Version of Record: https://www.sciencedirect.com/science/article/pii/S1364815219300362
Manuscript_45df1b57554597ffa9db0cee15777257

Using multivariate regression trees and multiobjective tradeoff sets to reveal fundamental insights about water resources systems

Rebecca Smith[a], Joseph Kasprzyk[a], Balaji Rajagopalan[a]

[a]Department of Civil, Environmental, and Architectural Engineering, University of Colorado, 607 UCB, Boulder, Colorado, 80309, USA

Abstract

This paper presents the use of Multivariate Regression Trees (MRTs) to analyze Multiobjective Evolutionary Algorithm (MOEA) tradeoff sets generated from a long-term water utility planning problem. MOEAs produce large sets of non-dominated solutions, where each solution represents an observation of how multiple predictor variables (decision levers) impact performance in multiple response variables (objectives). Because they explicitly accommodate multiple response variables, MRTs can preserve the relationships between objectives revealed through MOEA-assisted optimization. We generated MRTs for two tradeoff sets that resulted from optimizing the Eldorado Utility planning problem under two climate change scenarios. A single MRT helped identify the subset of core planning decisions that led to preferred performance and demonstrated how decision preferences impacted performance in different objectives. Comparing MRTs from two scenarios revealed decisions that performed well across scenarios. The systematic and repeatable MRT approach can help water managers understand large, high-dimensional tradeoff sets and prompt additional promising analyses.

Highlights

- MOEA tradeoff sets contain information that can be hard to extract heuristically
- MRTs offer an unbiased, repeatable method to analyze MOEA tradeoff sets
- MRTs can reveal core planning decisions that perform well across future scenarios

Keywords

Multivariate Regression Tree (MRT), Multiobjective Evolutionary Algorithm (MOEA), feature selection, long-term planning, Front Range, Colorado

1 Introduction

Many academic studies and, recently, several real-world applications of Multiobjective Evolutionary Algorithms (MOEAs) have established the tool's ability to produce innovative solutions and valuable performance information about water resources planning problems (CSU, 2017; Maier et al., 2014; Reed et al., 2013). Through intelligent search and evaluation of thousands of potential portfolios, MOEAs produce suites of nondominated alternatives that explicitly elucidate relationships between conflicting planning objectives and the complex interactions among decisions that affect those objectives (i.e. tradeoff sets). Thorough analysis of the information contained in such tradeoff sets generally requires working with hundreds or thousands of solutions and considering many performance and decision characteristics, or dimensions, simultaneously. Interpreting this magnitude and complexity of relational information is difficult, so it is important to develop tools that facilitate water managers' ability to understand causes, effects, and trends in decisions and performance that are embedded in the results generated by MOEAs.

To date, Water Resources Systems Analysis (WRSA) research applications of MOEAs have mostly relied on parallel axis plots and/or glyph plots (Kasprzyk et al., 2009; Kollat & Reed, 2007; Matrosov et al., 2015; Smith et al., 2016) or Cartesian plots (Mortazavi et al., 2012; Wu et al., 2017) for insights, performing relatively subjective assessments on the tradeoffs to frame assertions of different performance priorities. Subjective explorations are useful for gaining general familiarity with the tradeoffs and identifying individual portfolios of interest, but they may not result in insights about the fundamental system dynamics that drive performance.

The volume of solutions and large numbers of decisions and objectives that make MOEA tradeoff sets difficult to analyze heuristically also make them good candidates for employing feature selection – the process of systematically reducing the dimensionality of a data set by distinguishing the most sensitive features from those that are noisy, redundant, or irrelevant (Liu et al., 2010), thus identifying fundamental system properties. Since the 1990s, feature selection has been used as a pre-processing step to improve subsequent data mining applications in a wide range of fields such as bioinformatics (Saeys et al., 2007), satellite imagery classification (Jain and Zongker, 1997), social network exploitation analysis (Zheleva and Getoor, 2009), and financial fraud detection (Ravisankar et al., 2011).

Feature selection has been applied specifically to MOEA tradeoff sets in a number of fields, though the term is not common in such studies. In this context, various feature selection approaches are presented as the main data mining event (not preparation for data mining) and used for the purpose of "knowledge discovery". Bandaru et al (2017) provide a thorough review of this literature; here we will focus on work that has employed regression trees, which are the specific type of feature selection/data mining method in which we are interested.

Regression tree models are generated by recursively partitioning data into two mutually-exclusive sets in order to sort the data into groups that have similar attributes. They are a popular method of feature selection because they are very versatile and easy to interpret. The partitioning process does not require or assume any specific distributions within the data, and it can uncover hidden structures and interactions between hierarchical and nonlinear variables (Prasad et al., 2006; Verbyla, 1987). Among many predictor variables, the method can determine which have the greatest influence on response (Lawrence and Wright, 2001). The binary rules are easy to understand for users who do not have expertise in statistics, and the tree structure itself is an intuitive way to visualize a model.

The few studies that have used regression trees to perform feature selection on MOEA tradeoff sets have been limited to univariate regression trees, which relate a single objective (response) to multiple decision levers (predictors). Sugimura et al (2010) demonstrated regression tree analysis on the design of a centrifugal impeller. The case study optimized 2 objectives using 16 design parameters and then generated a regression tree for each objective individually. Dudas et al (2011) optimized process rules for a 3-objective automotive production line system using 22 variables, and like Sugimura et al, created a separate tree for each objective. In another study, Dudas et al optimized investments in a production line using three objectives and nine decisions, this time eliciting a preferred performance region from decision makers and generating trees based on different methods of measuring solutions' spatial relationships to the region (2014). While these univariate regression trees provide useful information, they either separate or collapse the relationships between the objectives, and thus do not capitalize on one of the primary benefits of using an MOEA.

The Multivariate Regression Tree (MRT) was developed to relate predictor variables to multiple response variables while maintaining the individual characteristics of the responses (De'Ath, 2002). It originated in the field of ecology and was designed to be able to relate populations of multiple species to a number of independent environmental conditions. Importantly, MRTs do not make any assumptions about the underlying relationships between the response variables. MRTs have been used previously for feature selection, for example by Questier et al (2005) to analyze how the presence or absence of various chemicals predicts certain types of smells. However, to our knowledge, they have not been applied to feature selection using an MOEA tradeoff set. The versatility of the method (i.e. there is no requirement of any sort of data structure) suggests it can be successfully used to analyze the complex dynamics found in such data.

This study makes two contributions. First, it builds on previous efforts to perform feature selection on MOEA tradeoff sets using regression trees by newly applying MRTs to the task. Second, it employs feature selection to analyze tradeoff sets generated from optimizing a complex water supply system. By applying MRTs to a long-term water resources planning study performed using the Eldorado Utility Planning Model, we demonstrate how the method can facilitate and expand on heuristically-derived insights by extracting latent information about how specific combinations of decisions impact different types of performance, and about which decisions are likely to perform well in a wide range of potential futures. Such insights may either not be discernable from heuristic analyses alone, or the process of discovering them may require applying preferences that are not agreeable to all parties involved in developing a plan.

In the following section we present information about our methods: MOEA tradeoff sets and regression trees. We then give background on the Eldorado Utility case study used in the optimization. Next are the regression tree results, followed by discussion of their implications for practical applications in water resources planning and also for future research. The last section offers concluding remarks.

2 Methods

2.1 MOEA tradeoff sets

Multiobjective Evolutionary Algorithms (MOEAs) are a search technology used to efficiently generate and evaluate alternative solutions to systems whose conflicting performance objectives are impacted by many decisions that exhibit complex interactions (Reed et al., 2013). In the context of long-term water supply planning, the MOEA intelligently designs and tests thousands of different combinations (or portfolios) of decisions such as reservoir sizes and conservation levels to optimize performance in objectives such as maximizing storage reliability and minimizing frequency of water-use restrictions. When attempting to optimize multiple conflicting objectives, improvement in one objective requires sacrificing performance in another, so there are tradeoffs. During optimization, the MOEA removes from the preferred group any portfolio whose performance is worse than another portfolio in all objectives; that is, the *dominated* portfolios are removed. The end product of MOEA-assisted optimization is a set of nondominated planning portfolios that quantitatively characterize the performance tradeoffs of a system.

The nondominated tradeoff set is valuable because it represents the system information learned by the MOEA as it refines combinations of decision levers to achieve better results in the objectives. Each portfolio within the MOEA tradeoff set is an observation of how multiple predictor variables (decision levers) affect a system's performance in multiple response variables (objectives). Framing the dataset in this way motivates the use of MRTs to help extract the relational information contained in the tradeoffs.

2.2 Multivariate Regression Trees (MRTs)

MRTs are an extension of the univariate Classification and Regression Tree (CART) algorithm (Breiman et al., 1984). (Trees generated from categorical data are termed classification trees; we are working with quantitative data and, as such, will limit our description to regression trees.) CART has been used for feature selection in several fields (Chebrolu et al., 2004; Gomez-Chova et al., 2003) and also for other data mining purposes in WRSA (Bryant and Lempert, 2010; Kwakkel and Jaxa-Rozen, 2016). For CART, the steps and calculations presented below would be performed on values of a single response variable; for MRTs, they are instead applied to the geometric centroid defined by the summed Euclidian distances to the means of all response variables. To convert from MRT to CART, one only needs to remove the innermost summation from Equations 1 through 3 found in Section 2.2.1.

Note that because objectives incorporated into MOEA optimization often measure very different types of quantities, all objective values need to be scaled before an MRT is generated to prevent objectives with large units and ranges from dominating the splitting. This preserves the equal weighting of objectives, which is a core concept that underpins the value of MOEAs. Though there are several ways of standardizing or normalizing data, we recommend scaling the observations for each objective to a range of 0 to 1 because this approach does not distort within-objective distribution or across-objective relationships. Decision lever values do not need to be scaled.

2.2.1 MRT algorithm

The steps of the MRT algorithm will be presented in terms of an MOEA tradeoff set: portfolios are observations, decision variables are the predictors, and objectives are the responses.

A. Calculate the error of the full data set at the root node:

Equation 1

$$E_{root} = \sum_{i=1}^{N} \sum_{j=1}^{J} \left(y_{ij} - \bar{y}_{j(N)}\right)^2$$

Where N is the number of portfolios in the tradeoff set, J is the number of objectives, y_{ij} is a portfolio's value of objective j, and $\bar{y}_{(N)}$ is the mean of all values of objective j.

B. For every split between values in every decision lever, sum the error (impurity) within and across each of the two child nodes that would result from splitting the data by that decision lever value:

Equation 2

$$\mathrm{E}_{split} = \sum_{k=1}^{2} \sum_{i=1}^{n} \sum_{j=1}^{J} (y_{ij(k)} - \bar{y}_{j(k)})^2$$

where $y_{ij(k)}$ is a portfolio's value of objective j, $\bar{y}_{j(k)}$ is the child node's mean value of objective j, k is the child node formed by the split, and n is the number of observations in child node k.

C. Split the parent node using the decision lever and value from Step B that resulted in the smallest value of E_{split}.

D. Repeat Steps B and C for each child node until a user-specified stopping criterion is met. When the criterion is met, splitting terminates and a node becomes a leaf. Stopping criterion determines the number of leaves (i.e. the size of the tree). Within-leaf error is defined as:

Equation 3

$$\mathrm{E}_{leaf} = \sum_{i=1}^{n} \sum_{j=1}^{J} (y_{ij(k)} - \bar{y}_{j(k)})^2$$

The explanatory power of a tree is traditionally captured by its relative error; this value represents how much of the root error was not resolved by sorting the portfolios via recursive splits in decision variables:

Equation 4

$$RE_{tree} = \frac{\sum_{l=1}^{L} \mathrm{E}_{leaf}}{E_{root}}$$

Where L is the total number of leaves on the tree, $y_{ij(l)}$ is a portfolio's value of objective j, $\bar{y}_{j(l)}$ is the leaf's mean value of objective j, and n is the number of observations in leaf l.

Though it offers the same technical information as relative tree error, we propose use of the complementary "explained variance" quantity (Cannon, 2012) to summarize the overall explanatory power of the tree:

Equation 5

$$EV(\%) = (1 - RE_{tree}) \times 100$$

Focusing on minimizing "error" is a misleading characterization in terms of understanding the tree's value. As such, EV will be reported for MRTs in this study.

2.2.2 Cross validation

The cross validation technique commonly used when generating MRTs (which is also used in this study) is 10-fold cross validation. For each "fold", a model is trained on 90% of the data and tested by calculating the RE_{tree} that results once the withheld 10% are placed in their respective leaves. The average of the 10 RE_{tree} values is reported back as the cross validated relative error (CVRE). This process is undertaken for every potential size of tree to enable comparison between RE_{tree} and CVRE, which can guide users in determining optimal tree size.

2.2.3 MRT selection

In any statistical modeling application, steps must be taken to ensure that the model appropriately represents the underlying data to the extent that the method can do so meaningfully. For regression trees, it is also important to ensure that the structure of the tree is understandable for users, so the most appropriate tree will need to balance descriptive power and interpretability. This can be achieved by using standard cross validation procedures and/or knowledge of the data to prune the tree (Murphy, 2012). One common way of determining the best tree is by examining the progression of its CVRE as size increases. If CVRE stagnates or starts to increase, the model is said to be “overfit”- there is no gain, or potentially a loss, of explanatory power as the tree increases in complexity. In such cases, a rule of thumb proposed by Breiman et al (1984) is to choose the smallest tree that achieves the minimum CVRE plus one standard error (the Min + 1SE rule). However, not all data sets exhibit divergent error behavior. Other approaches to determining appropriate tree size include requiring a minimum number of observations per node, defining the minimum amount of error reduction that must be met for a split to proceed, and pre-specifying tree size or depth.

2.2.4 Example MRT

To preview the conceptual and visual results of generating an MRT, we present a simple example problem and tree. The water utility for a growing city has two objectives: to minimize frequency of imposing annual water-use restrictions on customers and also minimize the amount of expensive and environmentally-disruptive new storage it has to build to meet increasing demands. The decision levers available to the city are to build (or not build) a reservoir with up to 100 million cubic meter (MCM) of capacity, and to enact (or not enact) conservation measures that would reduce demand by 20% (a binary decision; 0 = no conservation, 1 = conservation enacted). Figure 1 represents how an MRT generated for this problem might look.

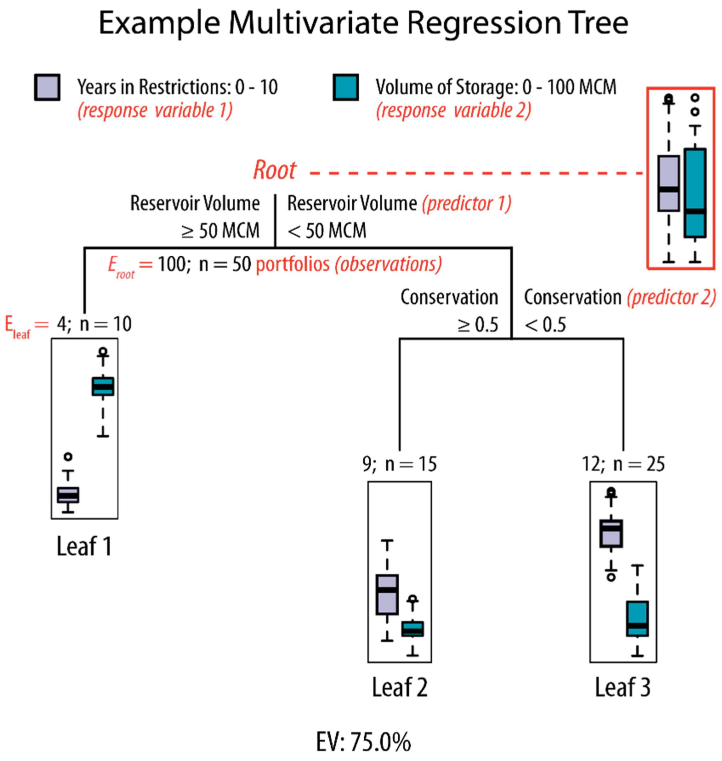

Figure 1. Example MRT. Red text and objects denote MRT concepts or terms that are not found on actual MRT plots but are provided here for clarity. All values are hypothetical and designed to explain MRT features.

The MRT has two splits and three leaves. Each leaf is comprised of a set of portfolios whose range of performance in each of the two planning objectives is characterized by a boxplot. At the root, the boxplot for each objective would span the entire plotting range (included in Figure 1 for demonstration purposes). Each of the three leaves has a different configuration of performance tradeoffs across the two objectives because the MRT divides the tradeoff set such that the variance in performance is reduced in one or more objectives.

The MRT shows how many portfolios are in each leaf and the error remaining within each leaf, calculated as the summed differences between each portfolio's actual performances in each objective compared to the mean performances in each objective of all in-leaf portfolios (see Equation 3). While "error" is appropriate in terms of the quantity calculated, it is actually just an indication of spread around the means and does not indicate that portfolios are incorrectly placed in the leaf. Leaves with larger numbers of portfolios tend to have higher error because there are more errors to sum, and leaves with a large distribution of performance in one or more objectives (i.e. large boxplot ranges) will also have higher in-leaf error. (Leaf 1 has the fewest portfolios and tightest ranges, so in-leaf error is smallest; Leaf 3 has the most portfolios and larger ranges for both objectives, so it has higher error.)

The branches of the tree are formed by splitting the portfolios based on their values of the two decision levers. The first split is based on reservoir size. This indicates that reservoir size explains the most variance across all objectives. Leaf 1 shows that the larger reservoirs require larger volume of new storage, as indicated by the height of the teal boxplot, and also that a larger reservoir will result in relatively fewer years of restrictions, as indicated by the low placement of the light purple boxplot. Portfolios with reservoirs less than 50 MCM in capacity (the right-hand tree branch) are further split based on whether conservation was enacted, and that distinction leads to two different ranges for years in restriction, as indicated by the different heights of the light purple boxplots in leaves 2 and 3.

The EV value at the bottom of the plot indicates how well the tree was able to organize the set of portfolios into groups of similar performance characteristics. This tree was able to explain 75% of the variance using two splits (the sum of error across leaves is 25 and the root error is 100). There is no threshold of EV that indicates whether a tree is valid or useful; if a user sees value in the percentage of variance explained, then the tree is valuable assuming the tree was appropriately pruned using cross-validation criteria and problem knowledge. Furthermore, it would not actually be desirable to generate a tree that explained 100% of the variance for two reasons: it would be unwieldy, and it would eliminate the opportunity to use human reasoning to explore flexibility in the decision space as opposed to relying on the MRT algorithm to exhaustively organize the portfolios.

2.2.5 Analyzing MRTs

Applications of MRTs in other fields are often motivated by understanding how things co-occur, which focuses on the relative relationships of response variables (generally species of plants or animals) within and across MRT leaves (e.g. two species may be very prominent in leaves characterized by certain environmental conditions but rarely found in others, so they are likely to be found together) (Davidson et al., 2010; De'Ath, 2002; Herzschuh and Birks, 2010; Larsen and Speckman, 2004). In this context studies often examine population composition at each split in the MRT to identify the species that are most influential in each partition and also perform within-leaf calculations to determine "indicator species". Another common use of MRTs is to use splits to delineate geographic regions either using latitude and longitude directly or via climate or ecosystem variables that can be mapped (Cannon, 2012; Hamann et al., 2011; Salonen et al., 2012). The first type of application is very leaves-focused and the second type gains most information from the splits.

A major difference between using MRTs to explain population composition or climate effects vs. applying them to MOEA tradeoff sets is that MOEA tradeoff sets exist to facilitate the elucidation and navigation of preferences. Users have values that determine which performance tradeoffs are acceptable and also have opinions about the decisions that comprise a portfolio. This suggests two approaches to analyzing MRTs generated from MOEA tradeoffs that are valuable because of the ability to navigate relationships between leaves and splits (objectives and decisions) iteratively.

2.2.5.1 Leaves-first analysis

After visual inspection of all leaves, users will be able to identify a subset that represent preferred patterns of performance tradeoffs and then review the splits (portfolio decisions) that led to the leaf. Referring back to Figure 1, a user that values minimization of water-use restrictions far more than avoidance of building new storage would focus on Leaf 1. Once Leaf 1 is identified, the user would learn that reaching the leaf requires at least a 50 MCM reservoir.

2.2.5.2 Root-first analysis

Without the benefit of MOEA tradeoff sets to facilitate in-depth discussion of performance tradeoffs, water utilities typically focus on decision preferences when crafting portfolios to test during long-term planning studies (Smith et al., 2018). In this paradigm, an MRT user would start at the root of the tree and at each split determine the preferred value of a decision. Following decision preferences down the tree to one or more leaves would reveal how decision preferences affect performance tradeoffs and potentially help users see where compromises are needed to avoid unacceptable performance. The MRT in Figure 1 would demonstrate to a user that if they wanted to avoid a large reservoir (right branch) and not enact conservation (right branch), they could expect relatively frequent incidence of water-use restrictions. This likely would not be considered desirable.

2.2.6 Software

This study used the *mvpart* R package (De'Ath, 2014, 2002; R Core Team, 2016), which is archived but still functional. Its primary function executes the algorithm described in Section 2.2.1, and the most important parameters are those that control cross validation and the complexity parameter (CP), the stopping criterion that defines the amount of error reduction that must be achieved by a split to continue growing the tree. We note our settings in the Results section. The only functionality we altered was the standard plotting included with the package; *mvpart* generates a set of bar plots for the mean objective values at each leaf, but we replaced the bars with boxplots to give more information. The repository of data, packages, and code necessary to reproduce this study's results can be found in the Acknowledgements.

3 Case study

3.1 Front Range, Colorado

The Front Range of Colorado is an urban corridor on the eastern slope of the Rocky Mountains that encompasses several mid-sized cities and many smaller communities. Water providers in the region rely heavily on runoff from highly variable annual mountain snowpack, so storage is critical for weathering intra- and interannual water supply fluctuations (Doesken, 2014; Rajagopalan et al., 2009). The long-term impact that climate change will have on Colorado's hydrology is unclear; temperatures are expected to continue increasing, but precipitation could increase or decrease (Lukas et al., 2014). However, despite the possibility of increased precipitation, there is likely to be less water in the future due to the dominance of higher temperatures (Udall and Overpeck, 2017; Woodbury et al., 2012). In addition to the natural supply variability and uncertainty from climate change, the Front Range is experiencing the compounding challenge of rapid population growth; the regional population is projected to increase by 40% by 2050 (State of Colorado, 2017).

Water management in Colorado is further complicated by the prior appropriation doctrine, a legal framework that bases the succession of streamflow access on date of first use ("first in time, first in right") (Hobbs, 2004). Farmers and energy companies own the vast majority of senior water rights in the state, and by 1900 most of the water in eastern slope rivers was fully appropriated (Eschner et al., 1983). This means that as cities grew, they collected a mixture of supplies from multiple locations (including the western slope of the Rockies) by acquiring junior streamflow diversion rights, building junior reservoirs, buying senior diversion rights from agriculture, or buying shares in other water companies. All long-term utility planning involves making many decisions and balancing conflicting objectives; on the Front Range, these inherent difficulties are exacerbated by rapidly increasing demand, highly uncertain impacts of climate change, complex regulations, and contentious social and environmental issues. This context is the basis of our MOEA case study, briefly described in the next section.

3.2 Eldorado Utility Planning Model

The Eldorado Utility Planning Model was designed based on input from 11 Front Range water managers to generically capture important regional management features and challenges (Smith et al., 2017). It encompasses the region surrounding a small municipal water provider called the Eldorado Utility. Eldorado is located on the eastern slope of a mountain range along with eight other water users that directly compete with the utility to divert and store water. Eldorado has mostly junior diversion rights, junior storage rights in two reservoirs that it owns, and also has shares in a water wholesale company that it takes out of a reservoir owned by that entity. One of Eldorado's diversion rights comes from the western slope, where an additional four users impede the utility's access to water.

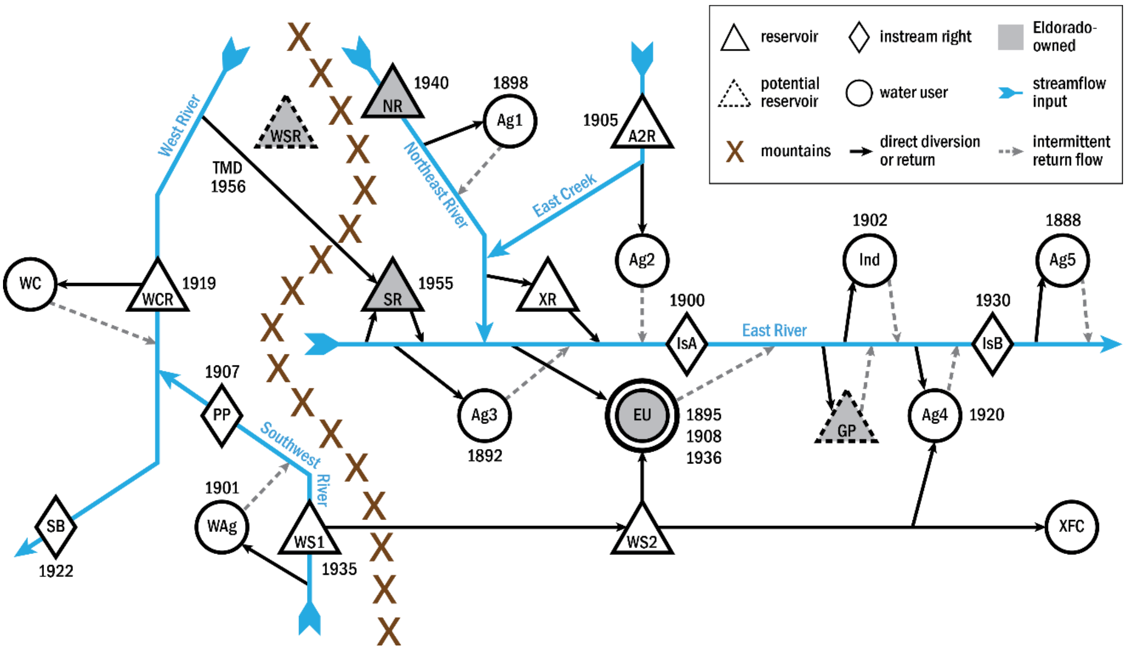

Figure 2. Schematic of the Eldorado Utility Planning Model. Many different users on both slopes of the mountain range impact Eldorado's (EU symbol at bottom center) ability to collect and divert water via their priority dates, the locations of their diversions, and the locations of their return flows (precise diversion and return flow locations are indicated by arrows). Each user or right in the diagram has a priority date associated with it where applicable and is listed in Table 1. Reprinted from Smith et al **(**2018**)** with permission from ASCE.

Table 1. Detail for water users in Eldorado Utility Planning Model. Abbreviations refer to those found in Figure 2. The order of users going down each table column corresponds approximately to reading left-to-right on the diagram. Bolded users are particularly relevant to the results presented in Section 4. Superscripts in the table are defined as follows: [A]Res = Reservoir; [B]MCM = million cubic meters; [C]KAF = thousand acre feet; [D]Ag = Agriculture; [E]cms = cubic meters per second; [F]cfs = cubic feet per second. Reprinted from Smith et al **(**2018**)** with permission from ASCE.

Abbr.	Name	Magnitude of Rights	Abbr.	Name	Magnitude of Rights
SB	Southern Basin	varying flow	XR	External Res	varying vol
WC	Western City	n/a	**A2R**	**Ag2 Irrigation Co. Res**	**24.7 MCM (20 KAF)**
WCR	Western City Res[A]	24.7 MCM[B] (20 KAF[C])	Ag2	Ag2 User	n/a
WAg	Western Ag[D] User	4.3 cms[E] (150 cfs[F]) seasonal	**EU**	**Eldorado Utility**	**0.28 cms (10 cfs); 0.34 cms (12 cfs); 0.42 cms (15 cfs)**
PP	Power Plant	varying flow	**WS2**	**Wholesaler Res 2**	**123.3 MCM (100 KAF)**
TMD	TransMtn Diversion	2.2 cms (80 cfs)	IsA	Instream Flow A	varying flow
WSR	**West Slope Res**	**varying vol; 2.2 cms (80 cfs)**	GP	Gravel Pit	1.0 MCM (800 AF)
WS1	**Wholesaler Res1**	**616.7 MCM (500 KAF)**	**Ind**	**Industrial User**	**varying flow**
NR	North Res	11.1 MCM (9 KAF)	Ag4	Ag User 4	1.4 cms (50 cfs) seasonal
SR	South Res	9.9 MCM (8 KAF)	IsB	Instream Flow B	0.42 cms (15 cfs)
Ag3	Ag User 3	1.4 cms (50 cfs) seasonal	XFC	External Farms & Cities	n/a
Ag1	Ag User 1	1.4 cms (50 cfs) seasonal	Ag5	Ag User 5	2.9 cms (100 cfs) seasonal

The model incorporates a wide range of water rights dates to capture the temporal complexity created by prior appropriation. It also has great spatial complexity to reflect the fact that in Colorado, water is constantly being diverted from and returned to the stream. Overall, there are five distinct basins in the model, each with a streamflow input site at its headwaters. The model was designed such that, under historic hydrology, Eldorado's existing system and sources could meet 100% of current demands with only rare need for restrictions. Different future streamflow scenarios that alter timing and volume of streamflow require the utility to take action in order to meet growing demands. These scenarios and demands are described in Section 3.4. For more detailed model and optimization information refer to Smith et al (2018).

The Eldorado Utility Planning Model was built using the RiverWare modeling software (Zagona et al., 2001). RiverWare's advanced capabilities facilitated our use of prior appropriation water allocation and enabled us to manage ownership of water through its accounting functionality. The model uses over 150 custom rules to operate the intricate relationships between objects, users, and accounts, and is an example of the kind of complex decision support system that many utilities have incorporated into their planning (Labadie, 2004).

3.3 Problem formulation

The problem formulation includes 13 decision levers and 7 objectives which are briefly described here.

3.3.1 Decision Levers

Eldorado Utility has a total of 13 decision levers available to enable it to meet growing demands with potentially more challenging streamflow conditions. Some increase the system's operational flexibility, some involve acquiring or freeing up water, and some develop new storage. They are briefly described below and summarized in Table 2. Where applicable, lever descriptions include a reference to the relevant user in Figure 2.

3.3.1.1 Enhancing operations

Certain water sources in Colorado are reusable; cities carefully monitor their return flows from unconsumed water so that they can re-divert reusable return flows to meet demands. This is only possible by legally acquiring the right to exchange the water from downstream to upstream and only works well with strategic storage options. Three levers help Eldorado take advantage of reusable return flows: *Exchange* determines whether the legal right is acquired to store reusable water in a reservoir owned by Eldorado; $LeaseVol_{XRes}$ determines the amount of dedicated exchange storage space Eldorado rent in the External Res (XR); and $Lease_{Ag2Res}$ determines whether Eldorado is allowed to use available space in Ag2 Irrigation Co. Res (A2R) to store reusable water.

3.3.1.2 Increasing supply

There are three ways that Eldorado can gain access to "new" supplies. The utility can acquire portions of water rights of other users in the model, it can buy shares of water companies in the model, and it can create water through conservation or increasing distribution efficiency. Eldorado may purchase up to 20% of the rights of Ag3 User (Ag3) ($Rights_{Ag3}$) and Industrial User (Ind) ($Rights_{Industrial}$). Ag3 rights are very senior and may be stored but are not available year-round; Industrial rights are mid-seniority and must be directly diverted from the stream, but are available year-round. Eldorado may buy shares from either Wholesaler (WS1, WS2) ($Shares_{Wholesaler}$) or Ag2 Irrigation Co. (A2R) ($Shares_{Ag2}$). Through $Shares_{Interruptible}$ the utility may also execute a contract for access to A2R shares that is triggered when Eldorado's storage is severely depleted. Acquiring water from any of these sources will draw water away from regional agriculture and industry and potentially disrupt those communities. Finally, Eldorado may enact none, moderate, or aggressive conservation measures (*ConsFactor*) or increase distribution efficiency (*DistEff*) by up to 3%.

3.3.1.3 Building storage

There are three opportunities for Eldorado to increase the amount of storage it owns. The utility may expand the existing South Res (SR) to help store both existing and new eastern slope and western slope water ($ExpandVol_{SouthRes}$). Eldorado can build a new West Slope Res (WSR) to store its existing western slope diversion right ($BuildVol_{WestSlopeRes}$); this is a very challenging proposition because of regulatory, social, and environmental considerations. Lastly, the utility can develop gravel pits (GP) downstream of its return point to capture reusable flows (*GP*).

Table 2. Summary of Eldorado Utiltiy decision levers. MCM = million cubic meters; AF = acre-feet.

Decision	Description	Units	Range	Increment
Enhancing Operations				
Exchange	Acquire right to exchange reusable return flows to NorthRes	---	0 - 1	Binary
LeaseVol$_{XRes}$	Pay owners of XRes to lease dedicated space that can facilitate Exchange	MCM (AF)	0 – 3.7 (0 - 3,000)	0.12 (100)
Lease$_{Ag2Res}$	Pay Ag2 Irrigation Co. to store water in any available unused space; 0 = off, 1 = on	---	0 - 1	Binary
Increasing Supply				
Rights$_{Ag3}$	Purchase a portion of Ag3's senior diversion right	%	0 - 20	1%
Rights$_{Industrial}$	Purchase a portion of Industrial user's mid-seniority diversion right	%	0 - 20	1%
Shares$_{Wholesaler}$	Purchase additional shares of Wholesaler water	shares	0 - 6,000	10
Shares$_{Ag2}$	Purchase shares of Ag2 Irrigation Co. water	shares	0 - 10,000	100
Shares$_{Interruptible}$	Negotiate agreement with Ag2 Irrigation Co. for optional supply leases	shares	0 - 10,000	100
ConsFactor	Reduce starting per capita demand through conservation measures; 0 = no change, 1 = 10% reduction, 2 = 20% reduction	---	0 - 2	1
DistEff	Improve distribution efficiency by reducing unaccounted-for water (e.g. fixing leaks, improving metering, etc.)	%	90 - 93	1%
Building Storage				
ExpandVol$_{SouthRes}$	Expand SouthRes	MCM (AF)	0 – 2.47 (0 – 2,000)	0.12 (100)
BuildVol$_{WestSlopeRes}$	Build West Slope Res	MCM (AF)	0 – 12.3 (0 - 10,000)	0.12 (100)
GP	Develop gravel pits to store reusable return flows downstream of the city; 0 = not developed, 1 = developed	---	0 - 1	Binary

3.3.2 Objectives

Brief qualitative descriptions of the seven objectives are given below. For further detail and equations, please refer to the Appendix.

The first three objectives, *RestLev1*, *RestLev2*, and *RestLev3*, seek to minimize the total number of years (within the 25-year simulation) that Eldorado goes into three restriction levels of increasing intensity. To comply with Eldorado’s current reliability policy, the utility can only go into each level 5, 1, and 0 years out of 25, respectively.

The fourth objective, *MissedOpp*, minimizes the average annual volume of water that the utility “misses”, i.e. when timing of demand or availability of storage space prevent Eldorado from capitalizing on the full amount of its water rights. Optimizing how efficiently Eldorado uses the water it has helps prevent wasteful acquisitions.

Objective five, *New Supply*, seeks to minimize the average annual volume of water Eldorado uses from new sources. Though the utility does need to acquire or create new water to meet growing demands, they do not want to take more than they need for future water security.

The sixth objective, *April1Storage*, maximizes carryover storage of the lowest storage-to-annual demand percentage recorded during the 25-year simulation. April 1 is the approximate date when reservoirs would be at their lowest levels before spring runoff begins to fill them again and is a measure of carry-over storage. Compared with the restrictions-based objectives, this captures a longer term reliability signal because it evaluates performance based on the worst-performing year of the simulation.

Finally, *NewStorage* minimizes the volume of newly-built storage within each portfolio. Because storage is difficult to permit and socially and environmentally contentious, Eldorado seeks to carefully consider the number and size of storage projects it pursues. The combination of this and the *NewSupply* objective provide a cost-like signal and allow the utility to consider planning policy on a broader level (Smith et al., 2018).

3.4 Scenarios

The optimization runs using the Eldorado Utility Planning Model assumed a buildout demand based on 40% population increase by 2050 (State of Colorado, 2017), when the simulation time horizon starts. The demands exhibit single family residential patterns, i.e. use increases substantially during summer months when lawns are irrigated. The irrigation demands go up slightly during dry years and are affected by conservation and distribution efficiency levers, but the baseline population demand does not change throughout the simulation.

Because future streamflow in Colorado is highly uncertain, the set of studies associated with this model use several hydrologic scenarios. The scenarios relevant to this article are the 1°C- and 4°C-warmer futures, which were chosen based on a Front Range climate change study (Woodbury et al., 2012). The perturbed hydrology used monthly changes (i.e. deltas) from that study and generated sets of stochastic headwater streamflow. The stochastic simulation first generates annual streamflow timeseries using a KNN resampling approach (Lall and Sharma, 1996), which are disaggregated to monthly flows using the proportional disaggregation method of Nowak et al (2010). The monthly deltas from for the warming scenario from Woodbury et al (2012) are then applied.

3.5 Optimization implementation

We used the Borg MOEA for this study (Hadka and Reed, 2013), which tests have shown to perform similarly or favorably compared to other state-of-the-art algorithms on difficult benchmark problems (Reed et al., 2013; Zatarain Salazar et al., 2016). The Eldorado Utility Planning Model embedded in the search loop simulates the supply and usage dynamics of Eldorado Utility and other regional water users over 25 years (from 2050 to 2075) at a monthly timestep. Portfolios were tested as fully-implemented configurations of Eldorado's system.

Performance of each portfolio was evaluated across ten hydrologic traces, each distributed to a separate computing core of an Amazon Web Services Elastic Compute Cloud (EC2) instance (Mathew and Varia, 2014). Each distributed simulation took approximately 20 seconds. This relatively long simulation time prompted us to limit search to 5,000 function evaluations. Though this number of evaluations is lower than that of many other MOEA studies, the resulting tradeoff set is sufficiently large and diverse to demonstrate the MRT method. We used the default Borg settings except for changing initial population size from 100 to 50.

3.6 Heuristic analysis of Eldorado optimization tradeoffs

3.6.1 Sample analysis using parallel axis plots

Figure 3 presents a set of Eldorado Pareto-optimal portfolios from a 1°C-perturbed optimization run. We will use the set to facilitate readers' understanding of MOEA tradeoff sets and describe how the Eldorado Utility Planning Model captures Front Range, Colorado, water management tradeoffs. The figure also offers an opportunity to demonstrate the challenge of heuristically analyzing the results of MOEA-assisted optimization. The performance and decision tradeoffs of the set of 961 portfolios are presented using parallel axis plots, which are a visual analytics technique commonly used in multiobjective optimization studies (Herman et al., 2014; Kasprzyk et al., 2013; Watson & Kasprzyk, 2017).

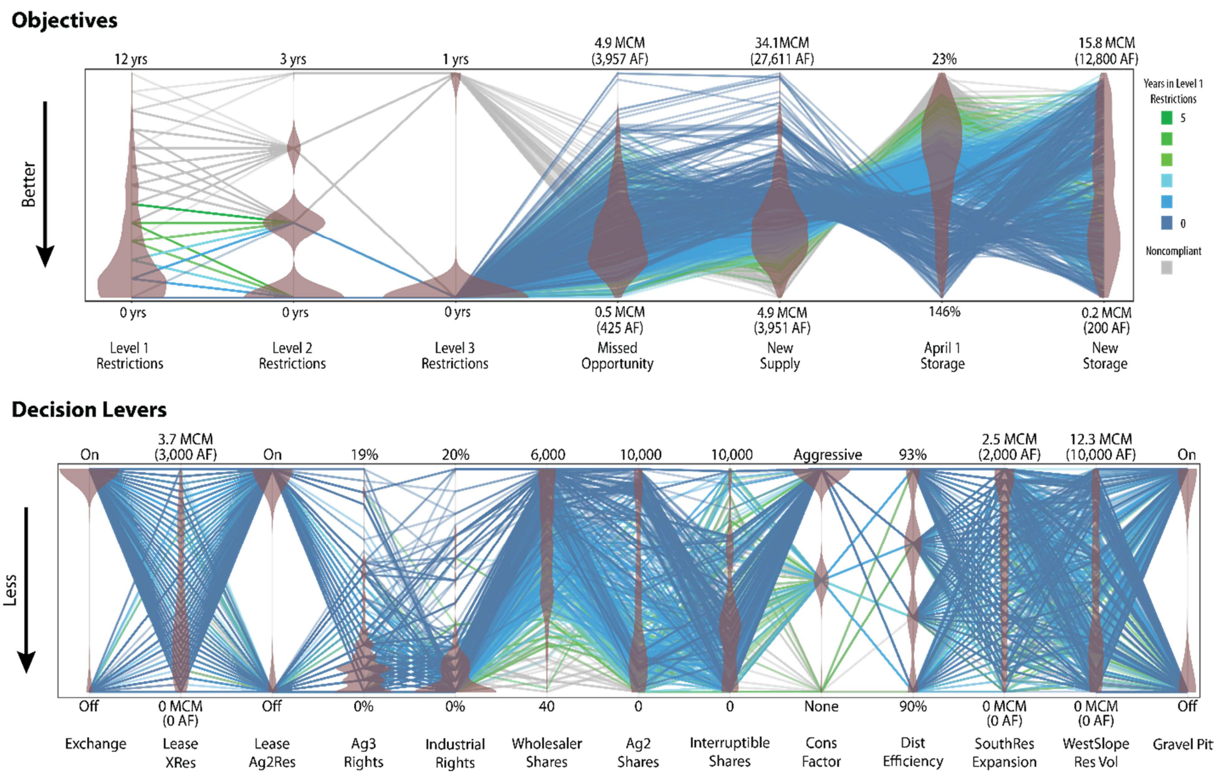

Figure 3. Parallel plots of the tradeoff set resulting from optimizing the Eldorado Utility Planning Model under 1˚C-warmer hydrology. Plot **(a)** shows the relationships between different performance objectives. Plot **(b)** shows the portfolios of decisions that resulted in the performance from plot **(a)**. Grey portfolios do not comply with the utility's reliability policy (which allows up to 5, 1, and 0 years in Levels 1, 2, and 3 restrictions, respectively); compliant portfolios are colored based on years in Level 1 restrictions.

In Figure 3a, each of the seven performance objectives is represented by a vertical axis. Each of the 961 portfolios is represented by a segmented line that crosses each axis at the level of performance it achieves in that objective, where crossing lower on an axis denotes better performance. (Note that April 1 Storage is the only maximization objective, so even though higher levels of storage are better, that is still represented by lower positioning on the axis.) The portfolio lines are colored based on whether they comply with Eldorado's reliability policy (grey portfolios are noncompliant) and then the number of years they were in Level 1 restrictions, with dark blue corresponding to zero years at the bottom of the leftmost axis. The "violins" on the axes show portfolio densities to clarify trends in performance that can be obscured due to overlapping lines. Figure 3b is oriented identically to Figure 3a except that there are 13 axes – one for each decision lever. Every portfolio line in Figure 3a has a corresponding line in Figure 3b that conveys the amounts or levels of all of the decisions within the portfolio. The lower a line crosses an axis in Figure 3b, the less of that decision has been chosen.

In Figure 3a we can see relationships between the objectives. Color enables us to tell that all of the dark blue portfolios with zero years in Level 1 restrictions have medium to high levels of New Supply (fifth axis from the left), medium to high levels of April 1 carryover storage, but may have anywhere from 0.2 to 15.4 MCM (200 to 12,500 AF) of New Storage (rightmost axis). This means that to minimize years in Level 1 restrictions, it is imperative that Eldorado obtain new water sources but may choose to build or avoid large amounts of new reservoir storage. However, portfolios that do not build much New Storage perform more poorly in April 1 Storage and tend to require greater volumes of New Supply. This shows an important tradeoff within the Eldorado model as well as on the Front Range: utilities often have to choose between meeting growing demands with new supplies that come from conservation and other users' shares and rights, which may be socially and economically disruptive to communities, and relying on contentious, expensive infrastructure that is difficult to permit.

Filtering the portfolios (here, through color) based on reliability compliance represents non-subjective criteria that a utility might use to begin learning how portfolios' decision attributes relate to performance. Working up from the bottom of the Wholesaler Shares axis in Figure 3b we can see that as these increase, performance in Level 1 restrictions improves. The same is true for increasing levels of conservation – only green portfolios with three or more years in restrictions plot at the bottom, indicating that no conservation was enacted.

3.6.2 Limitations of heuristic analysis

Figure 3 was designed to demonstrate the specific points described in the paragraphs below it. That is, the order of the axes and the use of color supported a narrative. This demonstrates a fundamental problem with how many-dimensional datasets are presented that is exacerbated by the existence of user preferences: shapes, colors, and data orientation all influence perception of underlying system dynamics.

Visually inspecting the parallel plots revealed useful information about tradeoffs and trends in two decision levers. Beyond this point in a heuristic analysis an issue arises: what is the next move? A manager could ask what would happen if the set was filtered to exclude portfolios that have *any* instances of restrictions. Another idea would be to exclude from the compliant set any portfolios that build a contentious West Slope Reservoir. Or we could try to focus on portfolios that have lower amounts of New Storage, but the violin in Figure 3a shows that there is no natural break point at which to segment the portfolios and thus the differentiation between levels of performance would be subjective.

The logic effectively applied to the early stages of the heuristic analysis is the result of the authors' years of expertise with MOEA tradeoff sets, parallel axis plots, and the Eldorado Utility system; not every MOEA user could do this. Despite having comfort and experience with the tradeoffs, the array of paths we suggested above shows how quickly the analysis can become ambiguous and potentially counterproductive: the questions that managers try to answer with the tradeoff set and the order in which they are asked will heavily influence perceptions which could then be difficult to dislodge despite re-ordering; the conditions on which the results are filtered may be based on preferences that are not shared by all parties involved in a planning process (e.g. minimizing New Storage may align with judgements about whether reservoirs are environmentally responsible). Even if analyses are iterative, relying heavily on conflicting preferences to orient filtering may increase users' focus on different positions.

Finally, because of computational limitations, most practical applications of MOEAs to WRSA problems will require users to cull hydrology. However, decisions that perform well in one set of potential hydrologic futures may or may not be robust given different conditions. When MOEAs are used for planning under deep uncertainty, multiple rounds of optimization can be used to address shortcomings inherent to scenario-specific optimization (Eker and Kwakkel, 2018; Watson and Kasprzyk, 2017). The issues described above compound when multiple tradeoff sets are generated.

The list below summarizes the general limitations of using a heuristic approach alone to analyze MOEA tradeoffs:

- visual representations of many-dimensional data involve subjective decisions that can influence perceptions of system dynamics;
- humans are not good at deciphering patterns across many dimensions that could include complex interactions;
- predetermined decision and performance preferences may heavily influence heuristic analyses and prevent users of MOEA tradeoff sets from seeking or learning fundamental system dynamics;
- different users' perceptions of system dynamics resulting from subjective heuristic analyses could exacerbate conflict; and
- without objective information about decision and performance dynamics it is more difficult to draw conclusions when working with multiple tradeoff sets.

MRTs are a simple, automated approach to analyzing MOEA tradeoffs that produce multiple insights simultaneously. Including them alongside heuristic analyses of tradeoff sets provides a neutral and repeatable foundation that can clarify results and orient additional investigations. In the next section we present the results of using MRTs for feature selection and describe how their structure and insights can enhance system and tradeoff understanding.

4 Results

We performed two separate optimizations of the Eldorado Utility case study – one for the 1°C-perturbed hydrology and one for the 4°C – and created an MRT for each set of tradeoffs. We first describe the results of an MRT generated from the 1°C-perturbed portfolios described in the previous section, and follow that discussion with an MRT from a 4°C-perturbed optimization. Generating two trees helps to validate the use of the MRT method on tradeoffs from the

Eldorado Utility case study and enables us to gain additional insights into the system behavior by comparing them.

4.1 MRT for 1°C-perturbed tradeoff set

The plot of CVRE vs. tree size shown in Figure 4 was produced by fitting an MRT to the Eldorado Utility 1°C-perturbed tradeoff set described in Section 3.6. We allowed the MRT algorithm to build a large tree based on a CP of 0.001. A very small CP value allowed us to analyze the progression of CVRE over the course of many splits. The minimum CVRE shown here is marked by the red dot (though the CVRE would likely continue to decrease very slowly as the tree grew) and the tree size that corresponds to the Min + 1SE rule is marked in yellow.

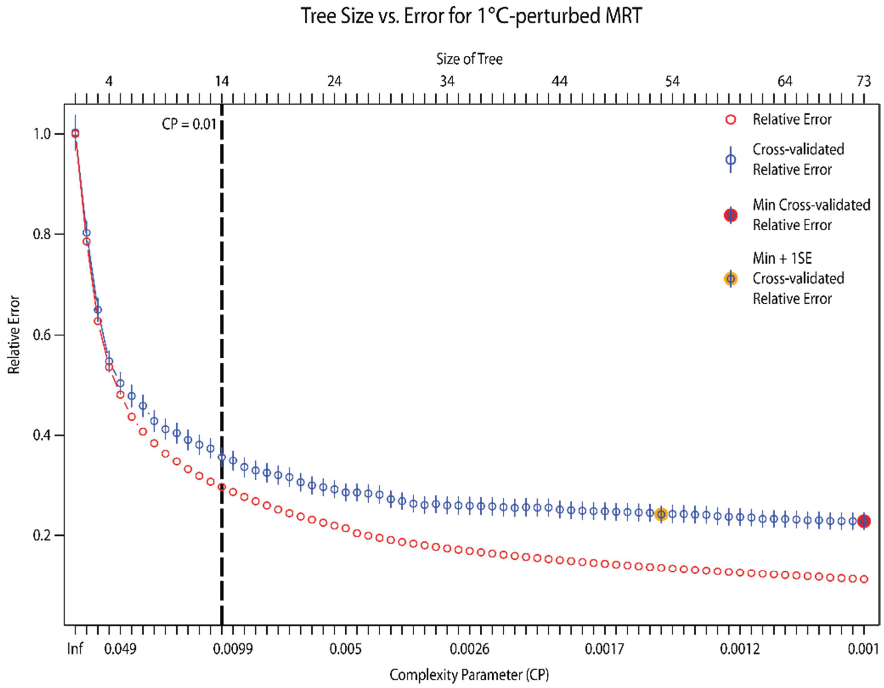

Figure 4. CVRE vs. size of MRT for 1°C results from Eldorado Utility optimization.

At no point does the CVRE start to increase or fully stagnate, so choosing the tree size for this data set is more subjective. The choice made for this study was to require that in order for a split to occur it must meet an error reduction threshold of at least 1% of the root error, so the CP was set to be 0.01. This corresponds to a tree with 14 leaves (marked by the vertical dashed line) and a maximum tree depth of 5 splits. The value was chosen heuristically by balancing simplicity, descriptive value, and meaningful interpretation of the criterion. The next “round” CP value would be 0.005 and result in an unwieldy 28-leaf tree.

Figure 5 and Figure 6 present the left and right branches, respectively, of the MRT generated from the 961 portfolios in the 1°C-perturbed tradeoff set described in Section 3.6. We will first orient the reader to the features of the tree and then discuss different approaches to analyzing it.

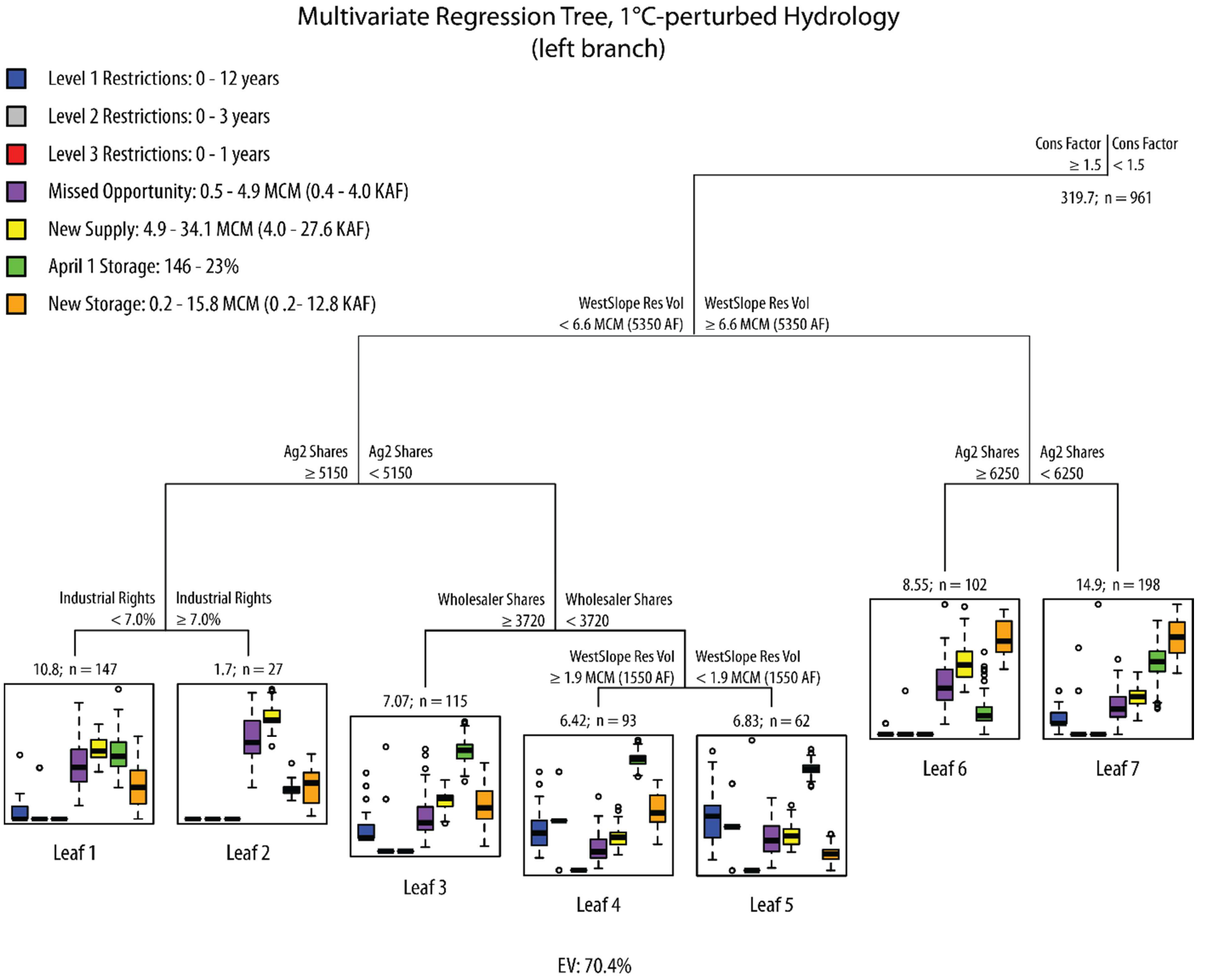

Figure 5. Left branch of the multivariate regression tree generated from the Eldorado Utility 1˚C optimization tradeoffs.

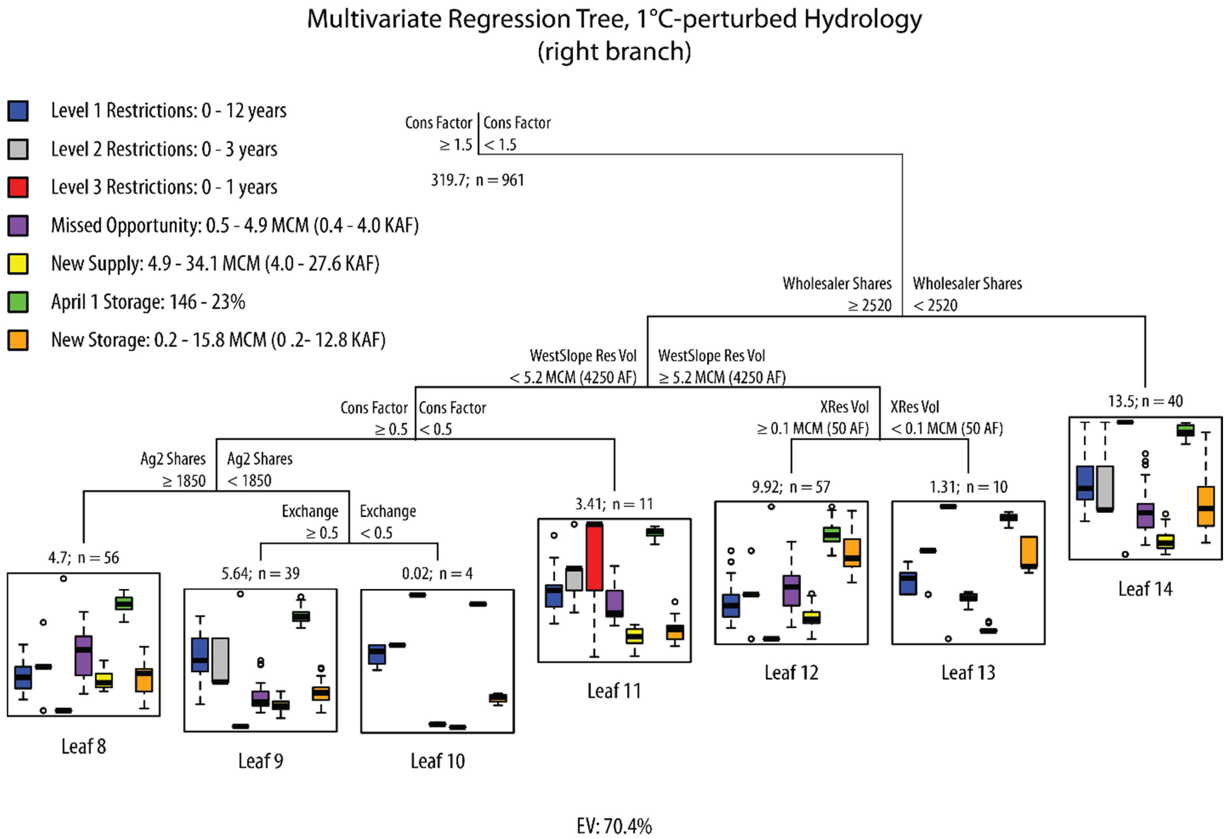

Figure 6. Right branch of the multivariate regression tree generated from the Eldorado Utility 1˚C optimization tradeoffs.

At the top of the tree halves in Figure 5 and Figure 6 we see that the first split is based on the conservation level incorporated into each portfolio. No conservation is ConsFactor = 0, moderate conservation is ConsFactor = 1, and aggressive conservation is ConsFactor = 2. The left branch includes portfolios where ConsFactor is greater than or equal to 1.5, i.e. portfolios that have aggressive conservation. The number reported is the average between the levels of decision above and below the split. As another example, following the left branch, the next split is on the volume of West Slope Res. To the left are portfolios that have reservoirs up to 6.5 MCM (5300 AF), and to the right go the portfolios that have reservoir volumes starting at 6.7 MCM (5400 AF). The granularity of the split value depends on the increment of a decision lever (presented in Table 2).

Following splits down to the leaves, each leaf has a set of boxplots: one for each of the seven objectives denoted by the color and ranges shown in the legend. The order of the boxplots is the same as the order in which the objectives were first described, which is also their order in Figure 3a. And, like the parallel plots, the lower a boxplot is positioned within the plot area, the better the performances of the portfolios within the leaf. The EV value at the bottom indicates that the tree explains 70.4% of the performance objective variance within the data set.

4.1.1 Analyzing the tree: leaves-first

Our first analysis will start at the leaves, consider the ranges of performance for the objectives, assert a set of priorities to direct focus on a single leaf, and then follow the branches up to the root to see what decision rules produced that leaf. For example, Eldorado Utility managers (and, by extension, managers in the Front Range who were the basis of the Eldorado model) may want to prioritize reliability-related objectives (Smith et al., 2019). Given that criteria, leaves that have boxplots that are very low with small ranges in the first three objectives (blue, grey, and red) would contain portfolios of interest. Examination of the leaves in Figure 5 and Figure 6 shows that there are three that meet that boxplot configuration- leaves 1, 2, and 6 (see Figure 5). Focusing on leaves 2 and 6, which are superior to Leaf 1 in years in Level 1 Restrictions, will help illustrate the value of MRTs and connect them to recognizable tradeoffs. Figure 7 provides a close-up comparison of the two sets of boxplots.

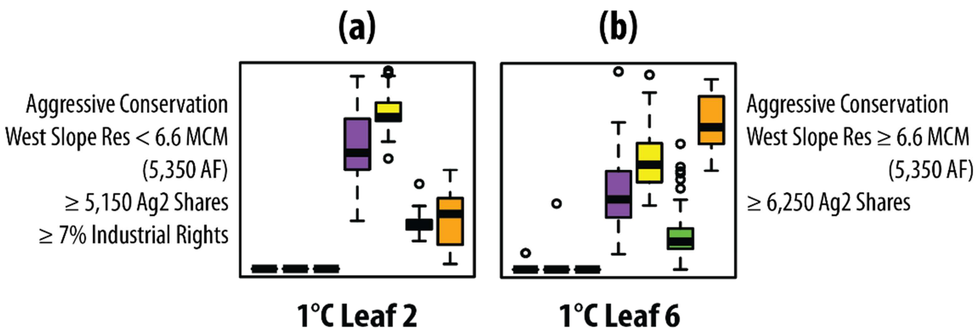

Figure 7. Comparison of two leaves from the 1˚C MRT. Note that both leaves incorporate Aggressive Conservation and have a very similar number of Ag2 Shares.

The decision rules that lead to Leaf 2 are: aggressive conservation; a West Slope Res smaller than 6.6 MCM (5350 AF); 5,200 or more shares of the Ag2 Irrigation Co.; and at least 7% of Industrial User's water rights. None of the portfolios have any incidence of any level of restrictions, they have moderate to high volumes of MissedOpp water, a very high range of NewSupply (the highest range of all the leaves), medium-high April 1 carryover storage, and moderate to low volumes of NewStorage. Despite having zero years in restrictions, the April 1 carryover storage objective is not as high (i.e. positioned as low) as might be expected because the portfolios within the leaf have relatively low amounts of NewStorage.

As shown in Figure 5 and emphasized in Figure 7, the path to Leaf 6 includes aggressive conservation, a West Slope Res 6.6 MCM (5350 AF) or larger, and at least 6300 Ag2 Shares. The performance ranges in Leaf 6 are notably different than in Leaf 2. Among the portfolios in Leaf 6, there is one occurrence of Level 1 restrictions and one occurrence of Level 2 restrictions, moderate volume of MissedOpp water, moderate to high NewSupply, high to moderate volumes of April 1 carryover storage, and high to very high volumes of NewStorage. Incorporating the larger West Slope Res reduced Leaf 6's reliance on NewSupply (e.g. via the Industrial Rights required in Leaf 2), but the portfolios are therefore more likely to have large amounts of NewStorage. The patterns in these two leaves echo the tradeoff between NewSupply and NewStorage discussed for the parallel plot in Figure 3a. The ability to confirm these MRT results with surface-level visual analysis provides more confidence in the MRT findings that are harder to deduce heuristically, such as the importance of a large amount of Ag2Shares, which shows up in both leaves.

Emphasizing leaves 2 and 6 as superior to others in reliability objectives does not preclude other leaves and other sets of decisions from containing portfolios that match Eldorado's performance priorities. The leaves simply indicate that after sequentially splitting the portfolios based on all of the relationships within the tradeoff set, these particular sets of decision levers are most likely to result in appealing portfolios. Furthermore, the decisions in the paths to highly reliable leaves must still be accompanied by actions in the other decision levers; there is just more flexibility in the values for the levers not represented in splits.

4.1.2 Analyzing the tree: root-first

Analyzing the tree starting from the leaves up as described in the previous section is a way of asserting *performance preferences* and understanding which decisions are likely to lead to good performance. Starting from the root and working down allows users to understand the impact on performance of *decision preferences*.

Using the tree branch in Figure 6, we can demonstrate the four steps of a path that an Eldorado manager might take down the tree if a general policy of new water sources but limited reservoir expansion was preferred.

1. At the first split, a manager may choose to go to the right because she or he does not want to have to rely on aggressive conservation to meet performance goals.
2. At the next split, a manager may choose to go left because Wholesaler Shares are a reliable water source that does not require infrastructure.
3. Next, a manager may go left to avoid a large West Slope Res because of cost, permitting, etc.
4. Finally, the manager may go right to see how bad the outcomes could be if no conservation was enacted.

Leaf 11 is the outcome of applying these decision preferences, and the boxplots reveal that they will likely result in decent performance in NewSupply and NewStorage but poor performance in the other objectives. This manager would have learned that the combination of decisions in this path will likely result in non-preferable performance regardless of the other 10 decisions in the portfolio.

4.1.3 Reviewing MRT insights

The insights gained from MRTs would likely have been difficult to obtain through heuristic approaches, but they should also be verified (because splits are not guaranteed to be meaningful) and built upon using different types of analyses. One option is to use interactive visual analytics software such as Tableau (Jones, 2014) to manually apply MRT splits and further explore portfolios within leaves of interest. As each split down a branch is applied to the dataset, for example by filtering on decision levers, ranges in one or more objectives should shift indicating that the split had meaningful impact. Figure 8 provides an example by revisiting the parallel plots used in Figure 3.

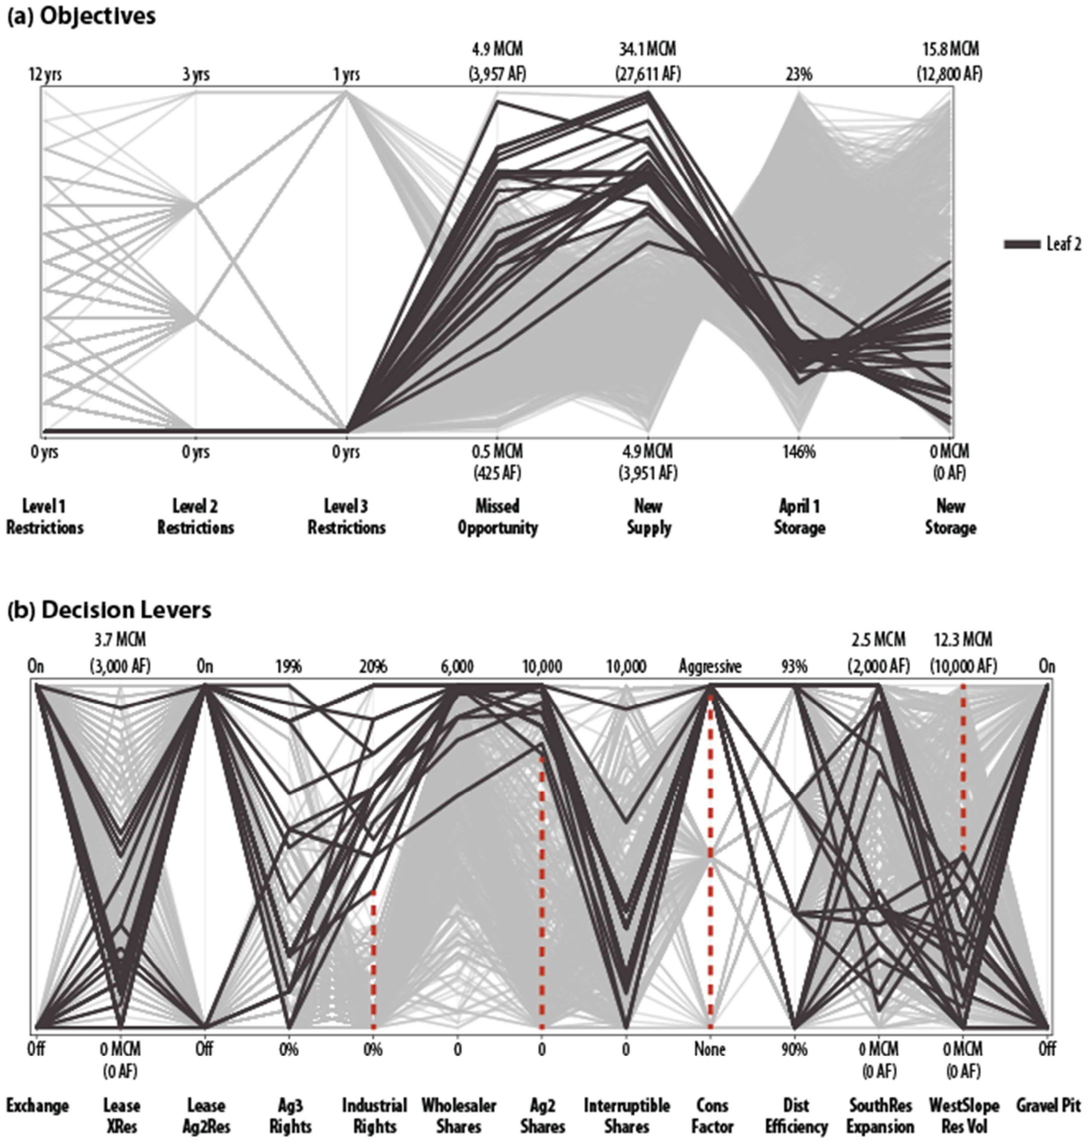

Figure 8. Parallel plots from the 1˚C Eldorado Utility tradeoff set, with the portfolios contained within Leaf 2 of the 1˚C MRT emphasized.

Figure 8a and Figure 8b are oriented almost exactly like the plots in Figure 3; the only difference is that in this figure, color is used to distinguish the set of 27 portfolios from Leaf 2. In Figure 8a, the pattern and ranges of the Leaf 2 portfolios' performance across the seven objectives matches the boxplots from Figure 7a. The ranges of ConsFactor, West Slope Res, Ag2Shares, and Industrial Rights in the decision levers in Figure 8b reflect the splits, and red dashed axis lines highlight the restricted ranges of those decisions. In eight of the remaining nine decision dimensions, there is considerable variety in potential values to accompany the constrained decisions. The levels of Wholesaler Shares are almost universally very high, though, so this decision lever correlated closely with a split based on another decision lever and the large numbers of Wholesaler Shares are contributing to the preferred performance though they were not explicitly represented in the MRT.

Other types of visualizations such as pair-wise scatter plots between individual decisions and objectives or correlation matrices between objectives and decision levers may also be helpful in MRT verification. If there are no identifiable relationships between the decision levers represented in splits and one or more objectives, users should be cautious in their interpretation of MRT results.

4.2 MRT for 4°C-perturbed tradeoff set

All previous discussions of tradeoffs, portfolios, and trees have referred to a set of portfolios generated from optimizing for a 1°C-warmer future. Planning in consideration of multiple possible future scenarios is beneficial in and of itself, and it also increases the impact of MOEA-based MRTs. Figure 9 and Figure 10 present an MRT generated from a set of portfolios optimized for 4°C-perturbed hydrology. Tree size was determined by using the same logic that was described for the 1°C tree in Section 4.1 and the 1% error reduction criteria was used again. After briefly describing a few features specific to the 4°C tree, we discuss findings from comparing the two trees.

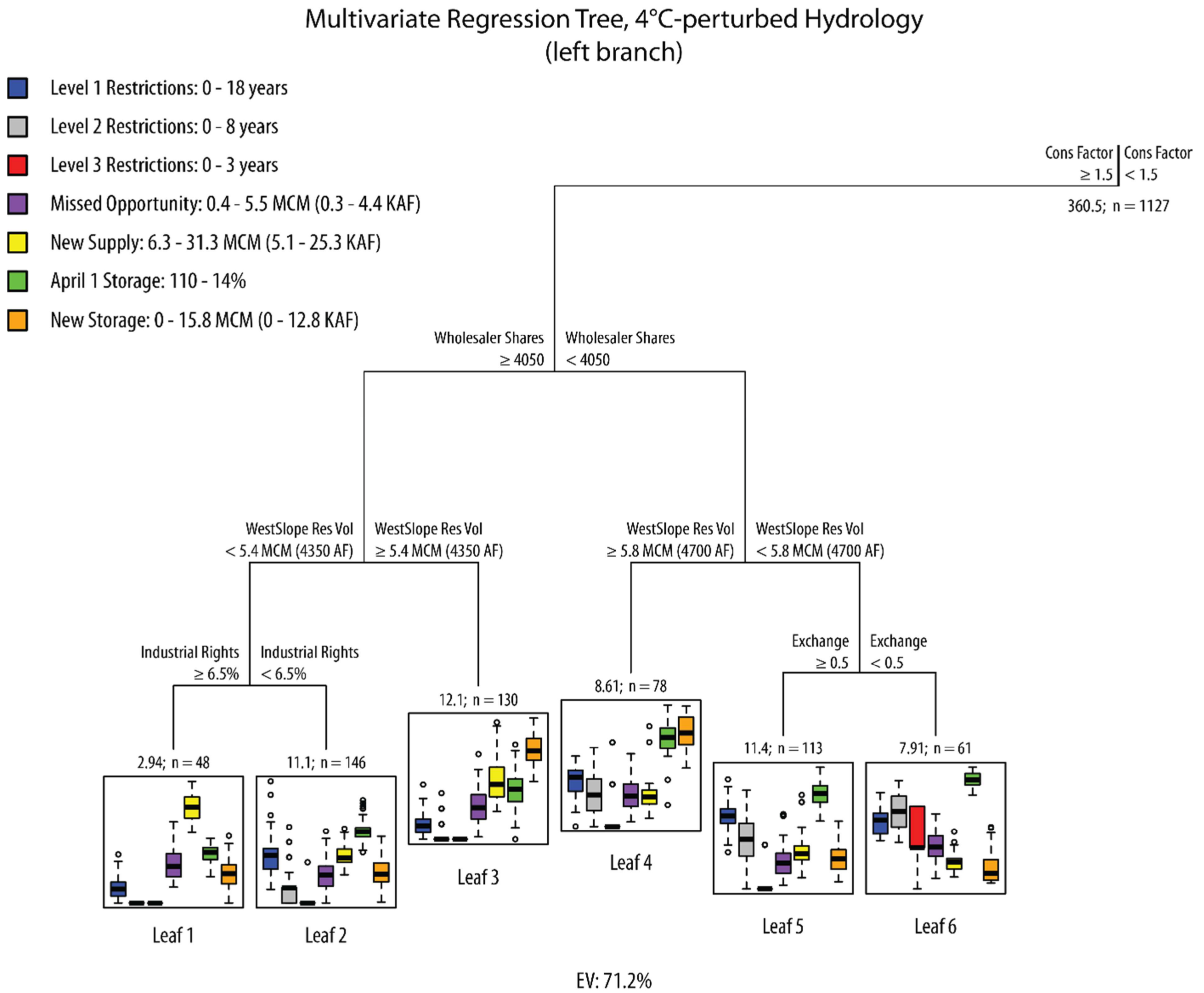

Figure 9. Left branch of the multivariate regression tree generated from the Eldorado Utility 4˚C optimization tradeoffs.

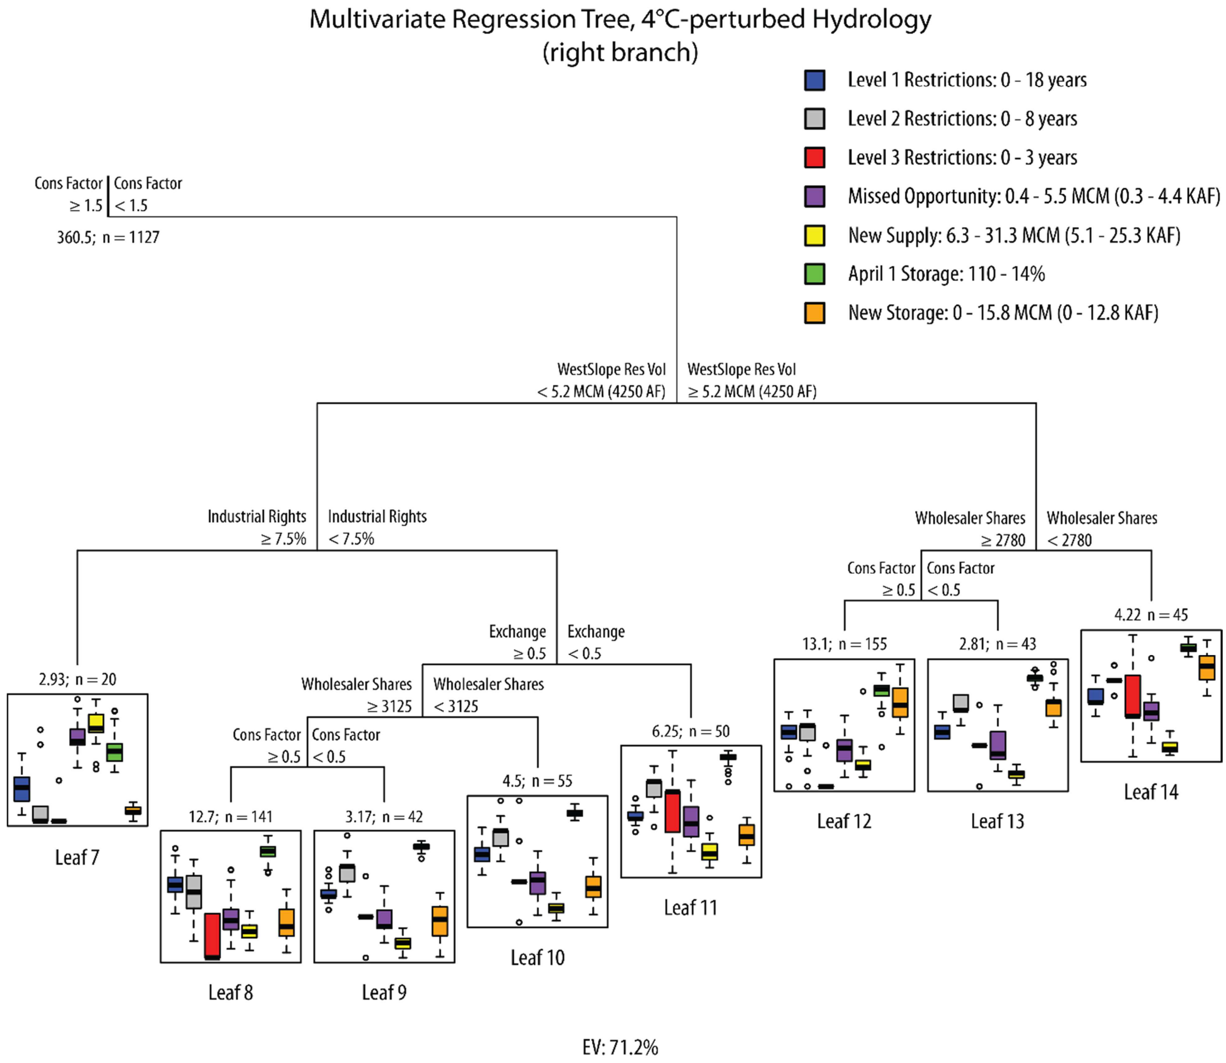

Figure 10. Right branch of the multivariate regression tree generated from the Eldorado Utility 4˚C optimization tradeoffs.

The root error and total number of portfolios are given at the root node of the MRT in Figure 9 and Figure 10. Splits, leaves, boxplots, colors, and objective ranges are all oriented the same as in Figure 5, but note that the objective ranges are different. This is especially relevant in the first three objectives (years in levels of restrictions); the more challenging hydrology resulted in more frequent restrictions and fewer portfolios with low incidence of curtailment. The EV indicates that the tree explains 71.2% of the performance variance found within the tradeoff set.

If we repeat the same leaf-first exercise from the 1°C tree, where we determined that the performance preference was to have minimal years in all three levels of restrictions, that criteria reduces viable leaves down to two: Leaf 1 and Leaf 3 from Figure 9. The decision path to Leaf 1 includes aggressive conservation, at least 4060 Wholesaler Shares, a West Slope Res less than 4.3 MCM (4350 AF), and at least 7% of Industrial Users rights. The splits for Leaf 3 are, like Leaf 1, aggressive conservation and at least 4060 Wholesaler Shares, but then instead of a small West Slope Res and a percentage of Industrial rights, Leaf 3 includes a West Slope Res at least 5.4 MCM (4350 AF) in volume. A comparison of the two leaves shows that they exhibit the same NewSupply-NewStorage tradeoff seen in the 1°C MRT and the original parallel plots of the 1°C tradeoffs. As noted for the 1°C tree, this agreement the parallel plots and the 4°C tree signals that the MRT is accurately capturing major dynamics while providing more detailed latent information.

4.3 Comparing MRTs

Comparing the broad characteristics of the two trees provides valuable information. First, we note that the decisions on which splits occur are very similar across both trees: ConsFactor, West Slope Res, and Industrial Rights are prominent in both trees. In the 1°C tree, Ag2 Shares are more important, while in the 4°C tree, Wholesaler Shares are more important. Since Wholesaler Shares are a western slope source and Ag2 Shares are eastern slope, this may be indicative of a shift in basin yields with warmer temperatures. The general agreement in splits suggests that these decisions are the most influential factors in a portfolio in either scenario, and this is a fundamental insight about the Eldorado system.

We can expand on this general decision lever agreement by comparing sets of leaves from the two trees. First we will compare Leaf 2 from the 1°C tree and Leaf 1 from the 4°C tree, as shown in Figure 11. The decisions that lead to these leaves with very similar objective tradeoffs include three nearly identical splits: aggressive conservation, a medium or smaller West Slope Res, and approximately 7% or more of the Industrial rights. Ag2 Shares in 1°C are traded for Wholesaler Shares in 4°C.

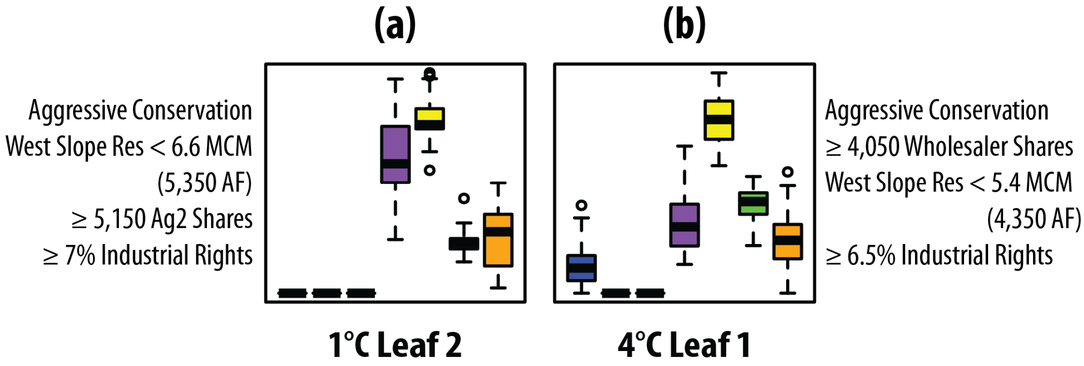

Figure 11. Comparison of Leaf 2 from the 1˚C MRT and Leaf 1 from the 4˚C MRT. Note that Aggressive Conservation, West Slope Res, and Industrial Rights have identical or similar values in both leaves.

Now compare Leaf 6 from the 1°C tree and Leaf 3 from the 4°C tree in Figure 12. Like the previous comparison, the patterns of objective performances are similar, and they share two almost identical splits: aggressive conservation and medium to large West Slope Res. Again, Ag2 Shares in 1°C are replaced by Wholesaler Shares in 4°C.

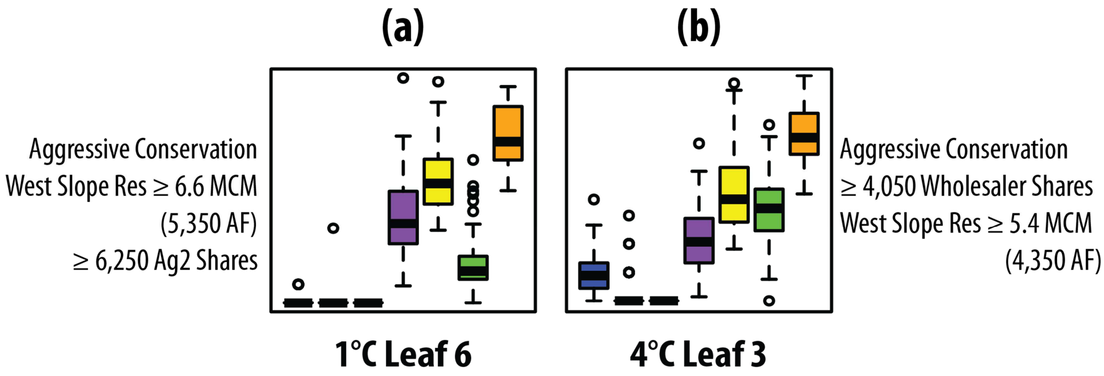

Figure 12. Comparison of Leaf 6 from the 1˚C MRT and Leaf 3 from the 4˚C MRT. Note that these leaves both incorporate Aggressive Conservation and a moderate-to-large West Slop Res.

The 1°C- and 4°C-perturbed hydrologies are substantially different in runoff timing, magnitude, and overall annual hydrograph shape. The presence of overlapping decisions across the two hydrologies for two fundamental planning strategies (minimizing NewStorage in Figure 11 vs. minimizing NewSupply in Figure 12) suggests that each strategy corresponds to a set of core decisions and that these decisions are robust to a wide range of futures.

5 Discussion

The previous section focused on interpretation of MRTs, but we also wish to address larger implications of their use in both WRSA practice and research.

5.1 Context for incorporating MRTs into tradeoff analyses

MRTs offer an unbiased and repeatable method of using MOEA tradeoff sets to learn about fundamental system dynamics and gain detailed information about which decisions are likely to have the most impact on system performance. While these are valuable contributions, an MRT is not a substitute for in-depth exploration of a tradeoff set and the information it provides should be combined with system knowledge to ensure accurate interpretation. Ultimately, MRTs are a promising tool to orient and enhance other types of tradeoff set analysis, all of which require technical skills and expert reflection to use properly.

5.2 Building on MRT results in practice

Generating and interpreting MRTs can result in fundamental system insights as demonstrated in Section 4, but the most important benefits are derived from the different ways that the predictive capabilities of the trees can be used. For example, once promising leaves are identified, managers can build on the core strategies without having to re-run an optimization. This can facilitate diving into unexplored parts of the decision space or answering policy questions that were not formally posed in the original formulation of the optimization in a computationally efficient manner. The trees can also provide quick answers to what-if scenarios or emerging contexts, such as if a new path is needed because a reservoir does not receive a critical permit or a source of water that was integral to a plan is no longer available.

The ability of MOEAs to generate a set of portfolios that have approximately-best performance for the many ways that tradeoffs can be balanced does not necessarily imply that the portfolio chosen will ultimately be contained within it; portfolios that are not strictly nondominated may be preferred for un-modeled reasons. The insights gained from nondominated MRTs can be used in a targeted exploration of dominated portfolios of interest, or MRTs could be generated using the full set of portfolios generated by the MOEA.

5.3 Future research

The findings and limitations of this study suggest several interesting areas of additional WRSA MOEA research. One idea is to use insights from MRTs to learn about a problem and constrain a new iteration of the problem formulation to target a specific region of objective performance. This would be a new approach to the de Novo planning framework demonstrated by Kasprzyk et al (2012). The identification of robust decisions across scenarios is also a promising result that warrants structured analysis to develop a framework to guide the use of MRTs for this purpose. Another beneficial direction would be to test MRTs on different types of water resources applications, e.g. reservoir operations, to determine whether the results are meaningful in non-planning contexts. We also suggest exploring the efficacy and value of combining different feature selection methods with WRSA tradeoff sets, and then assessing whether they are considered useful and usable by practitioners who have used or are interested in employing MOEAs in their planning processes (Smith et al., 2017).

6 Conclusion

The increasing prevalence of MOEA studies in WRSA research and practice calls for greater attention to developing tradeoff analysis tools. While tradeoff sets are complex and often challenging to analyze heuristically, the high dimensionality and large volume of results produced by MOEAs can be assets when combined with feature selection. Here we present MRTs, which relate performance variations within and across multiple objectives to distinct subsets of specific decisions, providing users with information about the most consequential decisions and their most productive ranges.

Using the Eldorado Utility Planning Model, we demonstrated multiple types of analysis that can be performed with MRTs. Starting with the MRT leaves, managers may identify groups of portfolios that correspond to their performance priorities and learn which decision splits were critical to arriving at promising leaves. Alternatively, starting at the top of the tree and following decision splits down based on decision preferences provides information about how these preferences impact performance across multiple objectives. Finally, comparing leaves from different trees may shed light on decisions that perform well across multiple futures. Insights gained from all of these approaches can inform in-depth exploration of the tradeoff set and prompt new policy questions.

MRTs are versatile, simple to generate, and present easily comprehensible insights that may not be apparent during heuristic analyses of tradeoff sets. They overcome mostly or entirely the limitations described in this paper: the only subjective choice required in the visual representation of the tradeoffs is the order in which performance boxplots are places; they find patterns across many dimensions in an objective and repeatable process that eliminates the possibility of interference from user preferences that can skew perceptions and exacerbate conflict; and they provide an objective basis on which to compare multiple tradeoff sets.

Appendix

This appendix presents the underlying equations for the seven objectives qualitatively described in Section 3.3.2.

MOEA-assisted optimization evaluates performance based on an objective function vector, **F**(**x**), where **x** is the portfolio defined by decision lever values (described and defined in Section 3.3.2). Each value in the vector results from calculating a separate objective, $f_{objective}$.

Equation A-1

$$\mathrm{F}(\mathbf{x}) = (f_{RestLev1}, f_{RestLev2}, f_{RestLev3}, f_{MissedOpp}, f_{NewSupply}, f_{April1Storage}, f_{NewStorage},)$$

$$\forall\ \mathbf{x} \in \Omega$$

$f_{RestLev1}$, $f_{RestLev2}$, and $f_{RestLev3}$ are restrictions-based reliability measures. Restriction levels are triggered based on April 1 storage levels, which are used by Front Range, Colorado, utilities to assess their system status for the upcoming year. In the model, restrictions are represented by reductions in outdoor water use (while indoor use is never curtailed). Table A-1 summarizes the restriction triggers and impacts.

Table A-1. Storage-based triggers and water-use impacts of restriction levels.

Current Storage-to-Long-Term Avg Annual Demand	Restriction Level	Resulting Indoor Use	Resulting Outdoor Use
> = 75%	0	100%	100%
< 75%	1	100%	80%
< 50%	2	100%	50%
< 25%	3	100%	0%

where "Current Storage-to-Long-Term Avg Annual Demand" is defined as

Equation A-2

$$RestLev = \frac{Total\ Water\ in\ Storage\ on\ April\ 1}{Long\ Term\ Unrestricted\ Annual\ Utility\ Demand} \times 100$$

The three restrictions objectives are calculated as follows:

Minimize the number of years that Eldorado spends in Level 1 Restrictions:

Equation A-3

$$f_{RestLev1}(\mathbf{x}) = E\left[\sum_{i=1}^{Y} y_{RestLev_i=1}\right]_t$$

Minimize the number of years that Eldorado spends in Level 2 Restrictions:

Equation A-4

$$f_{RestLev2}(\mathbf{x}) = E\left[\sum_{i=1}^{Y} y_{RestLev_i=2}\right]_t$$

Minimize the number of years that Eldorado spends in Level 3 Restrictions:

Equation A-5

$$f_{RestLev3}(\mathbf{x}) = E\left[\sum_{i=1}^{Y} y_{RestLev_i=3}\right]_t$$

where Y is the number of years simulated per t traces in the hydrologic ensemble. Expectation notation, $E[\quad]$, denotes that the average across the traces was used.

The optimization seeks to minimize the fourth objective, $f_{MissedOpp}$, which measures how efficiently Eldorado uses its supplies and system components to meet demands. It is affected by whether the utility can capitalize on reusable water and also whether Eldorado acquires an overabundance of Wholesaler or Ag2 shares.

Equation A-6

$$f_{MissedOpp}(\mathbf{x}) =$$

$$E\left[\frac{1}{Y}\sum_{i=1}^{Y}(Unused\ Shares_{Wholesaler_i} + Unused\ Shares_{Interruptible_i} + Lost\ Reusable\ Return\ Flows_i\right]_t$$

Objective five, $f_{NewSupply}$, is also minimized, and quantifies the amount of "new" water that the utility acquires form shares and other water users or creates through conservation.

Equation A-7

$$f_{NewSupply}(\mathbf{x}) = E\left[\frac{1}{Y}\sum_{i=1}^{Y}\begin{matrix}(yield\ from: \\ Rights_{Ag3}, Rights_{Industrial}, Shares_{Wholesaler}, Shares_{Ag2}, ConsFactor, DistEff)_i\end{matrix}\right]_t$$

The sixth objective, $f_{April1Storage}$, seeks to maximize the amount of water Eldorado has in carryover storage on April 1 of every year.

Equation A-8

$$f_{April1Storage}(\mathbf{x}) = E\left[y_{min}\left(\frac{Total\ Eldorado\ April\ 1\ Storage\ Vol}{Avg\ Long\ Term\ Annual\ Demand}\right) \times 100\right]_t$$

where y_{min} denotes that the objective is calculated using the minimum annual value over the course of the simulation.

The final objective, $f_{NewStorage}$, minimizes the total volume of new storage that Eldorado builds.

Equation A-9

$$f_{NewStorage}(\mathbf{x}) = \sum[ExpandVol_{SouthRes}, BuildVol_{WestSlopeRes}, (GP * 0.99\ MCM)]$$

Note that GP is multiplied by 0.99 MCM (800 AF) because the GP lever is on/off or 1/0, but the volume added is 0.99 MCM (800 AF).

The optimization was subject to a single constraint- there could be no instance of unmet indoor demand:

Equation A-10

$$c_{UnmetDemand} = 0$$

Acknowledgements

This work benefitted from three funding sources: National Oceanographic and Atmospheric Administration (NOAA) Sectoral Applications Research Program (SARP), grant # NA14OAR4310251; the University of Colorado Boulder; and the Center for Advanced Decision Support for Water and Environmental Systems (CADSWES). Data, packages, and code necessary to reproduce these results can be found at https://github.com/rebsmith/boxplot_MRTs.

References

Bandaru, S., Ng, A.H.C., Deb, K., 2017. Data mining methods for knowledge discovery in multi-objective optimization: Part A - Survey. Expert Syst. Appl. 70, 139–159. https://doi.org/10.1016/j.eswa.2016.10.015

Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. Classification and Regression Trees, Wadsworth Statistics/Probability. Chapman and Hall/CRC.

Bryant, B.P., Lempert, R.J., 2010. Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. Technol. Forecast. Soc. Change 77, 34–49.

Cannon, A.J., 2012. Köppen versus the computer: comparing Köppen-Geiger and multivariate regression tree climate classifications in terms of climate homogeneity. Hydrol. Earth Syst. Sci. 16, 217–229. https://doi.org/10.5194/hess-16-217-2012

Chebrolu, S., Abraham, A., Thomas, J.P., 2004. Hybrid Feature Selection for Modeling Intrusion Detection Systems, in: Pal, N.R., Kasabov, N., Mudi, R.K., Pal, S., Parui, S.K. (Eds.), Neural Information Processing, Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 1020–1025.

CSU, 2017. Integrated Water Resources Plan Final Report. Colorado Springs Utilities, Colorado Springs.

Davidson, T.A., Sayer, C.D., Perrow, M., Bramm, M., Jeppesen, E., 2010. The simultaneous inference of zooplanktivorous fish and macrophyte density from sub-fossil cladoceran assemblages: a multivariate regression tree approach. Freshw. Biol. 55, 546–564. https://doi.org/10.1111/j.1365-2427.2008.02124.x

De'Ath, G., 2014. mvpart.

De'Ath, G., 2002. Multivariate regression trees: a new technique for modeling species–environment relationships. Ecology 83, 1105–1117.

Doesken, N., 2014. Colorado Climate Update.

Dudas, C., Frantzén, M., Ng, A.H.C., 2011. A synergy of multi-objective optimization and data mining for the analysis of a flexible flow shop. Robot. Comput.-Integr. Manuf., Conference papers of Flexible Automation and Intelligent Manufacturing 27, 687–695. https://doi.org/10.1016/j.rcim.2010.12.005

Dudas, C., Ng, A.H.C., Pehrsson, L., Boström, H., 2014. Integration of data mining and multi-objective optimisation for decision support in production systems development. Int. J. Comput. Integr. Manuf. 27, 824–839. https://doi.org/10.1080/0951192X.2013.834481

Eker, S., Kwakkel, J.H., 2018. Including robustness considerations in the search phase of Many-Objective Robust Decision Making. Environ. Model. Softw. 105, 201–216. https://doi.org/10.1016/j.envsoft.2018.03.029

Eschner, T., Hadley, R., Crowley, K., 1983. Hydrologic and Morphologic Changes in Channels of the Platte River Basin in Colorado, Wyoming, and Nebraska: A Historical Perspective, in: Hydrologic and Geomorphic Studies of the Platte River Basin, Geological Survey Professional Paper 1277. United States Department of the Interior, Washington.

Gomez-Chova, L., Calpe, J., Soria, E., Camps-Valls, G., Martin, J.D., Moreno, J., 2003. CART-based feature selection of hyperspectral images for crop cover classification, in: Proceedings 2003 International Conference on Image Processing (Cat. No.03CH37429). Presented at the Proceedings 2003 International Conference on Image Processing (Cat. No.03CH37429), pp. III–589. https://doi.org/10.1109/ICIP.2003.1247313

Hadka, D., Reed, P., 2013. Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework. Evol. Comput. 21, 231–259. https://doi.org/doi:10.1162/EVCO_a_00075

Hamann, A., Gylander, T., Chen, P., 2011. Developing seed zones and transfer guidelines with multivariate regression trees. Tree Genet. Genomes 7, 399–408. https://doi.org/10.1007/s11295-010-0341-7

Herman, J.D., Zeff, H.B., Reed, P.M., Characklis, G.W., 2014. Beyond optimality: Multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty. Water Resour. Res. 50, 7692–7713. https://doi.org/10.1002/2014WR015338

Herzschuh, U., Birks, H.J.B., 2010. Evaluating the indicator value of Tibetan pollen taxa for modern vegetation and climate. Rev. Palaeobot. Palynol. 160, 197–208. https://doi.org/10.1016/j.revpalbo.2010.02.016

Hobbs, G.J., 2004. Citizen's Guide to Colorado Water Law, 2nd ed. Colorado Foundation for Water Education, Denver, Colorado.

Jain, A., Zongker, D., 1997. Feature selection: evaluation, application, and small sample performance. IEEE Trans. Pattern Anal. Mach. Intell. 19, 153–158. https://doi.org/10.1109/34.574797

Jones, B., 2014. Communicating Data with Tableau: Designing, Developing, and Delivering Data Visualizations. O'Reilly Media, Inc.

Kasprzyk, J., Nataraj, S., Reed, P., Lempert, R., 2013. Many objective robust decision making for complex environmental systems undergoing change. Environ. Model. Softw. 42, 55–71. https://doi.org/10.1016/j.envsoft.2012.12.007

Kasprzyk, J., Reed, P., Kirsch, B., Characklis, G., 2009. Managing population and drought risks using many-objective water portfolio planning under uncertainty. Water Resour. Res. 45. https://doi.org/10.1029/2009WR008121

Kasprzyk, J.R., Reed, P.M., Characklis, G.W., Kirsch, B.R., 2012. Many-objective de Novo water supply portfolio planning under deep uncertainty. Environ. Model. Softw. 34, 87–104. https://doi.org/10.1016/j.envsoft.2011.04.003

Kollat, J.B., Reed, P.M., 2007. A Framework for Visually Interactive Decision-making and Design using Evolutionary Multiobjective Optimization (VIDEO). Environ. Model. Softw. 22, 1691–1704.

Kwakkel, J.H., Jaxa-Rozen, M., 2016. Improving scenario discovery for handling heterogeneous uncertainties and multinomial classified outcomes. Environ. Model. Softw. 79, 311–321. https://doi.org/10.1016/j.envsoft.2015.11.020

Lall, U., Sharma, A., 1996. A nearest neighbor bootstrap for resampling hydrologic time series. Water Resour. Res. 32, 679–693.

Larsen, D.R., Speckman, P.L., 2004. Multivariate Regression Trees for Analysis of Abundance Data. Biometrics 60, 543–549. https://doi.org/10.1111/j.0006-341X.2004.00202.x

Lawrence, R.L., Wright, A., 2001. Rule-based classification systems using classification and regression tree (CART) analysis. Photogramm. Eng. Remote Sens. 67, 1137–1142.

Liu, H., Motoda, H., Setiono, R., Zhao, Z., 2010. Feature Selection: An Ever Evolving Frontier in Data Mining. JMLR Feature Sel. Data Min. 4–13.

Lukas, J., Barsugli, J., Doesken, N., Rangwala, I., Wolter, K., 2014. Climate Change in Colorado: A Synthesis to Support Water Resources Management and Adaptation. Western Water Assessment, Boulder, CO.

Maier, H.R., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L.S., Cunha, M.C., Dandy, G.C., Gibbs, M.S., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D.P., Vrugt, J.A., Zecchin, A.C., Minsker, B.S., Barbour, E.J., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., Reed, P.M., 2014. Evolutionary algorithms and other metaheuristics in

water resources: Current status, research challenges and future directions. Environ. Model. Softw. 62, 271–299. https://doi.org/10.1016/j.envsoft.2014.09.013
Mathew, S., Varia, J., 2014. Overview of amazon web services. Amaz. Whitepapers.
Matrosov, E.S., Huskova, I., Kasprzyk, J.R., Harou, J.J., Lambert, C., Reed, P.M., 2015. Many-objective optimization and visual analytics reveal key trade-offs for London's water supply. J. Hydrol. 531, Part 3, 1040–1053. https://doi.org/10.1016/j.jhydrol.2015.11.003
Mortazavi, M., Kuczera, G., Cui, L., 2012. Multiobjective optimization of urban water resources: Moving toward more practical solutions. Water Resour. Res. 48. https://doi.org/10.1029/2011WR010866
Murphy, K.P., 2012. Machine Learning: A Probabilistic Perspective. MIT Press.
Nowak, K., Prairie, J., Rajagopalan, B., Lall, U., 2010. A nonparametric stochastic approach for multisite disaggregation of annual to daily streamflow. Water Resour. Res. 46. https://doi.org/10.1029/2009WR008530
Prasad, A.M., Iverson, L.R., Liaw, A., 2006. Newer Classification and Regression Tree Techniques: Bagging and Random Forests for Ecological Prediction. Ecosystems 9, 181–199. https://doi.org/10.1007/s10021-005-0054-1
Questier, F., Put, R., Coomans, D., Walczak, B., Heyden, Y.V., 2005. The use of CART and multivariate regression trees for supervised and unsupervised feature selection. Chemom. Intell. Lab. Syst. 76, 45–54. https://doi.org/10.1016/j.chemolab.2004.09.003
R Core Team, 2016. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
Rajagopalan, B., Nowak, K., Prairie, J., Hoerling, M., Harding, B., Barsugli, J., Ray, A., Udall, B., 2009. Water supply risk on the Colorado River: Can management mitigate? Water Resour. Res. 45. https://doi.org/10.1029/2008WR007652
Ravisankar, P., Ravi, V., Raghava Rao, G., Bose, I., 2011. Detection of financial statement fraud and feature selection using data mining techniques. Decis. Support Syst. 50, 491–500. https://doi.org/10.1016/j.dss.2010.11.006
Reed, P.M., Hadka, D., Herman, J.D., Kasprzyk, J.R., Kollat, J.B., 2013. Evolutionary Multiobjective Optimization in Water Resources: The Past, Present and Future. Adv. Water Resour. 51, 438–456.
Saeys, Y., Inza, I., Larrañaga, P., 2007. A review of feature selection techniques in bioinformatics. Bioinformatics 23, 2507–2517. https://doi.org/10.1093/bioinformatics/btm344
Salonen, J.S., Seppä, H., Luoto, M., Bjune, A.E., Birks, H.J.B., 2012. A North European pollen–climate calibration set: analysing the climatic responses of a biological proxy using novel regression tree methods. Quat. Sci. Rev. 45, 95–110. https://doi.org/10.1016/j.quascirev.2012.05.003
Smith, R., Kasprzyk, J., Basdekas, L., 2018. Experimenting with water supply planning objectives using the Eldorado Utility Planning Model multireservoir testbed. J. Water Resour. Plan. Manag. 144. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000962
Smith, R., Kasprzyk, J., Dilling, L., 2019. Testing the potential of Multiobjective Evolutionary Algorithms (MOEAs) with Colorado water managers. Environ. Model. Softw. 117, 149–163.
Smith, R., Kasprzyk, J., Dilling, L., 2017. Participatory Framework for Assessment and Improvement of Tools (ParFAIT): Increasing the impact and relevance of water

management decision support research. Environ. Model. Softw. 95, 432–446. https://doi.org/10.1016/j.envsoft.2017.05.004
Smith, R., Kasprzyk, J., Zagona, E., 2016. Many-Objective Analysis to Optimize Pumping and Releases in Multireservoir Water Supply Network. J. Water Resour. Plan. Manag. 142. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000576
State of Colorado, 2017. Population Forecasts (2000 - 2050) [WWW Document]. Colo. Demogr. URL https://demography.dola.colorado.gov/population/population-totals-colorado-substate/ (accessed 8.4.17).
Sugimura, K., Obayashi, S., Jeong, S., 2010. Multi-objective optimization and design rule mining for an aerodynamically efficient and stable centrifugal impeller with a vaned diffuser. Eng. Optim. 42, 271–293. https://doi.org/10.1080/03052150903171084
Udall, B., Overpeck, J., 2017. The twenty-first century Colorado River hot drought and implications for the future. Water Resour. Res. 53, 2404–2418. https://doi.org/10.1002/2016WR019638
Verbyla, D.L., 1987. Classification trees: a new discrimination tool. Can. J. For. Res. 17, 1150–1152. https://doi.org/10.1139/x87-177
Watson, A.A., Kasprzyk, J.R., 2017. Incorporating deeply uncertain factors into the many objective search process. Environ. Model. Softw. 89, 159–171. https://doi.org/10.1016/j.envsoft.2016.12.001
Woodbury, M., Baldo, M., Yates, D., Kaatz, L., 2012. Joint Front Range Climate Change Vulnerability Study. Water Research Foundation, Colorado.
Wu, W., Dandy, G.C., Maier, H.R., Maheepala, S., Marchi, A., Mirza, F., 2017. Identification of Optimal Water Supply Portfolios for a Major City. J. Water Resour. Plan. Manag. 143, 05017007. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000811
Zatarain Salazar, J., 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. https://doi.org/10.1016/j.advwatres.2016.04.006
Zheleva, E., Getoor, L., 2009. To Join or Not to Join: The Illusion of Privacy in Social Networks with Mixed Public and Private User Profiles, in: Proceedings of the 18th International Conference on World Wide Web, WWW '09. ACM, New York, NY, USA, pp. 531–540. https://doi.org/10.1145/1526709.1526781