observations. The resulting baseline policy solution performance is denoted as $J_B(s) = J(\theta_B, s)$, a one-dimensional vector containing a single value for each objective, rather than a full Pareto set. Under the baseline policy, this performance is not optimized; it should be viewed as a simplified representation of the several performance considerations of a real-world system operator.

The upper bound performance is established by a perfect foresight optimization. In this case, the policy search is performed over each scenario individually to determine the objective values if the future were known exactly. We define the perfect foresight performance metric as $\boldsymbol{J_P}(s) = \boldsymbol{J_P}(\boldsymbol{\theta}^{*}, s)$, where the optimized parameters $\boldsymbol{\theta}^{*}$ are specific to the training set consisting of the single scenario s.

3.2.2. Hypervolume metric

We use a hypervolume metric to quantify the baseline regret of each scenario s in the ensemble. The hypervolume is defined as the volume in the objective space between the perfect foresight Pareto set $\boldsymbol{J_P}(s)$ and the baseline policy performance J_B, which is used as a reference point. While baseline regret is calculated in a five-dimensional objective space for this

application, the hypervolume concept is illustrated in two dimensions in the top row of Fig. 3. Solutions in the perfect foresight Pareto front $\boldsymbol{J_P}(s)$ are anticipated to dominate the baseline policy performance $J_B(s)$. The rare solutions for which this does not occur are not considered in the remainder of the calculations. In general, a larger hypervolume value indicates improvement over the baseline policy as well as a higher variety in alternatives among the Pareto set.

Initially disregarding the baseline solution, we normalize all objective values in $\boldsymbol{J_P}(s) \in [0, 1]$ to reduce scaling issues between the objectives. The baseline policy performance is normalized accordingly to $J_B \in (-\infty, 0]$ to allow for consistent comparison of baseline regret across scenarios. We then calculate the baseline regret $R(s)$ based on the hypervolume $h(\ldots)$ between the perfect foresight Pareto set and the baseline policy reference point, such that:

$$R(s) = h(\boldsymbol{J_P}(s), J_B(s)) \tag{13}$$

The baseline regret describes the performance of a perfect foresight optimization relative to the baseline for each climate scenario. Because the objective values are normalized, it provides an upper bound performance metric that can be directly compared across scenarios for a given policy.

3.3. Scenario clustering

Unsupervised clustering provides the basis for separating climate projections into training and test sets for the policy search. The clustering is based on three features: averaged annual streamflow, averaged peak annual snow-water equivalent, and the baseline regret metric described above. These features are calculated on the time horizon 2070–2100, which is chosen as the period of analysis due to its large deviation from historical hydrologic conditions and thus high regret (Fig. 2 in the Supplementary Material). This time period also contains much more variability in hydrologic properties than do earlier periods in the projected time horizon (Fig. 2).

The three features are calculated for each of the 97 scenarios in the ensemble and clustered using the K-means algorithm with $K = 3$, equal to the number of features. This allows for a minimally complex characterization of scenario properties based on cluster centroids. The resulting clusters are denoted as C_1, C_2, and C_3.

3.4. Training and test sets

We first split each cluster C_i randomly into roughly equal training and test subsets, S and S_t, respectively. Various combinations of these training and test subsets make up the overall training and test sets, S_i and S_{tj}, respectively (Table 1). While the goal of this division is to ensure that test information is never used in training, we acknowledge the possibility for interdependence among the ensemble of climate scenarios, for example using the same model or emissions scenario, or different models relying on the same components (Steinschneider et al., 2015b).

3.4.1. Training and testing

Policy search runs separately for each training set to identify the Pareto set of policies $\boldsymbol{\theta}^*_{S_i}$ corresponding to each training set of scenarios S_i:

$$\boldsymbol{\theta}^*_{S_i} = \text{argmax}_\theta J(\theta, S_i) \tag{14}$$

In order to increase the extent and continuity of the Pareto-optimal solutions, three trials of each optimization are run using varying random seeds. The use of baseline regret as a scenario property links the perfect foresight optimization to various combinations of training scenarios without explicitly using perfect foresight to inform the choice of all training scenarios.

We next re-evaluate the policies trained to set S_i over each scenario in the test set S_{tj}, resulting in a set of objectives $\boldsymbol{J}_{S_i}(s)$ for each test scenario:

$$\boldsymbol{J}_{S_i}(s) = J\left(\boldsymbol{\theta}^*_{S_i}, s\right) \quad \forall s \in S_{tj} \tag{15}$$

We consider only solutions that also dominate the baseline policy for all scenarios in the test set. This is achieved via a filtering step which identifies the solutions that will at a minimum outperform the status quo in all re-evaluations. We identify these particular policies and solutions as policy set $\boldsymbol{\theta}_{i,j}$ and solution set $\boldsymbol{J}_{i,j}$. This process results in a total of 28 pairwise combinations of training and test scenarios.

3.4.2. Set diversity

While the training and test sets are delineated via K-means clustering, the diversity of each set should also be considered for the analysis. This can help determine if policy performance across train-test set combinations is influenced by the scenario diversity in each set. Diversity is determined via Equations (16) and (17), adapted from Carlsen et al. (2016) and Eker and Kwakkel (2018). In Equation (16), D_S represents the diversity of set S. w, the weight assigned to the extent the mean distance, is 0.5 in this case. $d_{l,k}$ is the Euclidean distance based on the $m = 3$ features (FNF, SWE, baseline regret) for scenarios l and k. In Equation (17) $\overline{f_{m,l}}$ and $\overline{f_{m,k}}$ are the values for these features. To ensure equal weighting of all features, $\overline{f_{m,l}}$ and $\overline{f_{m,k}}$ are normalized from 0 to 1 across all scenarios.

$$D_S = (1-w)\min_{\forall l,k \in S}\{d_{l,k}\} + w\,\text{mean}_{\forall l,k \in S}\{d_{l,k}\} \tag{16}$$

$$d_{l,k} = \sqrt{\sum_m \left(\overline{f_{m,l}} - \overline{f_{m,k}}\right)^2} \tag{17}$$

3.5. Policy robustness

We would like to evaluate the robustness of policies trained to set S_i when re-evaluated over each scenario s in set S_{tj}, for all combinations of (i, j). Because the performance is multi-objective across a range of hydrologic scenarios, we use a hypervolume metric normalized by the baseline regret (Section 3.2) to represent the robustness of the policy set as a whole.

Table 1
Outline and properties of training and test sets.

Training Sets S_i	# of scenarios	Diversity	Test Sets S_{tj}	# of scenarios	Diversity
S_1 High-regret	13	0.241	S_{t1} High-regret	12	0.206
S_2 Low-regret wet	17	0.199	S_{t2} Low-regret wet	16	0.162
S_3 Low-regret dry	20	0.130	S_{t3} Low-regret dry	19	0.131
S_4 High-regret/low-regret wet ($S_1 \cup S_2$)	30	0.233	S_{t4} All test scenarios ($S_{t1} \cup S_{t2} \cup S_{t3}$)	47	0.202
S_5 High-regret/low-regret dry ($S_1 \cup S_3$)	33	0.230			
S_6 Low-regret wet/dry ($S_2 \cup S_3$)	37	0.197			
S_7 All training scenarios ($S_1 \cup S_2 \cup S_3$)	50	0.233			

3.5.1. Hypervolume robustness metric

Robustness is represented by a normalized hypervolume metric for each solution set. The hypervolume for a solution set of train-test set combination (i,j) applied to scenario s is defined as that between the baseline reference point $J_B(s)$ and solution set $\boldsymbol{J}_{i,j}$: $h(\boldsymbol{J}_{i,j}(s), J_B(s))$. This is normalized by the baseline regret $R(s)$, giving the hypervolume robustnesss metric $HR_{i,j}(s)$:

$$HR_{i,j}(s) = \frac{h\left(\boldsymbol{J}_{i,j}(s), J_B(s)\right)}{R(s)} \tag{18}$$

Thus the hypervolume robustness metric will always be a fraction of the baseline regret $R(s)$, ensuring that it can be appropriately compared across train-test combinations (see fourth row in Fig. 3). A higher normalized hypervolume metric denotes a more robust policy set, with a value of 1 equaling the performance of perfect foresight policies. This ensures that the robustness of a policy set is not measured only by its ability to improve system performance relative to the baseline, but also the extent to which the policies are able to reach the maximum attainable level of system performance.

For each train test-combination, we can obtain a set of hypervolume robustness metric values $HRT_{i,j}$, where:

$$HRT_{i,j} = \left\{HR_{i,j}(s_1), HR_{i,j}(s_2), \ldots, HR_{i,j}(s_n)\right\} \ \forall s \in S_{tj} \tag{19}$$

In general, larger hypervolume robustness metrics will indicate two properties of the objective outputs. The first is that as hypervolumes increase, the distance between the baseline policy $J_B(s)$ and policy performance set $\boldsymbol{J}_{i,j}(s)$ will increase, indicating higher performance improvements compared with the baseline policy. Additionally, a larger hypervolume indicates a higher diversity of solutions across the Pareto front.

3.5.2. Rank-sum tests

For each pair of train-test combinations with identical test sets, we perform a one-sided Mann-Whitney U test (Mann and Whitney, 1947) to determine if the hypervolume of a given training set exceeds that of a second training set when evaluated on the same test set. This test aims to determine if policies trained to a test set with particular properties are significantly more robust. With $p \leq 0.05$, we reject the null hypothesis and conclude that the distribution of hypervolume across test scenarios S_{tj} in sample $HRT_{1,j}$ is greater than that in sample $HRT_{2,j}$ with statistical significance.

3.6. Policy analysis

Finally, we analyze individual policies chosen from the most robust training sets by considering tradeoffs between the objective values. The decision variables of these policies are compared to the baseline policy to understand what combinations of adaptations to system operations could be promising under a range of future climates. We then compare the dynamics in terms of reservoir storage and water supply exports to those obtained by simulating the baseline policy on the same hydrologic inputs, and then relate key differences to the decision variables interpreted in the context of the system.

4. Results and discussion

4.1. Scenario clusters

Scenarios are divided into three clusters based on their streamflow, snowpack, and baseline regret, as shown in Fig. 4. Based on the cluster centroids, we define them as high-regret, low-regret wet, and low-regret dry. The high-regret scenarios contain a mix of streamflow and snowpack values distributed throughout their respective ranges, indicating that baseline regret does not solely depend on annual hydrologic properties. The clear separation between the high-regret and low-regret clusters suggests the possible utility of this metric in determining combinations of training scenarios in policy search experiments.

The low-regret scenarios occur in both wet and dry clusters. However, the ranges of streamflow and snowpack values overlap across these two clusters. (Fig. 4a and b). Specifically, this occurs in two cases: first, some wetter scenarios may also show high levels of snowpack decline due to severely warmer temperatures; second, there exist dry scenarios with relatively higher snowpack values than other low-flow scenarios due to less warming. This overlap, along with the much clearer separation between high-regret and low-regret clusters, supports the choice

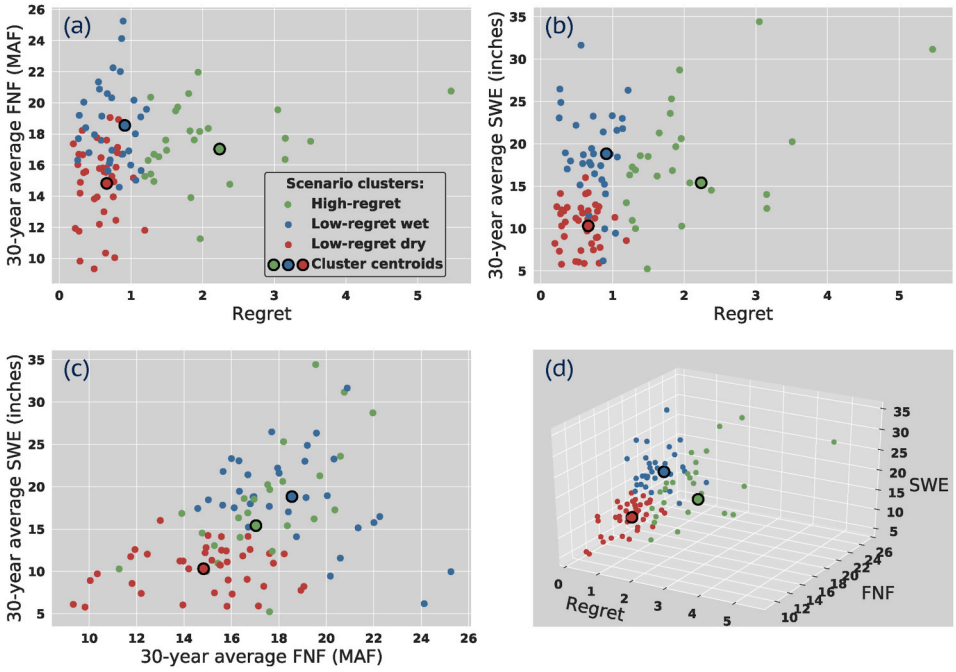

Fig. 4. Scenario clusters in (a,b,c) two-dimensional projections and (d) all three properties: full natural flow (streamflow, FNF), snow-water equivalent (snowpack, SWE), and baseline regret.

of $K = 3$ clusters to minimize complexity.

4.2. Training set robustness comparison

The three clusters, each split randomly into training and test subsets, are combined to create different train-test splits. The overall proportion of training to test sets out of all available scenarios is 50:47 (Table 1). From all possible combinations combined to create different train-test splits created in this process, a total of seven training sets and four test sets are chosen to demonstrate the training-testing process. These are described in Table 1.

These sets are used to determine the performance of policies optimized to each training set when re-evaluated in each test set, measured according to the hypervolume robustness metric. Fig. 5 shows the distributions of the resulting hypervolume metric for each train-test split, plotted as cumulative distributions.

Distributions shifted further right indicate higher robustness of the policy sets over the test set. While these distributions support the interpretation of the performance differences between policy sets trained on different scenarios, the rankings of policy sets must be shown to be statistically significant. These conclusions are made using the Mann-Whitney U test between each pair of train sets over each test set, with results shown in Fig. 6.

For the high-regret test set (Fig. 6a), the most robust policies are those optimized to the high-regret and all-scenario training sets, where the latter contains the former. Neither of these significantly outperforms the other. This finding is not surprising, as the policies trained to scenarios with similar properties demonstrate the best out-of-sample performance. However, this result does not always hold for the other test sets. For example, in the test set consisting of low-regret wet scenarios (Fig. 6b), the best-performing set of policies are those trained to a mix of high-regret and low-regret wet scenarios (S_4), which ranks higher than every other training set. The training sets containing dry scenarios and lacking wet scenarios (S_3 and S_5) perform worst for the high-regret test set. Set S_1, consisting of only high-regret scenarios, outperforms set S_6, which consists of wet and dry low-regret scenarios. This indicates that training to only high-regret scenarios may be more effective than training to low-regret scenarios regardless of the variability in scenarios' hydrologic properties. This result shows that adding high-regret scenarios to the training set—whether they are wet or dry—improves the robustness of the optimized policies when tested in out-of-sample wet scenarios. Additionally, including low-regret dry scenarios in training

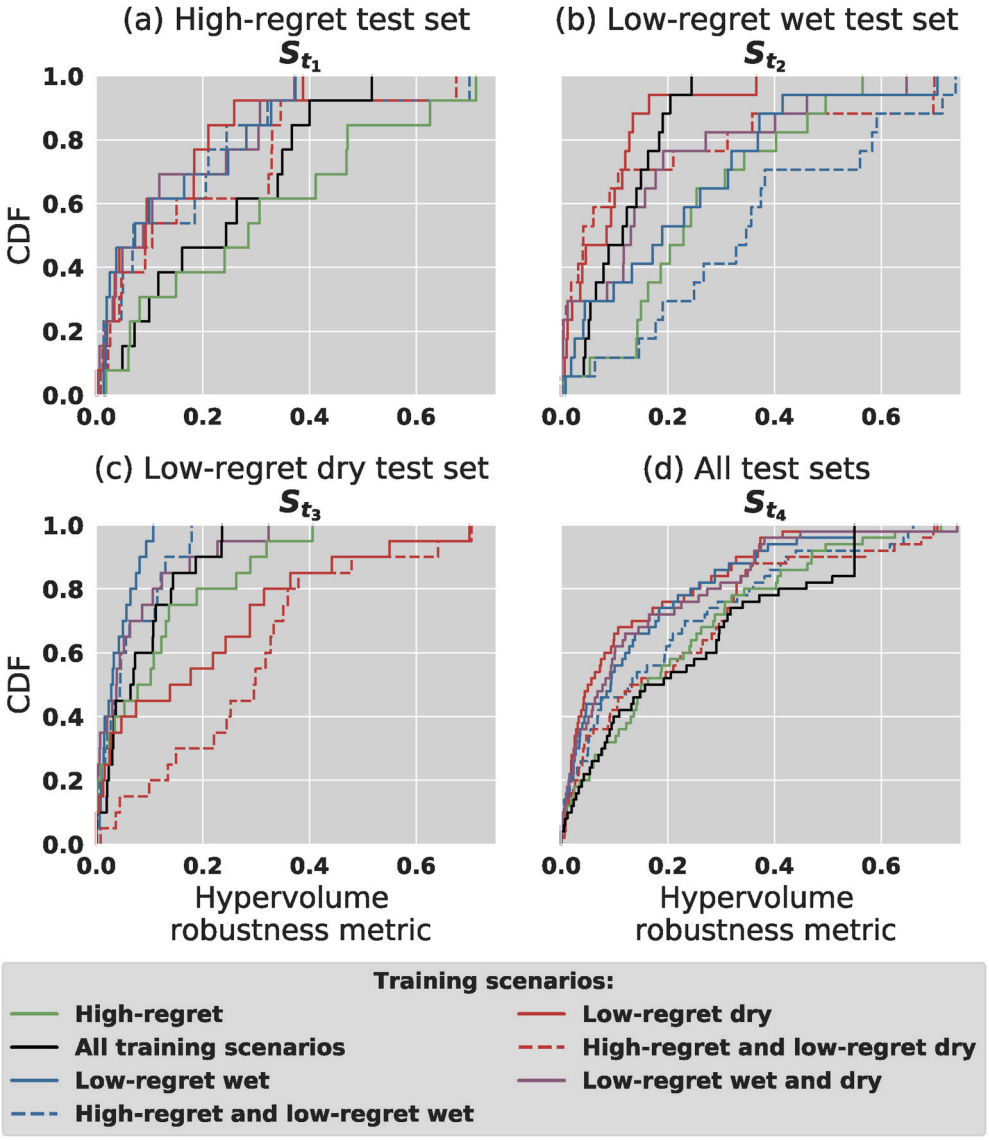

Fig. 5. Cumulative distributions of the hypervolume metric evaluated on each test set (A–D). Each CDF represents the distribution of performance over all scenarios in the test set for the Pareto front of policies trained to the scenarios identified by the line style.

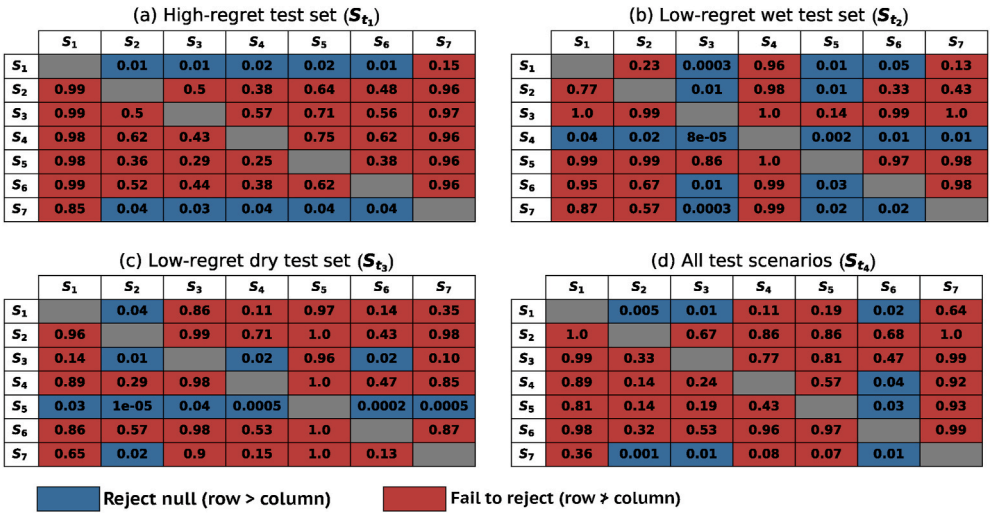

(a)	S_1	S_2	S_3	S_4	S_5	S_6	S_7
S_1		0.01	0.01	0.02	0.02	0.01	0.15
S_2	0.99		0.5	0.38	0.64	0.48	0.96
S_3	0.99	0.5		0.57	0.71	0.56	0.97
S_4	0.98	0.62	0.43		0.75	0.62	0.96
S_5	0.98	0.36	0.29	0.25		0.38	0.96
S_6	0.99	0.52	0.44	0.38	0.62		0.96
S_7	0.85	0.04	0.03	0.04	0.04	0.04	

(b)	S_1	S_2	S_3	S_4	S_5	S_6	S_7
S_1		0.23	0.0003	0.96	0.01	0.05	0.13
S_2	0.77		0.01	0.98	0.01	0.33	0.43
S_3	1.0	0.99		1.0	0.14	0.99	1.0
S_4	0.04	0.02	8e-05		0.002	0.01	0.01
S_5	0.99	0.99	0.86	1.0		0.97	0.98
S_6	0.95	0.67	0.01	0.99	0.03		0.98
S_7	0.87	0.57	0.0003	0.99	0.02	0.02	

(c)	S_1	S_2	S_3	S_4	S_5	S_6	S_7
S_1		0.04	0.86	0.11	0.97	0.14	0.35
S_2	0.96		0.99	0.71	1.0	0.43	0.98
S_3	0.14	0.01		0.02	0.96	0.02	0.10
S_4	0.89	0.29	0.98		1.0	0.47	0.85
S_5	0.03	1e-05	0.04	0.0005		0.0002	0.0005
S_6	0.86	0.57	0.98	0.53	1.0		0.87
S_7	0.65	0.02	0.9	0.15	1.0	0.13	

(d)	S_1	S_2	S_3	S_4	S_5	S_6	S_7
S_1		0.005	0.01	0.11	0.19	0.02	0.64
S_2	1.0		0.67	0.86	0.86	0.68	1.0
S_3	0.99	0.33		0.77	0.81	0.47	0.99
S_4	0.89	0.14	0.24		0.57	0.04	0.92
S_5	0.81	0.14	0.19	0.43		0.03	0.93
S_6	0.98	0.32	0.53	0.96	0.97		0.99
S_7	0.36	0.001	0.01	0.08	0.07	0.01	

Fig. 6. Results of pairwise Mann-Whitney U rank-sum tests. Each test has the null hypothesis that the hypervolume metric associated with the training set in each row is less than or equal to that associated with the column. The subplots correspond to the four test sets. A row where the null hypothesis is rejected (blue) for each cell denotes a robust policy set that ranks highest for the particular test set.

sets for policies tested on low-regret wet scenarios degrades policy performance. Similar results are shown for the low-regret dry test set (Fig. 6c), where the highest ranking training set is again not only the dry scenarios (S_3), but also a mix of high-regret and low-regret dry (S_5) scenarios. In addition, training sets including low-regret wet scenarios (S_2 and S_4) have the lowest ranking when their corresponding policies are simulated over the low-regret dry test set. Lastly, the high-regret training set continues to outperform the low-regret wet/dry training set for S_{t4}, further highlighting the good training value of high-regret scenarios.

For the final test set S_{t4}, which includes all testing scenarios (Fig. 6d), the majority of Mann-Whitney U tests fail to reject the null hypothesis. However, results indicate that the high-regret training set (S_1) outperforms the low-regret training sets (S_2, S_3, and S_6), as does the all-scenarios training set (S_7). Especially notable is that the high-regret training set S_1 outperforms the combined wet-dry low regret training set S_6 when testing to all scenarios. Since both of these training sets have wide ranges for the hydrologic properties, this further highlights the benefit of high-regret training scenarios over low-regret scenarios.

The diversity of the training sets can be analyzed in tandem with these results. The high-regret (S_1) and all training scenario (S_7) sets are the most diverse (Table 1). Based on our specific quantification of set diversity, this is an artifact specifically of the high-regret values, which contain more outliers and a more skewed distribution across all scenarios (Fig. 4a and b). The high mean distances that occur from this cause the diversity values to be larger whenever the high-regret scenarios are included in a set. This leads to the fact that the three sets which do not contain the high-regret scenarios (S_2, S_3, and S_6) are the least diverse of the sets. It could then be concluded that the larger diversity of the high-regret set influences its good performance. However, since this value is skewed by just a few outliers, it should not be considered the only reason for the effective training value of the high-regret sets.

Training sets S_1 (high-regret only) and S_6 (low-regret wet/dry) have similar ranges across both hydrologic properties. The lower diversity of set S_6 is influenced by its small range in baseline regret values, as well as the fact that is has several scenarios in close proximity in terms of hydrologic properties (Fig. 4c), leading to a skewed minimum distance value in the diversity calculation. Set S_6 has almost three times the number of scenarios as set S_1, which contributes to its low diversity calculation. Several close-proximity scenarios could be omitted to make set S_6 more diverse. This would not improve the performance as the set would lose valuable training data and potential for overfitting would increase. Therefore, the high diversity of set S_1 is not the only factor controlling the set's good performance. Its high baseline regret values will enable the policy search to find solutions more robust to vulnerable conditions. Additionally, there may be many other scenario properties that are not examined in this study which contribute to set performance and scenario training value. These include hydroclimatic properties such as temperature rise, flood frequencies, flow timing, precipitation, drought patterns, soil moisture, and evapotranspiration.

Because the high-regret training set performs no worse than training to all scenarios, the strategy of designing a training set around scenarios with high baseline regret may serve to reduce the computational cost of policy search for large-ensemble cases, and/or to reserve more scenarios for testing. To support this point, Table 2 compares the computational cost for different aspects of policy training in this study. Training to scenarios with high baseline regret (which includes the perfect foresight optimizations) required 9733 computing hours, roughly three times less than training to all scenarios. Training to scenarios with high baseline regret improves the efficiency of policy search without sacrificing robustness relative to the case of training to all scenarios. This denotes the benefit of analyzing the hydrology and baseline regret of scenarios before a train/test split is determined.

Thus, it is also possible to determine the conditions under which a high baseline regret set will give computational benefits by generalizing the requirements outlined in Table 2. This condition is described as:

$$f_P\rho_P + \eta f_r\rho < f_A\rho_A \tag{20}$$

Where η represents fraction of overall scenarios which are in the high-regret set, f_P, f_R, and f_A denote the number of function evaluations, and ρ_P, ρ_R, and ρ_A denote the number of random seeds for each of the perfect foresight, high-regret only, and all training scenario sets, respectively. This generalization can potentially be applied to other planning problems in which the baseline regret is determined a priori, and where there is a choice about how many high-regret solutions to include in the training set.

This analysis has important implications for the generalizability of this approach. Several variables may be degrees of freedom, for instance numbers of random seeds ρ and function evaluations f necessary for convergence to diverse and near-optimal Pareto-solutions sets will vary across models. The fraction of high regret scenarios η may differ based on the number of clusters chosen. In some instances, if the level of baseline regret is not a significant source of variation among scenarios, it may not provide a way of separating different training sets using a

Table 2
Description of computing requirements for several optimizations included in this study. Note that less NFEs are required for a perfect foresight, as these optimizations are quicker to converge. Times per function evaluation and total computing hours are specific to the UC Davis HPC1 computing cluster.

	Max NFE per optimization	Scenarios per optimization(n)	Time per function evaluation	# of random seeds (ρ)	Total computing hours
Perfect foresight (P)	10,000	1	12s	97 (individual scenario trials)	3,233 h
High-regret only (R)	50,000	13	156s	3 (random seed trials)	6,500 h
All training scenarios (A)	50,000	49	588s	3 (random seed trials)	24,500 h

clustering approach. Furthermore, differences in performance among training sets may be due to confounding factors not reflected in the abstracted scenario properties, especially for hydrologic timeseries which can be summarized in a number of different ways. However, the proposed clustering and train/test methodology is still generalizable across environmental planning applications to pinpoint the most important scenario properties for policy training and out-of-sample performance, therefore discovering conditions for computational benefits.

Furthermore, results presented in Figs. 5 and 6 must be interpreted in light of the fact that the future climate trajectory is uncertain. It is likely that more information about future hydrology will be collected over time, and this process could complement policy search methods in the context of dynamic planning (e.g., Hui et al., 2018; Fletcher et al., 2019). Therefore, in this study the methodology aims to identify a training strategy that leads to robust outcomes to both uncertain and clustered future climate, measured according to multi-objective performance bounded by the baseline policy and perfect foresight cases. We find that training to scenarios with high baseline regret is competitive with training to all scenarios across a range of future climates, and often leads to the best out-of-sample performance. This is likely due to higher inter-annual variability in these scenarios. Based on a higher diversity of extreme events across individual scenarios and potential poor baseline performance in the high-regret cluster, solutions will give both a wider variety of tradeoffs in objectives and improvements relative to baseline policy performance. These findings extend to both wet and dry futures, where the inclusion of high-regret scenarios in the training set outperforms using exclusively either wet or dry training scenarios. This

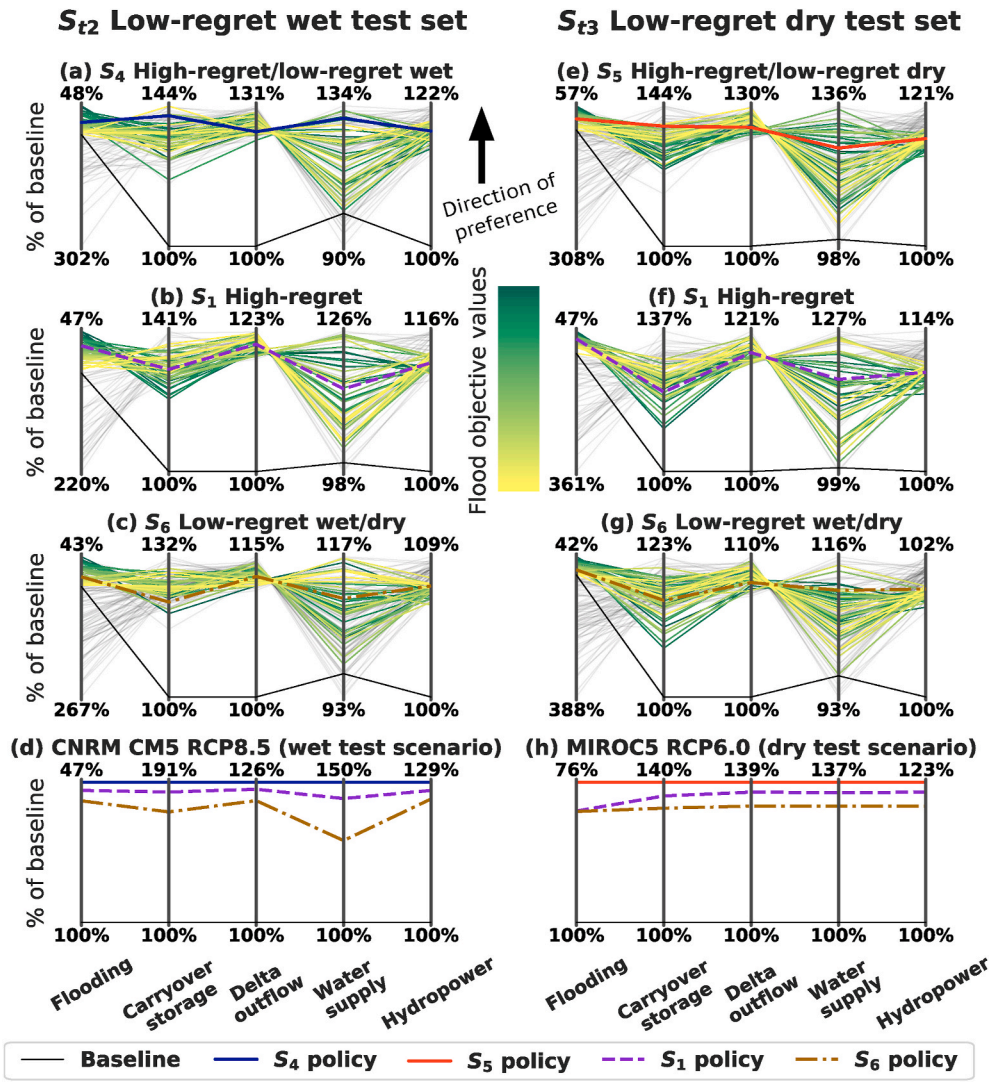

Fig. 7. Parallel axis plots displaying various results of train-test combinations. (a,b,c) policies trained on sets S_4, S_1, and S_6, respectively, tested on set S_{t2}. (e,f,g) policies trained on sets S_5, S_1, and S_6, respectively, tested on set S_{t3}. Solutions highlighted by the yellow-green gradient represent solutions for which the expected value of solutions across the testing scenarios dominate the baseline policy solution. This gradient represents the solutions' ranks for the flood objective in their particular Pareto set. Grey solutions dominate the baseline in terms of expected value, but do not dominate the baseline for each individual scenario in the testing set. The individual highlighted solutions denote four different compromise policies (S_4, S_5, S_1, and S_6 policies) that are analyzed in detail. Subplots (d) and (h) show robust performance of the four compromise policies over an individual scenario in the corresponding test set. These particular scenarios are also highlighted in Fig. 2.

result links to the importance of evaluating perfect foresight policies in individual scenarios when designing the training set to establish an upper bound for system performance.

4.3. Policy analysis

The final step of the analysis is to determine what specific adaptations are implemented by the robust policies. This analysis focuses on six specific train-test splits, chosen based on their high-ranking performance: (1) policies trained on set S_4 and tested on set S_{t2}, (2) policies trained on the set S_5 tested on set S_{t3}, (3,4) policies trained on set S_1 tested on sets S_{t2} and S_{t3}, and (5,6) policies trained on set S_6 tested on sets S_{t2} and S_{t3}. The average performance measures across all scenarios for these sets are shown by the highlighted solutions on the parallel-axis plots in Fig. 7. While the expected value of all highlighted solutions dominates the baseline policy, there are still several significant tradeoffs between the objectives, indicated by their nonlinear correlations (see Section 3 in Supplementary Material). For the S_4/S_{t2} train-test combination, these include statistically significant tradeoffs between hydropower and flooding ($\rho = 0.53$) and water supply and Delta outflow ($\rho = 0.97$). The same tradeoffs exist in the dry test scenarios, which also exhibit tradeoffs between carryover storage and flooding ($\rho = 0.49$). In general, these relationships reflect the fact that higher storage levels benefit several of the proposed objectives, although they can be detrimental to the flooding objective, which is to be minimized. While this high water elevation benefits the hydropower and carryover storage objectives, it can induce larger releases if large storms occur later in the spring.

We have shown that training set S_4 (high-regret/low-regret wet) will yield the best performing policies for the low-regret wet test set S_{t2} (Fig. 6b). Likewise, training set S_5 (high-regret/low-regret dry) will yield the best performing policies for the low-regret wet test set S_{t3} (Fig. 6b). This is reflected in Fig. 7(a,b,c), (e,f,g) where the highlighted Pareto solutions for S_4 and S_5 are shifted higher than S_1 and S_6 over their particular test sets, as shown by the higher maximum percent of baseline values (for the flood objective lower minimum) in Fig. 7a, e. These ranges in Fig. 7(a,b), (e,f) also reflect the better performance of the high-regret training set S_1 over the low-regret wet/dry training set S_6 for both test sets.

We next examine the four compromise policies that balance the tradeoffs in performance measures, denoted as the S_4, S_5, S_1, and S_6 policies in Fig. 7. The S_4 and S_5 policies, coming from the most robust training sets for the respective test sets, also give the best performance on the individual scenarios (Fig. 7d,h). The alternatives that these policies employ are shown in Fig. 8 along with a comparison to the decision variables of the baseline policy. Each column in the tables represents the decision variable which occurs for that specific water year type.

In the S_4 policy, Shasta and Folsom reservoirs have higher maximum allowable curtailments $c^r_{max,wyt}$ than in the baseline policy. These higher maximum curtailment levels will allow for increased hedging of releases. The curtailments for Oroville reservoir are higher in wet, above and below normal years, but lower in dry and critical years. All three reservoirs also have a flood pool shift of at least 10 days forward in the water year for the S_4 policy. In wet and above normal years, Shasta and Oroville use low $\overline{Z}^r_{wyt}$ values, indicating a very conservative forecast with a high exceedance level. In drier water year types, the $\overline{Z}^r_{wyt}$ values are generally close to or greater than the baseline exceedance levels. For Folsom reservoir, these values vary much more across water year types. The differences between operational adaptations at each reservoir highlight the complexity of managing the multi-reservoir system, and the potential to design adaptations for system-wide benefit.

Fig. 9 shows the system dynamics of the baseline policy compared to the compromise policies in a time series over one scenario from each corresponding test set: an RCP 8.5 scenario (CNRM-CM5) for low-regret wet, and an RCP 6.0 (MIROC5) scenario for low-regret dry. Under the baseline policy, reservoir storage levels are vulnerable to snowmelt loss regardless of water year type, evidenced by low storage levels in the

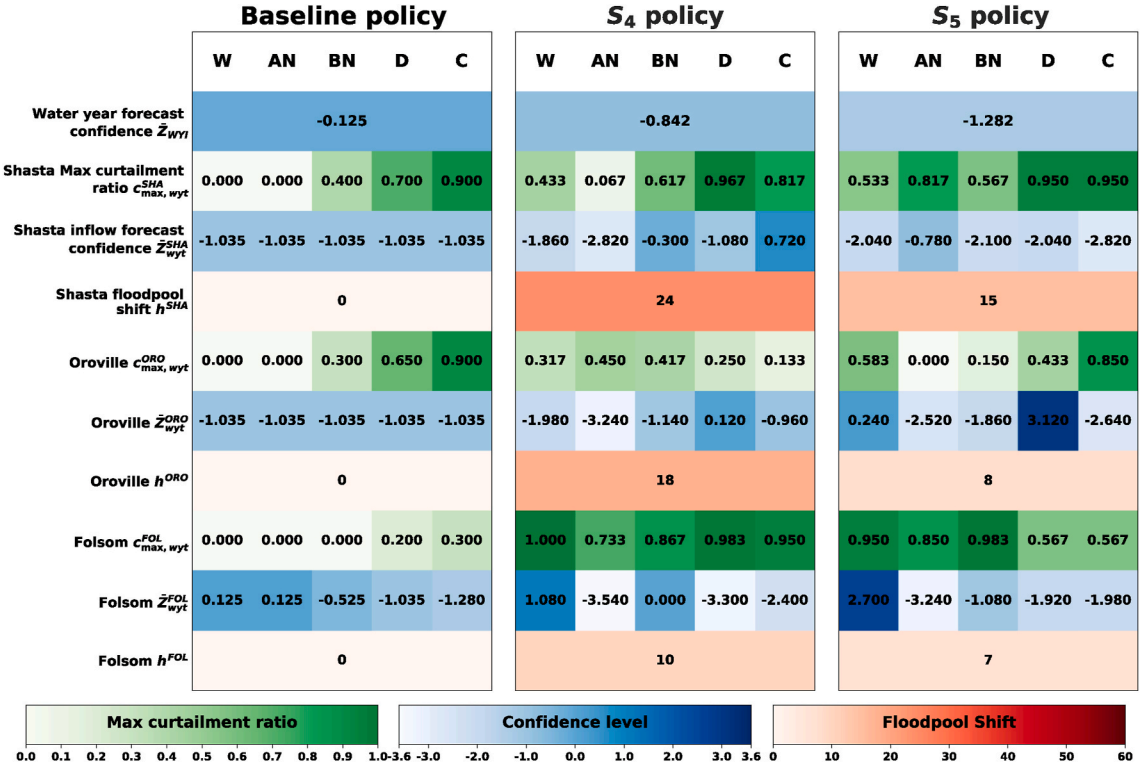

Fig. 8. Policy tables showing decision variables for: the baseline policy, the S_4 policy (high-regret/low-regret wet set), and the S_5 policy (high-regret/low-regret dry set). The columns denote water year type classifications associated with each decision variable, corresponding to wet, above normal, below normal, dry, and critical. Policy tables for the S_1 and S_6 policies can be found in Section 4 of the Supplementary Material.

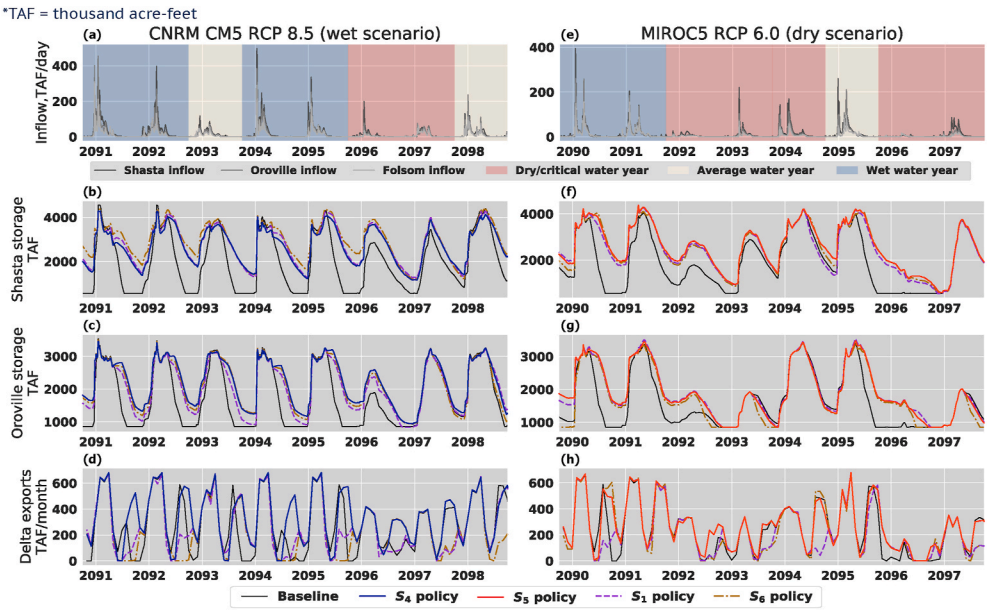

Fig. 9. Time series of system states and flows for baseline and adaptation policies in individual test scenarios. The left and right columns show results from the individual wet and dry scenarios, respectively. (a,e) Daily inflows to each reservoir with water year types highlighted; (b,c,f,g) Reservoir storage; (d,h) monthly Delta exports, primarily for agricultural and municipal water supply.

irrigation season even in wetter years. The S_4 policy mitigates this vulnerability via an intra-annual hedging, resulting in higher reservoir storage during the early irrigation season (May–June); The S_5 policy functions similarly. For both policies, this intra-annual hedging dynamic is supported by the adapted snowpack-to-streamflow forecasts, where underpredictions will cause some release curtailments to conserve for potential low inflows later in the season. However, curtailments can be partially avoided with higher carryover storage due to the flood pool shift. This seasonal shift is also reflected in the Delta exports (Fig. 9d,h), which maximize total volume by shifting throughout the year. The remaining S_6 and S_1 policies also exhibit the intra-annual hedging strategy (see Section 4 in Supplementary Material for these policies' decision variables). However, given that reservoir storage becomes higher in the flood season and carryover storage drops lower when these policies are deployed, they are slightly less effective (Fig. 9b,c,f,g). Additionally, they often will have periods of low Delta exports (Fig. 9d, h). This highlights that a policy from the best performing training set for a particular test set may be more likely to give better performance for scenarios in that test set.

There are two major differences between the S_4 and S_5 policies stemming from the hydrologic properties of their respective training scenarios. The first is that the S_5 policy tends to curtail releases more during dry and critical years, reflected in its conservative forecasts and high maximum curtailment allowances. The S_4 policy hedges less during dry and critical years, and instead relies on larger storage brought about by intra-annual hedging. This is further driven by the low maximum curtailment allowances for Oroville during these water year types. The second difference is that the S_4 policy tends to hold less storage during the flood season than the baseline policy, while the S_5 policy does not. For the S_4 policy, this makes curtailment less necessary later in the mid- to late summer, and reduces flood vulnerabilities. The fact that policies exist that can improve upon both of these objectives via the same policy parameters is the main reason why flood control and carryover storage do not have a significant tradeoff in the wet test set. In summary, analysis of these two compromise policies shows how training to scenarios with high baseline regret can yield policies with improved performance on out-of-sample hydrology to balance conflicting objectives.

5. Conclusions

This study advances the design and testing of robust control policies as an adaptation to uncertainty in environmental planning problems, contributing an experimental design to better understand the influence of the forcing scenario properties and baseline regret of training scenarios on the robustness of resulting policies. We demonstrate this approach for the northern California reservoir system to determine how transient downscaled climate scenarios impact tradeoffs between water supply, flood control, environmental flows, and hydropower generation. Results indicate that policies trained to scenario sets with high baseline regret tend to outperform those generated with other training sets in both wetter and drier futures. Additionally, the policies adapted under these conditions develop an intra-annual hedging strategy to mitigate the effects of snowpack decline under rising temperatures. The approach highlights the general importance of considering the specific properties of training scenarios in the design of robust control policies.

Beyond the pairwise comparison of train-test splits, this analysis also highlights the general difficulty of maintaining out-of-sample performance for reservoir control policies. This is driven primarily by extreme events that occur infrequently by definition and which may be the result of natural variability rather than anthropogenic change, creating a risk of overfitting to the training set. The baseline regret, based on perfect foresight optimization, provides a measure of regret to place this performance degradation in context. Unlike the traditional minimax regret strategy, where the alternative that minimizes the maximum regret across all scenarios is chosen, our approach uses a regret metric to choose training scenarios rather than optimal alternatives. We show that optimal policies benefit from training to sets of scenarios with a high regret for the baseline solution. Our methodology also provides a way to group ensembles of scenarios using an unsupervised learning approach, along with other hydrologic properties including streamflow and snowpack, to create an experiment which maps the relationship between training and test scenarios to the outcome of policy robustness considering both the performance and diversity of solutions. The latter is particularly important given the concern with reversible adaptations to operations which can be changed over time (Herman et al., 2020).

While this study considers uncertainty in hydrology due to climate change across downscaled model projections, it could further test the

robustness of the resulting policies against more realizations of sampling variability from a synthetic generator, or supplement the training set with the same. Increasing the number of scenario realizations would allow for additional hydrologic variables to be included in clustering, such as changes in flood and drought frequencies and intra-annual streamflow shifts. Additionally, policy training might be improved with a more flexible policy structure beyond parameterizing the existing system, such as a neural network—though this may also increase the potential for overfitting due to increased degrees of freedom. Policy training can also be coupled with infrastructure design (e.g. Bertoni et al., 2020), which in many regions will be required to cope with the more extreme projections of hydrologic change. Lastly, while our approach is demonstrated with an exampled from the water resources management field, it can generalize to any environmental, natural resources, or infrastructure planning problem which includes a no action case, an optimization component, and a forcing scenario ensemble. Future work should explore the impacts of these additional experimental components in combination with the analysis of the training scenarios properties presented here to further improve robust policy search under uncertainty.

Software availability

Code for Operations of Reservoir in California (ORCA), a python simulation model, is available at https://github.com/jscohen4/orca Simulation code, data analysis and figure scripts for this manuscripts are available at: https://github.com/jscohen4/orca/tree/cohen-2021-properties-training-scenarios.

Processed CMIP5 climate projection data files used in this study are available at https://github.com/jscohen4/orca_cmip5_inputs.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially supported by the U.S. National Science Foundation grant CBET-1803589 and INFEWS grant CNS-1639268. Any opinions, findings, and conclusions are those of the authors and do not necessarily reflect the views or policies of the NSF. We further acknowledge the World Climate Research Program's Working Group on Coupled Modeling and the climate modeling groups listed in the supplement of this paper for producing and making available their model output.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.envsoft.2021.105047.

References

Anderson, J., Chung, F., Anderson, M., Brekke, L., Easton, D., Ejeta, M., Peterson, R., Snyder, R., 2008. Progress on incorporating climate change into management of California's water resources. Climatic Change 87 (1), 91–108.

Asadieh, B., Krakauer, N.Y., 2017. Global change in streamflow extremes under climate change over the 21st century. Hydrol. Earth Syst. Sci. 21 (11), 5863.

Beh, E.H., Maier, H.R., Dandy, G.C., 2015. Adaptive, multiobjective optimal sequencing approach for urban water supply augmentation under deep uncertainty. Water Resour. Res. 51 (3), 1529–1551.

Bertoni, F., Giuliani, M., Castelletti, A., 2020. Integrated design of dam size and operations via reinforcement learning. J. Water Resour. Plann. Manag. 146 (4), 04020,010.

Brekke, L., Wood, A., Pruitt, T., 2014. Downscaled Cmip3 and Cmip5 Hydrology Projections: Release of Hydrology Projections, Comparison with Preceding Information, and Summary of User Needs. National Center for Atmospheric Research.

Brekke, L.D., Maurer, E.P., Anderson, J.D., Dettinger, M.D., Townsley, E.S., Harrison, A., Pruitt, T., 2009. Assessing reservoir operations risk under climate change. Water Resour. Res. 45 (4).

Brodeur, Z.P., Herman, J.D., Steinschneider, S., 2020. Bootstrap aggregation and cross-validation methods to reduce overfitting in reservoir control policy search. Water Resour. Res. 56 (8), e2020WR027, 184.

Brown, C., Ghile, Y., Laverty, M., Li, K., 2012. Decision scaling: linking bottom-up vulnerability analysis with climate projections in the water sector. Water Resour. Res. 48 (9).

Carlsen, H., Lempert, R., Wikman-Svahn, P., Schweizer, V., 2016. Choosing small sets of policy-relevant scenarios by combining vulnerability and diversity approaches. Environ. Model. Software 84, 155–164.

Cayan, D.R., Kammerdiener, S.A., Dettinger, M.D., Caprio, J.M., Peterson, D.H., 2001. Changes in the onset of spring in the western United States. Bull. Am. Meteorol. Soc. 82 (3), 399–416.

CDEC, 2018. California Data Exchange Center. California Department of Water Resources.

Cohen, J.S., Zeff, H.B., Herman, J.D., 2020. Adaptation of multiobjective reservoir operations to snowpack decline in the western United States. J. Water Resour. Plann. Manag. 146 (12), 04020,091.

Culley, S., Noble, S., Yates, A., Timbs, M., Westra, S., Maier, H., Giuliani, M., Castelletti, A., 2016. A bottom-up approach to identifying the maximum operational adaptive capacity of water resource systems to a changing climate. Water Resour. Res. 52 (9), 6751–6768.

Deb, K., Jain, H., 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. IEEE Trans. Evol. Comput. 18 (4), 577–601.

Dessai, S., Hulme, M., 2004. Does climate adaptation policy need probabilities? Clim. Pol. 4 (2), 107–128.

Dottori, F., Szewczyk, W., Ciscar, J.-C., Zhao, F., Alfieri, L., Hirabayashi, Y., Bianchi, A., Mongelli, I., Frieler, K., Betts, R.A., et al., 2018. Increased human and economic losses from river flooding with anthropogenic warming. Nat. Clim. Change 8 (9), 781–786.

Draper, A.J., Lund, J.R., 2004. Optimal hedging and carryover storage value. J. Water Resour. Plann. Manag. 130 (1), 83–87.

Eker, S., Kwakkel, J.H., 2018. Including robustness considerations in the search phase of many-objective robust decision making. Environ. Model. Software 105, 201–216.

Fletcher, S., Lickley, M., Strzepek, K., 2019. Learning about climate change uncertainty enables flexible water infrastructure planning. Nat. Commun. 10 (1), 1782.

Fletcher, S.M., Miotti, M., Swaminathan, J., Klemun, M.M., Strzepek, K., Siddiqi, A., 2017. Water supply infrastructure planning: decision-making framework to classify multiple uncertainties and evaluate flexible design. J. Water Resour. Plann. Manag. 143 (10), 04017,061.

Giudici, F., Castelletti, A., Giuliani, M., Maier, H.R., 2020. An Active Learning Approach for Identifying the Smallest Subset of Informative Scenarios for Robust Planning under Deep Uncertainty. Environmental Modelling & Software, p. 104681.

Giuliani, M., Castelletti, A., 2016. Is robustness really robust? how different definitions of robustness impact decision-making under climate change. Climatic Change 135 (3–4), 409–424.

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 Resour. Res. 50 (4), 3355–3377.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., Reed, P.M., 2015a. Curses, tradeoffs, and scalable management: advancing evolutionary multiobjective direct policy search to improve water reservoir operations. J. Water Resour. Plann. Manag. 142 (2), 04015,050.

Giuliani, M., Pianosi, F., Castelletti, A., 2015b. Making the most of data: an information selection and assessment framework to improve water systems operations. Water Resour. Res. 51 (11), 9073–9093.

Giuliani, M., Quinn, J.D., Herman, J.D., Castelletti, A., Reed, P.M., 2017. Scalable multiobjective control for large-scale water resources systems under uncertainty. IEEE Trans. Contr. Syst. Technol. 26 (4), 1492–1499.

Gleick, P.H., 2002. Water management: soft water paths. Nature 418 (6896), 373.

Groves, D.G., Fischbach, J.R., Bloom, E., Knopman, D., Keefe, R., 2013. Adapting to a Changing Colorado River: Making Future Water Deliveries More Reliable through Robust Management Strategies. RAND corporation.

Haasnoot, M., Kwakkel, J.H., Walker, W.E., ter Maat, J., 2013. Dynamic adaptive policy pathways: a method for crafting robust decisions for a deeply uncertain world. Global Environ. Change 23 (2), 485–498.

Hadka, D., 2015. Platypus-multiobjective Optimization in python.

Hamarat, C., Kwakkel, J.H., Pruyt, E., Loonen, E.T., 2014. An exploratory approach for adaptive policymaking by using multi-objective robust optimization. Simulat. Model. Pract. Theor. 46, 25–39.

Herman, J.D., Giuliani, M., 2018. Policy tree optimization for threshold-based water resources management over multiple timescales. Environ. Model. Software 99, 39–51.

Herman, J.D., Reed, P.M., Zeff, H.B., Characklis, G.W., 2015. How should robustness be defined for water systems planning under change? J. Water Resour. Plann. Manag. 141 (10), 04015,012.

Herman, J.D., Quinn, J.D., Steinschneider, S., Giuliani, M., Fletcher, S., 2020. Climate Adaptation as a Control Problem: Review and Perspectives on Dynamic Water Resources Planning under Uncertainty. Water Resources Research, e24389.

Huang, X., Hall, A.D., Berg, N., 2018. Anthropogenic warming impacts on today's sierra Nevada snowpack and flood risk. Geophys. Res. Lett. 45 (12), 6215–6222.

Hui, R., Herman, J., Lund, J., Madani, K., 2018. Adaptive water infrastructure planning for nonstationary hydrology. Adv. Water Resour. 118, 83–94.
Kapnick, S., Hall, A., 2010. Observed climate–snowpack relationships in California and their implications for the future. J. Clim. 23 (13), 3446–3456.
Karamouz, M., Goharian, E., Nazif, S., 2013. Reliability assessment of the water supply systems under uncertain future extreme climate conditions. Earth Interact. 17 (20), 1–27.
Kasprzyk, J.R., Nataraj, S., Reed, P.M., Lempert, R.J., 2013. Many objective robust decision making for complex environmental systems undergoing change. Environ. Model. Software 42, 55–71.
Klos, P.Z., Link, T.E., Abatzoglou, J.T., 2014. Extent of the rain-snow transition zone in the western us under historic and projected climate. Geophys. Res. Lett. 41 (13), 4560–4568.
Knowles, N., Dettinger, M.D., Cayan, D.R., 2006. Trends in snowfall versus rainfall in the western United States. J. Clim. 19 (18), 4545–4559.
Knowles, N., Cronkite-Ratcliff, C., Pierce, D., Cayan, D., 2018. Responses of unimpaired flows, storage, and managed flows to scenarios of climate change in the san francisco bay-delta watershed. Water Resour. Res. 54 (10), 7631–7650.
Koutsoyiannis, D., Economou, A., 2003. Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems. Water Resour. Res. 39 (6).
Kwakkel, J.H., Haasnoot, M., Walker, W.E., 2015. Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world. Climatic Change 132 (3), 373–386.
Kwakkel, J.H., Haasnoot, M., Walker, W.E., 2016. Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. Environ. Model. Software 86, 168–183.
Lempert, R.J., Collins, M.T., 2007. Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. Risk Anal.: Int. J. 27 (4), 1009–1026.
Liang, X., Lettenmaier, D.P., Wood, E.F., Burges, S.J., 1994. A simple hydrologically based model of land surface water and energy fluxes for general circulation models. J. Geophys. Res.: Atmospheres 99 (D7), 14,415–14,428.
Maier, H.R., Guillaume, J.H., van Delden, H., Riddell, G.A., Haasnoot, M., Kwakkel, J.H., 2016. An uncertain future, deep uncertainty, scenarios, robustness and adaptation: how do they fit together? Environ. Model. Software 81, 154–164.
Mann, H.B., Whitney, D.R., 1947. On a test of whether one of two random variables is stochastically larger than the other. Ann. Math. Stat. 50–60.
Mateus, M.C., Tullos, D., 2017. Reliability, sensitivity, and vulnerability of reservoir operations under climate change. J. Water Resour. Plann. Manag. 143 (4), 04016,085.
McCabe, G.J., Clark, M.P., Hay, L.E., 2007. Rain-on-snow events in the western United States. Bull. Am. Meteorol. Soc. 88 (3), 319–328.
Medellín-Azuara, J., Harou, J.J., Olivares, M.A., Madani, K., Lund, J.R., Howitt, R.E., Tanaka, S.K., Jenkins, M.W., Zhu, T., 2008. Adaptability and adaptations of California's water supply system to dry climate warming. Climatic Change 87 (1), 75–90.
Nayak, M.A., Herman, J.D., Steinschneider, S., 2018. Balancing flood risk and water supply in California: policy search integrating short-term forecast ensembles with conjunctive use. Water Resour. Res. 54 (10), 7557–7576.
Prudhomme, C., Wilby, R.L., Crooks, S., Kay, A.L., Reynard, N.S., 2010. Scenario-neutral approach to climate change impact studies: application to flood risk. J. Hydrol. 390 (3–4), 198–209.
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.
Quinn, J.D., Reed, P.M., Keller, K., 2017. Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. Environ. Model. Software 92, 125–141.
Quinn, J.D., Reed, P.M., Giuliani, M., Castelletti, A., Oyler, J.W., Nicholas, R.E., 2018. Exploring how changing monsoonal dynamics and human pressures challenge multireservoir management for flood protection, hydropower production, and agricultural water supply. Water Resour. Res. 54 (7), 4638–4662.
Ray, P., Wi, S., Schwarz, A., Correa, M., He, M., Brown, C., 2020. Vulnerability and risk: climate change and water supply from California's central valley water system. Climatic Change 1–23.
Reclamation, 2013. Downscaled Cmip3 and Cmip5 Climate and Hydrology Projections: Release of Downscaled Cmip5 Climate Projections, Comparison with Preceding Information, and Summary of User Needs.
Rhoades, A.M., Jones, A.D., Ullrich, P.A., 2018. Assessing mountains as natural reservoirs with a multimetric framework. Earth's Future 6 (9), 1221–1241.
Russell, S., Norvig, P., 2002. Artificial Intelligence: a Modern Approach.
Salazar, J.Z., Reed, P.M., Herman, J.D., Giuliani, M., Castelletti, A., 2016. A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. Adv. Water Resour. 92, 172–185.
Savage, L.J., 1951. The theory of statistical decision. J. Am. Stat. Assoc. 46 (253), 55–67.
Steinschneider, S., McCrary, R., Wi, S., Mulligan, K., Mearns, L.O., Brown, C., 2015a. Expanded decision-scaling framework to select robust long-term water-system plans under hydroclimatic uncertainties. J. Water Resour. Plann. Manag. 141 (11), 04015,023.
Steinschneider, S., McCrary, R., Mearns, L.O., Brown, C., 2015b. The effects of climate model similarity on probabilistic climate projections and the implications for local, risk-based adaptation planning. Geophys. Res. Lett. 42 (12), 5014–5044.
Steinschneider, S., Ray, P., Rahat, S.H., Kucharski, J., 2019. A weather-regime-based stochastic weather generator for climate vulnerability assessments of water systems in the western United States. Water Resour. Res. 55 (8), 6923–6945.
Surfleet, C.G., Tullos, D., 2013. Variability in effect of climate change on rain-on-snow peak flow events in a temperate climate. J. Hydrol. 479, 24–34.
Trindade, B., Reed, P., Herman, J., Zeff, H., Characklis, G., 2017. Reducing regional drought vulnerabilities and multi-city robustness conflicts using many-objective optimization under deep uncertainty. Adv. Water Resour. 104, 195–209.
Turner, S.W., Marlow, D., Ekström, M., Rhodes, B.G., Kularathna, U., Jeffrey, P.J., 2014. Linking climate projections to performance: a yield-based decision scaling assessment of a large urban water resources system. Water Resour. Res. 50 (4), 3553–3567.
Watson, A.A., Kasprzyk, J.R., 2017. Incorporating deeply uncertain factors into the many objective search process. Environ. Model. Software 89, 159–171.
Weaver, C.P., Lempert, R.J., Brown, C., Hall, J.A., Revell, D., Sarewitz, D., 2013. Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks. Wiley Interdiscipl. Rev.: Clim. Change 4 (1), 39–60.
Wilby, R.L., Dessai, S., 2010. Robust adaptation to climate change. Weather 65 (7), 180–185.
Zeff, H.B., Herman, J.D., Reed, P.M., Characklis, G.W., 2016. Cooperative drought adaptation: integrating infrastructure development, conservation, and water transfers into adaptive policy pathways. Water Resour. Res. 52 (9), 7327–7346.