

Evaluating the choice of radial basis functions in multiobjective optimal control applications

Zatarain Salazar, Jazmin; Kwakkel, Jan H.; Witvliet, Mark

DOI

[10.1016/j.envsoft.2023.105889](https://doi.org/10.1016/j.envsoft.2023.105889)

Publication date

2024

Document Version

Final published version

Published in

Environmental Modelling and Software

Citation (APA)

Zatarain Salazar, J., Kwakkel, J. H., & Witvliet, M. (2024). Evaluating the choice of radial basis functions in multiobjective optimal control applications. *Environmental Modelling and Software*, 171, Article 105889. <https://doi.org/10.1016/j.envsoft.2023.105889>

Important note


To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Position Paper

Evaluating the choice of radial basis functions in multiobjective optimal control applications

Jazmin Zatarain Salazar^{a,*}, Jan H. Kwakkel^a, Mark Witvliet^b^a Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, Delft, 2628 BX, The Netherlands^b Axians, Meidoornlaan 24, Delft, 2612TR, The Netherlands

ARTICLE INFO

Keywords:

Direct policy search
 Many Objective Evolutionary Algorithms
 Water resources management
 Global approximators

ABSTRACT

Evolutionary Multi-Objective Direct Policy Search (EMODPS) is a prominent framework for designing control policies in multi-purpose environmental systems, combining direct policy search with multi-objective evolutionary algorithms (MOEAs) to identify Pareto approximate control policies. While EMODPS is effective, the choice of functions within its global approximator networks remains underexplored, despite their potential to significantly influence both solution quality and MOEA performance. This study conducts a rigorous assessment of a suite of Radial Basis Functions (RBFs) as candidates for these networks. We critically evaluate their ability to map system states to control actions, and assess their influence on Pareto efficient control policies. We apply this analysis to two contrasting case studies: the Conowingo Reservoir System, which balances competing water demands including hydropower, environmental flows, urban supply, power plant cooling, and recreation; and The Shallow Lake Problem, where a city navigates the trade-off between environmental and economic objectives when releasing anthropogenic phosphorus. Our findings reveal that the choice of RBF functions substantially impacts model outcomes. In complex scenarios like multi-objective reservoir control, this choice is critical, while in simpler contexts, such as the Shallow Lake Problem, the influence is less pronounced, though distinctive differences emerge in the characteristics of the prescribed control strategies.

1. Introduction

Effective control in environmental systems is essential to meet competing demands for resources and to cope with challenging system states. Evolutionary Multi-Objective Direct Policy Search (EMODPS) has proven to be a flexible and generalizable approach to designing effective operating policies in multi-purpose reservoir management due to its ability to find trade-offs across multiple competing objectives, with heterogeneous, non-linear objective function formulations (Giuliani et al., 2016). EMODPS combines Direct Policy Search (DPS), to parameterize the control policy using global nonlinear approximators, with multi-objective evolutionary algorithms (MOEAs) to find the set of approximate Pareto optimal control policies (Zatarain Salazar et al., 2016; Giuliani et al., 2016). A major benefit of this combination is that the set of Pareto optimal control policies can be attained in a single run (Giuliani et al., 2016) due to the MOEA's population-based search via the use of stochastic search operators (Zatarain Salazar et al., 2016). The suitability of several popular MOEAs within the EMODPS framework has been explored extensively (Zatarain Salazar et al., 2016; Gupta et al., 2020), but the choice of the nonlinear approximators has not received similar scrutiny.

Various nonlinear approximators can in principle be used. A flexible structure that is capable of capturing nonlinear relationships is required, and, so far, RBFs have been successfully applied for multi-purpose reservoir control (Giuliani et al., 2014, 2016; Zatarain Salazar et al., 2016; Gupta et al., 2020; Doering et al., 2021), for carbon mitigation policies (Marangoni et al., 2021), and to control the amount of pollution released into a lake (Quinn et al., 2017).

Various activation functions can be used within an RBF network. Buzoniu et al. (2009) recommends the use of a Gaussian distribution for continuous RBF parameters due to its unbounded support. Giuliani et al. (2014) used Gaussian RBFs for direct policy search within the EMODPS framework. Subsequent studies using EMODPS (e.g., Zatarain Salazar et al., 2016; Marangoni et al., 2021; Quinn et al., 2017) also used the same RBF as suggested by Giuliani et al. (2014). However, the performance of alternative activation functions requires further exploration. Specifically, we aim to understand to what extent the choice of the activation function within the EMODPS framework affects the quality of the Pareto optimal solutions identified.

To this end, we first introduce a set of commonly used activation functions in Section 2, we then delineate the metrics we use to assess the performance of the different functions. Next, in Section 3, we

* Corresponding author.

E-mail address: J.ZatarainSalazar@tudelft.nl (J. Zatarain Salazar).

introduce two distinct cases used to assess the efficacy of the set of activation functions. The first case is the Conowingo reservoir in the Lower Susquehanna River Basin (LSRB) (Zatarain Salazar et al., 2016). The second case is The Shallow Lake Problem (Quinn et al., 2017; Ward et al., 2015; Singh et al., 2015). Section 4 presents the results for both cases focusing on the quality of the final Pareto approximate set as well as the dynamics of the search process. In Section 5, we discuss the results, the trade-offs, performance metrics and policies encountered across each RBF. Finally, Section 6 presents our main conclusions and suggestions for future research.

2. Methods

Consider a basic control problem for which the control policy is described in Eq. (1). Here, u_k is the sum of basis functions defined by the input vector x , and the k th node in the output layer (with $k = 1, \dots, N_u$):

$$u^k = \sum_{i=1}^n w_i^k \phi_i(x_i) \quad (1)$$

where n is the number of RBFs and w_i is the weight of the i th RBF ϕ_i . The weights are non-negative (i.e., $w_i \geq 0 \forall i$) and their sum equals one (i.e., $\sum_{i=1}^n w_i = 1$). This normalization ensures that the weights form a convex combination of RBFs, allowing for a mixture of the various basis functions without amplifying or diminishing the overall output. The $\phi_i(x_i)$ is an activation function that transforms the input vector x in a non-linear manner. In the Susquehanna case, the activation function is represented as:

$$\phi_i(x) = \exp \left[- \sum_{j=1}^m \frac{(x_j - c_{j,i})^2}{b_{j,i}^2} \right] \quad (2)$$

Here, m denotes the number of input variables in vector x , while c_i and b_i are the m -dimensional center and radius vectors for the i th RBF, respectively. To ensure the applicability of this function, the centers of the RBF c_i are constrained to lie within the bounded input space (i.e., $c_i \in [-1, 1]$) and the radii b_i are required to be strictly positive (i.e., $b_i \geq (0, 1]$). The parameter vector θ is therefore composed of the RBFs centers, radii, and the corresponding weights for each output node, structured as $\theta = [c_{i,j}, b_{i,j}, w_i^k]$, with $i = 1, \dots, n, j = 1, \dots, m$, and $k = 1, \dots, N_u$ (Giuliani et al., 2016).

2.1. Activation functions (Kernels)

An activation function is classified as a radial basis function when it satisfies two key properties: it must be semi-positive definite and isotropic (Williams and Rasmussen, 2006). Semi-positive-definite functions have the property that the function of x is greater or equal to zero for all x . Isotropic functions only depend on the difference between $x - x'$ (i.e. the Euclidean distance). In this study, we explore a selection of well-established activation functions that are definite positive and isotropic functions detailed in Table 1 (Fasshauer, 2007; Schaback, 2007; Askari and Adibi, 2015; Zhang et al., 2014).

2.2. Assessment metrics

To evaluate the performance of the different RBF configurations, we analyze the trade-offs in the objective space, and the convergence dynamics of the MOEA (Zitzler et al., 2003) using generational distance, additive ϵ -indicator, hypervolume, epsilon progress, and the archive size.

Generational distance measures the Euclidean distance between the points in an approximation set and their closest counterparts in the reference set. The metric is then computed as the average of these distances. Generational distance is considered to be an easy metric to meet because it often requires that only one solution be close to the reference set to achieve good performance.

The additive epsilon $\bar{\epsilon}$ indicator (Zitzler et al., 2003) assesses the consistency of the approximate Pareto set. That is, the ability to capture all regions of the trade-off space. The metric is calculated as the largest distance that an approximation set must shift in order to dominate the reference set, making it extremely sensitive to gaps in trade-offs (Zatarain Salazar et al., 2016; Reed et al., 2013; Hadka and Reed, 2012).

The hypervolume indicator (Zitzler et al., 2003) provides a measure of convergence and diversity by examining the multidimensional volume attained by each approximation set in relation to a reference set. This metric calculates the difference in hypervolume between the reference set and the Pareto approximation set (Reed et al., 2013).

ϵ -progress is a computationally efficient indication of search progress and stagnation. ϵ -progress occurs when the current solution sits in a different ϵ -box that dominates the previous solution. The ϵ -box divides the objective space into several boxes with the size ϵ . If two solutions reside in the same ϵ -box, the solution closest to the optimal solution will be kept, while the other solution will be eliminated. ϵ -progress thus indicates that the optimizer is able to find solutions in a part of the objective space that was not seen before.

The archive size is the number of non-dominated solutions held by the archive. ϵ MOEAs utilize ϵ values to limit the size of the archive. All solutions that are ϵ -dominated are eliminated. This helps to avoid deterioration, indicating that the ability of the MOEA to find new solutions is diminishing. The final number of non-dominated solutions at the end of all model iterations is used to compute the performance metrics. A larger archive size can comprehensively represent the trade-off space; however, this can only be argued when both convergence and diversity are also high. The size of the archive can also give information about microevolution in different parts of the Pareto front.

3. Case studies

3.1. The Conowingo Reservoir System

The Conowingo Reservoir is an interstate water body shared by the states of Pennsylvania and Maryland. The reservoir needs to meet the demands for hydroelectric power, urban water supply to Chester (PA) and Baltimore (MD), cooling water for the Peach Bottom nuclear power station, and recreation. The downstream releases of the dam are subject to minimum flow requirements, which were set by the Federal Energy Regulatory Commission (FERC) to preserve fishing resources. The Conowingo Dam objectives are modeled over a simulation horizon of one year. This yearly simulation horizon was selected due to the system's limited regulatory capacity and low dependence on the reservoir levels at the start of the simulation (Zatarain Salazar et al., 2016).

Hydropower revenue (maximized): Hydropower revenue is defined as the economic revenue gained from hydropower production at the Conowingo hydropower dam in US\$/MWh defined in Eq. (3). Energy prices are defined by the seven-hour moving average of the energy price trajectory in the Pennsylvania, New Jersey, Maryland (PJM) energy market. The hourly energy production (MWh) is defined by Eq. (4), where η is the turbine efficiency, g is the gravitational acceleration (9.81 m/s^2), γ_w is the water density (1000 kg/m^3), \bar{h}_i is the net hydraulic water level difference (head) in meters and q_i^{Turb} is the turbine flow in m^3/s .








$$J^{hyd} = \sum_{i=1}^H (H P_i \cdot \rho_i) \quad (3)$$

$$H P_i = \eta g \gamma_w \bar{h}_i q_i^{Turb} \cdot 10^{-6} \quad (4)$$

Water supply reliability for the Atomic Power Plant, Chester, and Baltimore (maximized): The daily average volumetric reliability J^{VR} is computed to ensure water supply reliability to the Atomic Power Plant, Chester and Baltimore represented in Eq. (5).

$$J^{VR,i} = \frac{1}{H} \sum_{i=1}^H (Y_i^i / D_i^i) \quad (5)$$

Table 1
Selection of radial basis functions used in this study.

Activation function	$\phi(x)$	Reference	Shape
Modified squared exponential	$\exp(-\frac{(x-x')^2}{\sigma^2})$	Giuliani et al. (2014)	
Squared exponential	$\exp(-\frac{\ x-x'\ ^2}{2\sigma^2})$	Williams and Rasmussen (2006)	
Inverse quadratic	$\frac{1}{1+(\sigma*\ x-x_i\)^2}$	Fasshauer (2007)	
Inverse multiquadric	$\frac{1}{\sqrt{1+(\sigma*\ x-x_i\)^2}}$	Fasshauer (2007)	
Exponential	$\exp(-\frac{\ x-x'\ }{\sigma})$	Fasshauer (2007)	
Matern(3/2)	$(1 + \frac{\sqrt{3}\ x-x_i\ }{\sigma})\exp(-\frac{\sqrt{3}\ x-x_i\ }{\sigma})$	Williams and Rasmussen (2006)	
Matern(5/2)	$(1 + \frac{\sqrt{5}\ x-x_i\ }{\sigma} + \frac{5*\ x-x_i\ ^2}{3\sigma^2})\exp(-\frac{\sqrt{5}\ x-x_i\ }{\sigma})$	Williams and Rasmussen (2006)	

Here, $Y_i^i (m^3)$ denotes the daily delivery, $D_i^i (m^3)$ is the corresponding demand, and i represents the water supply to either Baltimore, Chester, or the Atomic Power Plant.

Environmental Shortage (minimized): This objective aims to minimize the average shortage index relative to the Federal Energy Regulatory Commission (FERC) minimum flow requirements. It is specified as the daily average shortage index and is calculated using the following equation:

$$J^{SI} = \frac{1}{H} \sum_{t=1}^H \left(\frac{\max(Z_t - Y_t, 0)}{Z_t} \right)^2 \tag{6}$$

Here, $Y_t (m^3)$ is the daily release and $Z_t (m^3)$ is the corresponding FERC flow requirement. The quadratic formulation is intended to penalize substantial deficits in a single time step while allowing for more frequent, minor shortages (Hashimoto et al., 1982).

Recreation (maximized). This objective aims to maximize the reliability of storage on weekends during the peak tourist season. It is quantified by evaluating the proportion of weekend days in the peak season when the water level is maintained at or above a target level, necessary for enabling various recreational activities. The target water level for recreational activities is set at 32.5 m (106.5 ft).

The storage reliability for recreation, J^{SR} , is computed using the following equation:

$$J^{SR} = 1 - \frac{n_F}{2N_{we}} \tag{7}$$

Here n_F denotes the number of weekend days in the peak season where the water level falls below the intended target level of 32.5 m (106.5 ft), and N_{we} is the total number of weekends in the tourist season.

3.2. The Shallow Lake Problem

The Shallow Lake Problem is a stylized decision problem in which a town must decide the amount of pollution to release into a nearby

shallow lake over time. This hypothetical problem involves two sources of pollution: anthropogenic pollution generated by the town through industrial and agricultural waste, and natural inflows that are uncontrollable and come from the environment. There is also a natural outflow process based on the capability of the lake to recycle resources that are capable of naturally reducing pollution over time (Carpenter et al., 1999; McInerney et al., 2012; Hadka et al., 2015).

Pollution levels are determined through Eq. (8), where X represents the concentration of pollution in the lake, a is the anthropogenic pollution input for the time period, Y refers to the natural inflow of pollution which is described using a lognormal distribution, q refers to the rate at which pollution is recycled into the lake's sediment, and b refers to the loss of pollution from the lake through natural outflows. The exact specification for each of the parameters is based on the lake model developed by Quinn et al. (2017).

$$X_{t+1} = X_t + a_t + Y_t + \frac{X_t^q}{1 + X_t^q} - bX_t \tag{8}$$

The behavior of the lake problem exhibits a tipping point. If the critical threshold of pollution concentration is surpassed, the lake's state transitions towards a eutrophic equilibrium, making it impossible to return to a healthier oligotrophic equilibrium without active human intervention to reduce the pollution in the lake. Here, we use the direct policy search variant of The Shallow Lake Problem as proposed by Quinn et al. (2017). In this setup, the aim is to find a state-based closed control loop where the anthropogenic release at a given time step is based on the observed pollution level of the lake in the previous time step.

There are four conflicting objectives: to minimize the maximum pollution level while maximizing the utility of the release policy to the town, the reliability of the policy, and policy inertia. The multi-objective formulation of this problem was introduced by Singh et al. (2015) and further developed by Ward et al. (2015), with the goal of introducing objectives that exemplify the conflicts that occur with

a diverse group of decision-makers and a problem characterized by both stochastic uncertainty (*that is*, the stochastic natural inflow) and deep uncertainty. To address the stochastic uncertainty, the model is executed for N stochastic realizations, and descriptive statistics are computed based on these replications.

Maximum Pollution (minimized): Decision makers, such as environmental regulators, are seeking to ensure that the maximum pollution level reached in the lake is kept as low as possible (Singh et al., 2015).

$$f_{max\ pollution} = \max_{t \in \{1, \dots, T\}} \frac{1}{N} \sum_{n=1}^N X_{t,n} \quad (9)$$

where $X_{t,n}$ is the concentration of the pollution in year t for stochastic realization n .

Reliability (maximized): Reliability captures the desire of decision-makers to keep the lake below the critical pollution threshold. At the same time, in contrast to the maximum pollution objective, a high-reliability policy also accepts a small amount of pollution, as long as it remains below the critical threshold (Singh et al., 2015). The reliability of a policy is the average reliability for each time step over all realizations N , shown in Eq. (10) (Ward et al., 2015).

$$f_{reliability} = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{T} \sum_{t \in T} \theta_{t,n} \right), \text{ where } \theta_{t,n} = \begin{cases} 1 & X_{t,n} < P_{crit} \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Utility (maximized): In contrast to the objectives that align with the goals of environmental regulators, utility represents the interests of the town's agriculture and industry. The aim here is to maximize the utility of a policy for these stakeholders. Here, α is the utility generated by one unit of anthropogenic pollution, while δ is the discount rate. This objective naturally conflicts with the objective of minimizing the pollution level in the lake, providing a valuable dynamic for robust decision support analysis (Ward et al., 2015).

$$f_{utility} = \frac{1}{N} \sum_{n=1}^N \left(\sum_{t \in T} \alpha a_{t,n} \delta^t \right) \quad (11)$$

Inertia (maximized): This objective captures the undesirability of large year-over-year changes to the anthropogenic pollution input. The aim is to maximize the average inertia of a policy. Similar to the utility objective, the inertia of a policy is first calculated for every time step involved, and across different stochastic realizations. Next, the mean of the resulting vector of values is used to determine inertia-based robustness. Inertia for a single time step in an experiment is determined with Eq. (12).

$$f_{inertia} = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{T} \sum_{t \in T} \phi_{t,n} \right), \text{ where } \phi_{t,n} = \begin{cases} 1 & |a_{t,n} - a_{t-1,n}| < 0.01 \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

3.3. Experimental setup

Our goal is to understand how the shape of the RBF affects the performance of the trade-offs and the convergence dynamics of the MOEA when using EMOODPS. In our experimental configuration, we adhere to the recommendations presented by Giuliani et al. (2016) for the RBF network design. As per their guidelines, the number of Radial Basis Functions (RBFs) is determined by adding two to the number of inputs. Specifically, for the Conowingo Reservoir System, which involves two inputs (namely, time of year and storage level), we employ a total of four RBFs, calculated as the number of inputs (2) plus two. On the other hand, in the case of The Shallow Lake Problem, which has a single input (pollution concentration in the lake), we utilize three RBFs, obtained by adding two to the number of inputs (1). The centers, radii, and weights were searched by the MOEA with $\mathbf{c}_i \in [-1, 1]$, $\mathbf{b}_i \in [0, 1]$, and $\mathbf{w}_i \in [0, 1]$. These ranges were preserved across all the tested activation functions in Table 1. For both cases, each RBF shape

was optimized using 10 random seeds to account for variability in the initial population. After preliminary assessment for convergence, and based on previous studies (see Zatarain Salazar et al. (2016) and Quinn et al. (2017)), for the Conowingo Reservoir System, we used 250 k function evaluations per seed, while for The Shallow Lake Problem we used 100 k function evaluations for each seed. For the MOEA, we used ϵ -NSGA2 (Kollat and Reed, 2005, 2006). This is a population-based MOEA that uses ϵ -archiving (see Section 2.2). It has shown robust performance attributed to its utilization of the simulated binary crossover operator. Moreover, it has been effective in maintaining diversity and keeping a bounded search across a large number of objectives due to its epsilon-dominance archiving mechanism (Zatarain Salazar et al., 2016). It is readily available in various software packages and is easy to parallelize since the evaluation of a given population is embarrassingly parallel, where the function evaluations are distributed and do not interact with each other.

4. Results

We evaluated the performance of each of the activation functions specified in Table 1 through a visual analysis of the trade-offs and through multi-objective performance metrics that indicate their ability to converge and diversify.

4.1. Trade-offs and release strategies for the Conowingo Reservoir System

Fig. 1 shows the trade-offs attained by the different activation functions. Each RBF configuration is depicted by a different color, and each axis contains the objective values, where the preferred solutions lie at the top of each axis. If two lines cross, this indicates that a trade-off was encountered.

Inspecting the trade-offs for each RBF configuration, we note that the modified squared exponential RBF (Fig. 1a), the squared exponential RBF (Fig. 1b), and the inverse multiquadric RBF (Fig. 1d) attain a diverse solution space with high performance across each objective. This is reflected by the lines reaching the upper bounds of each objective axis, indicating near-optimal solutions for those objectives. Conversely, the inverse quadratic RBF (Fig. 1c) performs slightly worse, yet it still finds a diverse set of solutions. In contrast, the Exponential RBF (Fig. 1e), Matern 3/2 (Fig. 1f), and Matern 5/2 kernel (Fig. 1g), find a narrow trade-off space, reflecting limited diversity of the approximation set. Fig. 2 illustrates the water release strategies designed to satisfy the water demands for the atomic power plant, for the cities of Baltimore and Chester, and to meet environmental flow requirements. The plots show daily releases within the annual planning horizon. Fig. 2 shows that the choice of RBFs substantially impacts both the objective space and the attributes of the release strategies. In this context, the intra-annual variability is particularly notable within the release policies.

In panels (a), (b), and (d), we observe release policies that consistently meet the water demands throughout the entire year. Specifically, there is a close alignment between the projected demands and the actual water releases, indicating effective water management. An exception to this tight alignment is noticeable in the case of Baltimore, where a slight gap between the city's water demand and the actual release can be observed.

This divergence becomes more pronounced in panel (c), where, towards the end of the year, we witness a widening gap between the water demands of both Baltimore and Chester and the actual releases. This suggests a potential shortfall in water supply for both cities.

Furthermore, panels (e) through (g) show more dramatic discrepancies within the intra-annual release dynamics. These panels reveal considerable gaps between demand and release at various points throughout the year. In the context of reservoir management, these fluctuations could have substantial implications, particularly for objectives that are sensitive to the timing of water releases. For example, significant gaps in releases during critical periods could compromise environmental health or the reliable supply of water to urban centers.

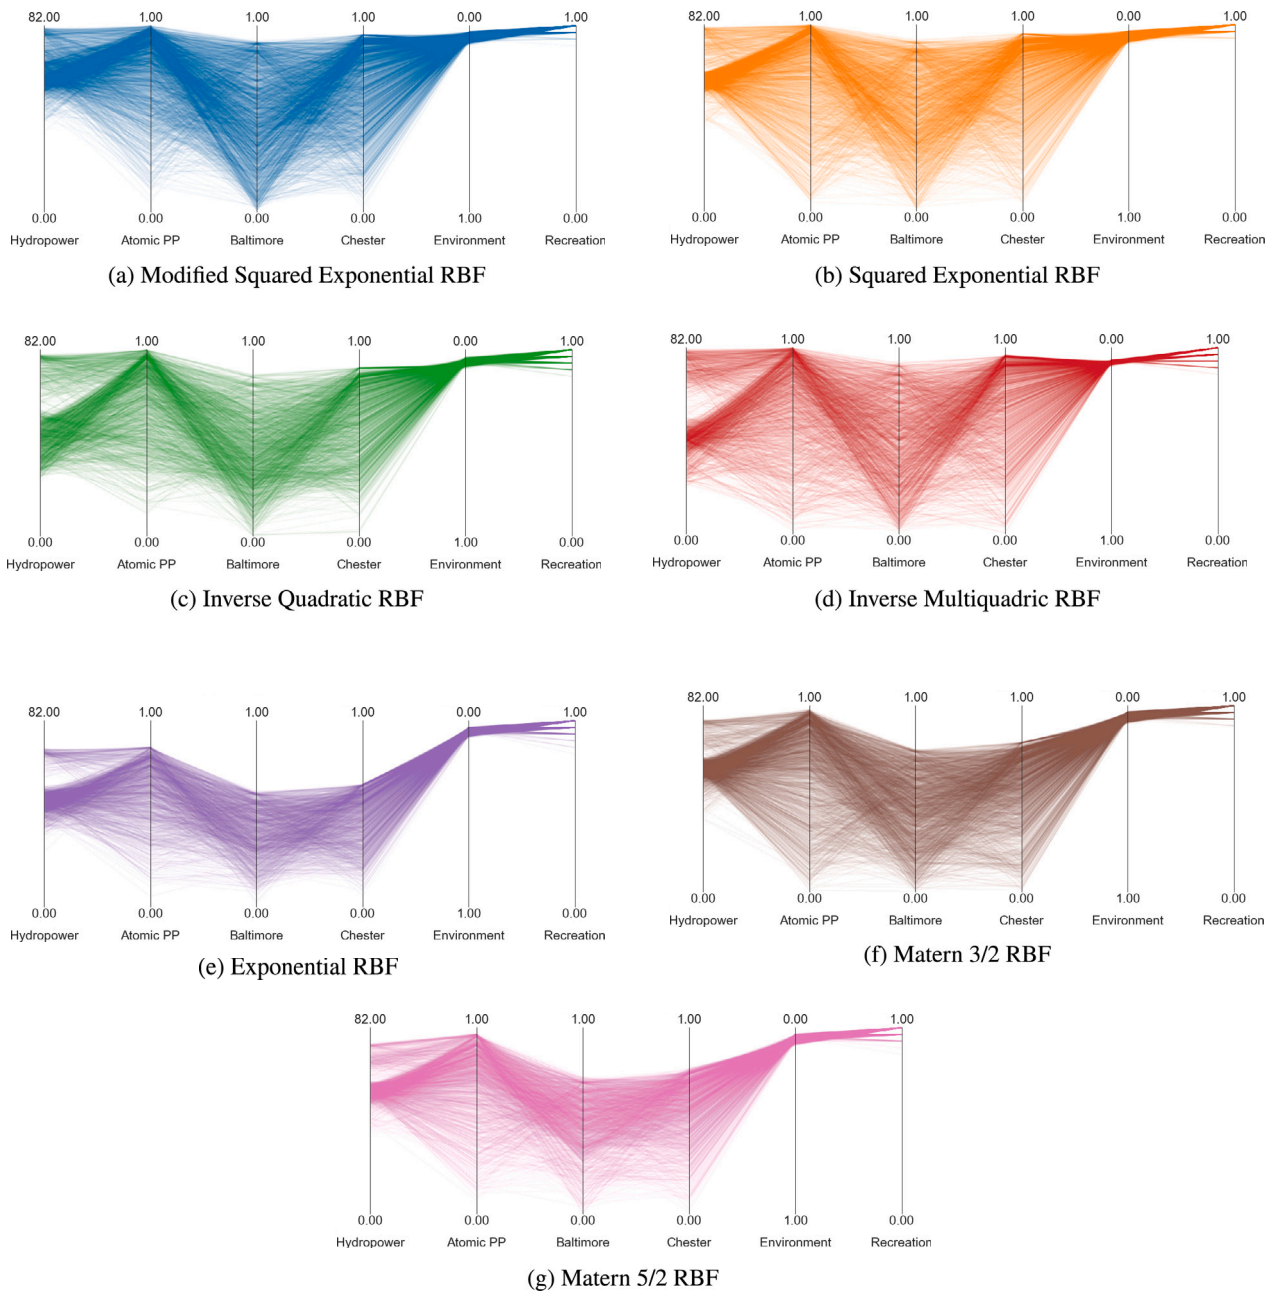


Fig. 1. Trade-offs in the Conowingo Reservoir System attained by each RBF configuration: In this plot, each axis represents a distinct objective, and each colored line corresponds to a unique policy. The direction of preference is upwards, indicating that a more desirable policy would be represented as a straight line situated at the top of the axes.

4.1.1. Performance across objectives for the Conowingo Reservoir System

In Fig. 3, we inspect in detail the performance across each of the LSRB objectives. Each panel shows the objective distributions attained by different activation functions. The boxplots show consistent results with the parallel coordinate plots in Fig. 1, with the benefit of detecting detailed differences across each objective. To supplement the results of Fig. 3, the detailed statistics can be found in the appendix, in Table 4 for the Susquehanna objectives and in Table 5 for the Lake problem objectives.

We observe that the modified squared exponential and the squared exponential RBFs attain the highest median values for the hydropower objective (panel a), with a difference of approximately 10 M USD/year reflected in their interquartile range, with several low-performing outliers. In contrast, the inverse quadratic and inverse multiquadric RBFs exhibit a median value of approximately 45 M USD/year, with an interquartile range difference of approximately 30 M USD/year.

Nevertheless, their whiskers indicate that these functions span the entire range of hydropower values, with this being particularly true for the inverse multiquadric RBF. Conversely, the exponential and Matern 3/2 RBFs reveal tightly clustered values around 45–50 and 50–55 M USD/year, respectively. The Matern 3/2 and Matern 5/2 achieve similar median values, but the Matern 5/2 RBF demonstrates a higher upper quartile.

Moving on to panel b, all the activation functions, except for the exponential RBF, achieved a median value higher than 75% reliability for the atomic power plant. It is worth noting that the squared exponential variants, the inverse quadratic and inverse multiquadric RBFs found solutions close to 100% reliability for this objective. The exponential RBF displayed the lowest performance with a median value at 60%.

Next, panel c shows the reliability results for water supply to Baltimore. We observe that the modified squared exponential and squared exponential RBFs show high performance. However, the interquartile

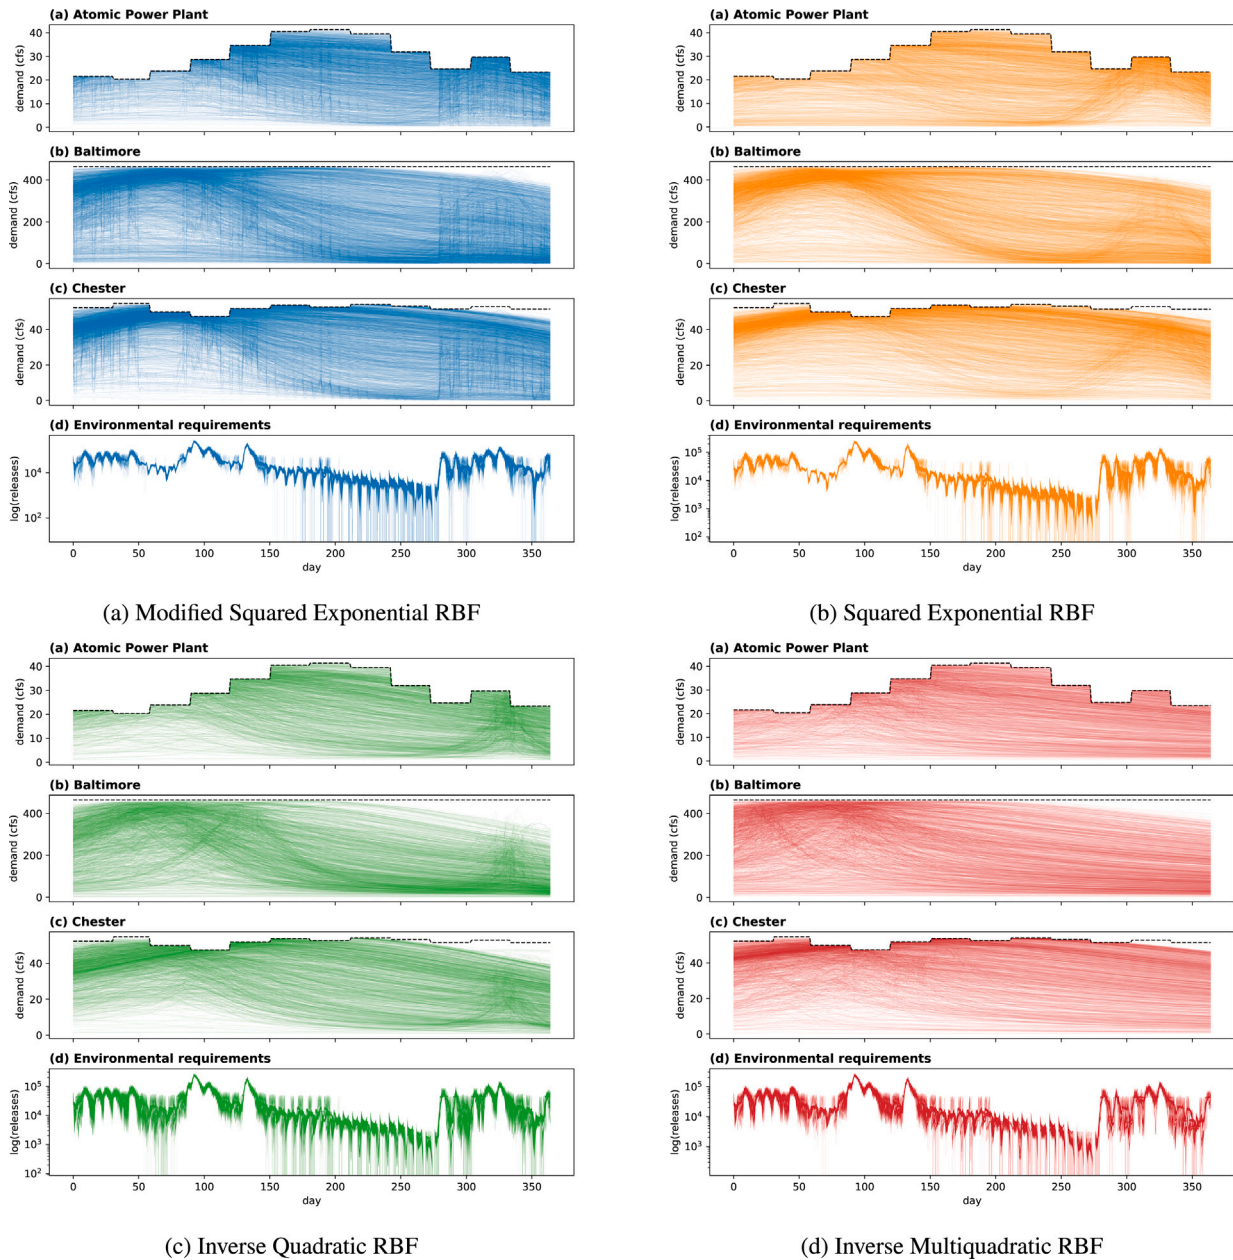


Fig. 2. Release policies for the Conowingo Reservoir System: In this figure, each subplot represents a distinct RBF configuration. From top to bottom, the subplots indicate the releases for the Atomic Power Plant, Baltimore, Chester, and the Environmental requirements. The lines within each subplot represent the release trajectory over a one-year period, with values expressed in cubic feet per second (cfs). The bottom plot across all subplots shows the log-transformed releases for the environmental requirements.

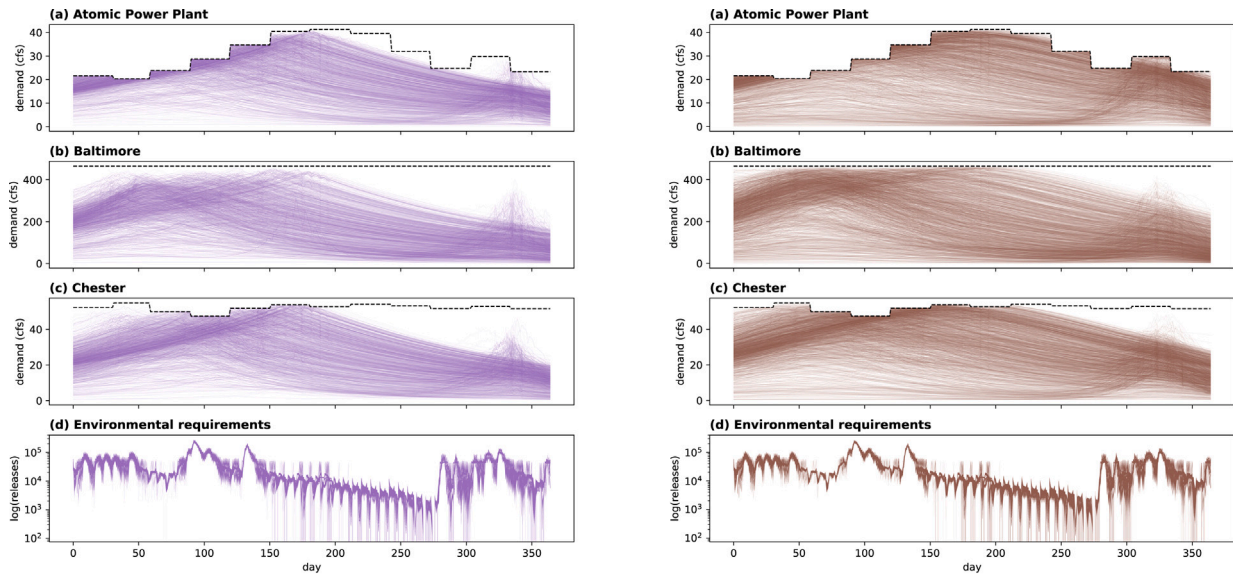
range for the modified squared exponential is broader in this case. The inverse multiquadratic and the exponential RBFs attain the lowest median reliability for Baltimore at roughly 40%. Remarkably, the former finds solutions across the entire reliability range, including the absolute highest value for this objective. This stands in contrast to the exponential RBF which has a tight boxplot and performance ranging between 0.15 and 0.62. The Matern 3/2 and Matern 5/2 attain median values around 0.5, with both whiskers skewed towards lower reliabilities.

Panel d details the reliability results for water supply to Chester. Here, the modified squared exponential and the squared exponential RBFs maintain their lead reflected by the median values, upper quartile ranges, and highest attained values. These two configurations are followed by the inverse multiquadratic RBF with a slightly lower median value and lower 25% score values. Once more, the exponential RBF is

outperformed by all the other tested configurations for water supply reliability.

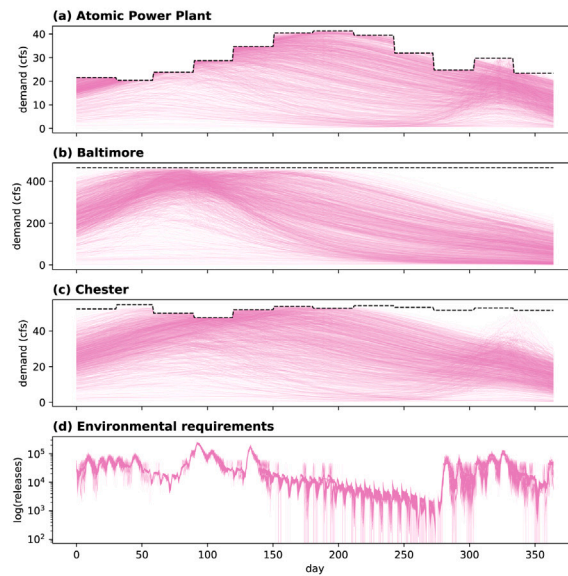
Panel e shows the environmental objective values. The modified squared exponential, and the squared exponential RBFs find the largest range for the environmental objective. Interestingly, the Matern 3/2 and Matern 5/2 find a large range of solutions for this objective, where the Matern 3/2 outperforms the squared exponential variants, which attained consistently high performance across the other objectives. This divergence could be explained by the stark trade-off between the environmental objective and the other objectives highlighted in Section 4.1.

Finally, in panel f, all the RBFs attain high scores for the recreation objective, with the exception of the inverse quadratic and the exponential RBF, most of the low scores for the other RBFs are outliers.



(e) Exponential RBF

(f) Matern 32 RBF



(g) Matern 52 RBF

Fig. 2. (continued).

4.1.2. Performance metrics for the Conowingo Reservoir System

Fig. 4 shows the performance metrics for each RBF. The metrics were computed relative to a global reference set generated from the non-dominated solutions across the 7 activation functions over 10 random seed trials. The x -axis shows the number of function evaluations (nfe). Each row shows a different metric value, and each column shows a different RBF. Finally, the colored plots depict the runtime dynamics for each configuration.

All the activation functions reach a comparable generational distance to that of the global reference set within the first 25,000 nfe. In this case, a low metric value is desired as it indicates the average distance between the global reference set and the Pareto approximation. Generational distance is used in this study mainly to detect absolute failure in the optimization.

The additive epsilon indicator measures gaps in the Pareto front, hence, it is a harder metric to meet than generational distance. Similarly to generational distance, this metric is computed relative to a

global reference set, and a low value is desired as it measures the distance that an approximation set needs to be translated in order to dominate the global reference set. We notice that the different random seed trials for the modified squared exponential RBF start to stabilize before 25,000 nfe. Further, some configurations are more sensitive to the random seed trials (*i.e.*, the initial population). This is particularly stark for the exponential RBF and the Matern 5/2 RBF, where there is a clear split between high and low seed performance. Conversely, the squared exponential, inverse quadratic, and inverse multi-quadratic RBF show less random seed variability but result in epsilon values far from the global reference set.

Hypervolume provides the most thorough measure of convergence and diversity, making it the hardest metric to meet. Hypervolume measures the volume in a multi-dimensional space attained by an approximation set relative to the global reference set. Hence, a high metric value is desired. The runtime dynamics show that the squared

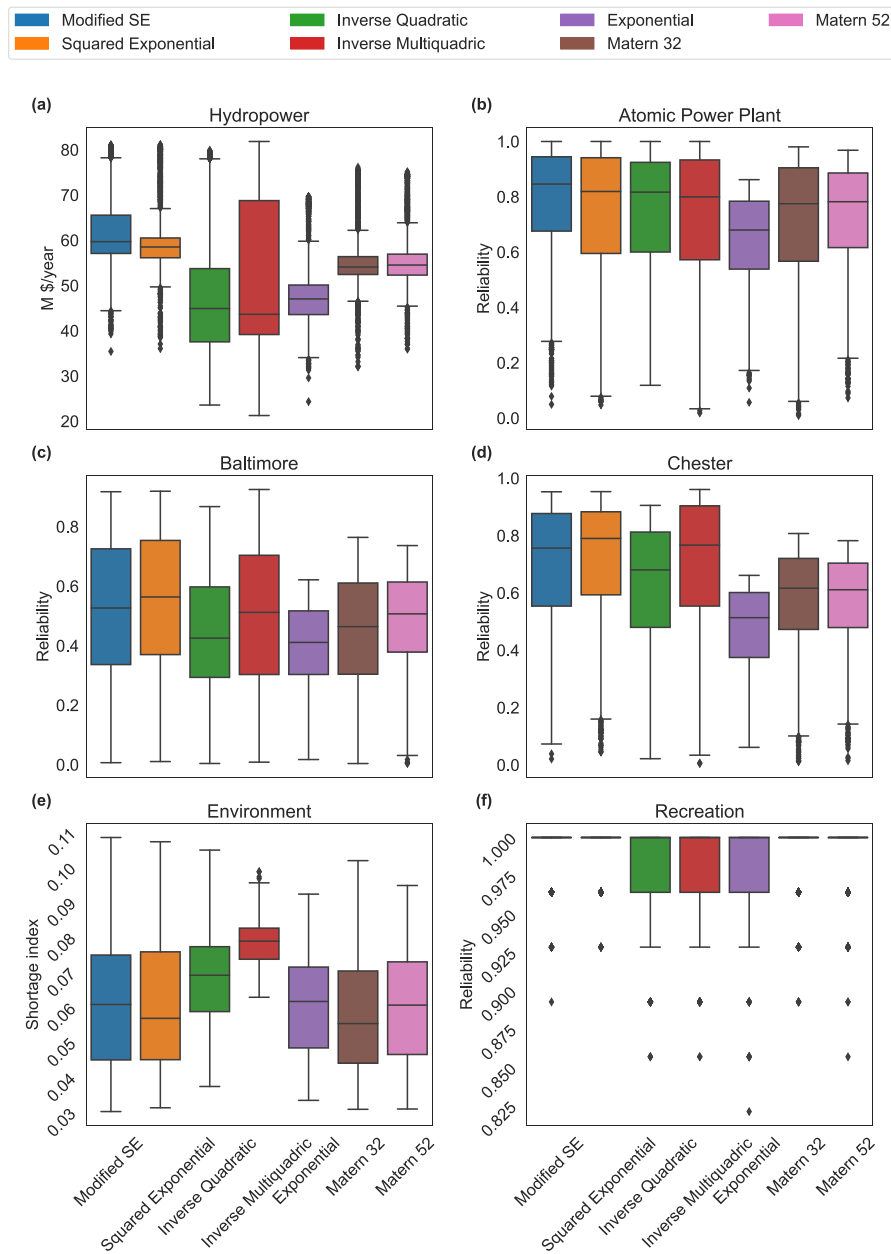


Fig. 3. Comparison of objective values for the Conowingo Reservoir System per RBF. The figure illustrates the distribution of objective values for various RBF configurations. Each subplot represents a distinct objective: (a) through (d) and (f) feature upward preferences, corresponding to higher hydropower revenue (Panel A) and reliability values. In contrast, for Panel (e), a downward preference is desired, indicating the necessity for a smaller shortage index. Within each subplot, individual box plots the objectives' statistics obtained using different RBF configurations.

Table 2
Average performance metrics and reference set contribution per RBF.

Activation function	Non-dominated solutions	Mean generational distance	Mean epsilon indicator	Mean hypervolume	Set contribution
Modified SE	2675	0.022556	0.349520	0.716880	0.743
Squared Exponential	2142	0.020800	0.411422	0.641916	0.183
Adapted Inverse Quadratic	1684	0.030182	0.472414	0.333763	0.020
Adapted Inverse Multiquadric	1636	0.023650	0.582708	0.350160	0.052
Exponential	1585	0.042522	0.363626	0.151511	0.000
Matern 3/2	2268	0.032786	0.356777	0.374623	0.001
Matern 5/2	1774	0.031683	0.369322	0.334237	0.002

exponential variants (in blue and orange) attain the highest hypervolume. In contrast, the inverse quadratic, the inverse multiquadric and the exponential RBFs achieve hypervolume values far from the global reference set. Interestingly, all the RBFs, with the exception of the squared exponential variants, seem to stabilize after 25,000 nfe and do not show signs of further hypervolume improvements beyond 100,000 nfe. These results are contrary to the trends observed for the modified and the squared exponential RBFs, which may indicate that further hypervolume improvements are possible by continuing exploration with a larger number of function evaluations. These two RBF configurations could be suitable in combination with asynchronous evolutionary optimization to capitalize on parallel function evaluations.

Epsilon progress indicates the ability to escape local optima and to find continued improvements to the non-dominated archive. Specifically, the epsilon value indicates the user-specified threshold for which the search algorithm needs to produce at least one solution above this threshold at a certain frequency to avoid stagnation. To this effect, the search dynamics for the modified squared exponential RBF show practically a linear relationship between nfe and epsilon progress. In contrast, the squared exponential RBF shows milder progress throughout the search and a wider performance range across random seed trials. This is also true for the exponential, for the Matern 3/2, and for the Matern 5/2 RBFs, whose high-performing seeds show fast epsilon progress, whereas the seeds that perform poorly flatten quickly, indicating stagnation, or infrequent or no improvements to the archive.

The archive contains the population of non-dominated solutions. The size is adapted based on the epsilon-dominance criterion, which guarantees simultaneous diversity and convergence of the set. This explains the similar trends observed between hypervolume and archive sizes. In essence, a large archive size can contribute to the diversity of the approximation sets. In these results, the modified squared exponential and squared exponential RBFs contain between 500–1000 members in the archive at the end of the run, depending on the seed. These two RBFs also achieved the highest hypervolume values at the end of the run. Evidently, the archive sizing is highly dependent on the starting populations reflected by large random seed variability across all the activation functions.

Overall, this analysis suggests that the modified squared exponential and the squared exponential RBFs are able to escape local optima, while the other configurations either suffer stagnation or have two modes of performance between seeds; some are able to make epsilon progress while others get stuck in local optima.

Table 2 shows an overview of the statistical differences observed among the assessed RBFs. The table presents the averaged outcomes across various seed runs, encompassing both the maximum and minimum metric values illustrated in Fig. 4. Furthermore, it outlines the individual contributions of each RBF to the global reference set. The global reference set refers to a collection of non-dominated solutions, and the contribution of each activation function to this set represents the proportion of these non-dominated solutions that were produced using that specific activation function. As highlighted in the table, the modified squared exponential RBF was the most influential activation function, contributing significantly to the global reference set with a substantial 56%. Following this, the squared exponential RBF also

showed a notable contribution of 28%. In contrast, the exponential RBF exhibited minimal impact, contributing only 0.1% to the global reference set.

4.2. Trade-offs and emission policies for the Shallow Lake Problem

Fig. 5 shows the parallel coordinate plots for each of the RBFs for The Shallow Lake Problem. Again, each RBF configuration is depicted by a different color, and each axis contains the objective values, where the preferred solutions lie at the top of each axis. If two lines cross, this indicates that a trade-off was encountered. Fig. 7 shows the boxplots across the RBFs for each objective.

Looking across the results for the different RBFs, we see largely the same trade-offs and similar ranges for the individual outcomes of interest. Only on the inertia objective do we see a slight deviation where the Matern 5/2 RBF is able to find a solution with high inertia and higher reliability, while the other RBFs generally have higher inertia at the expense of reliability. The large similarity across RBFs can also be seen in Fig. 7, where the boxplots and median values are largely the same across the different RBFs for each outcome of interest. Only for inertia do we see some small differences in the outliers. Examining Figs. 5 and 7 together, we observe that, for The Shallow Lake Problem, there is no major difference in performance across the various RBFs. However, a closer inspection of the emission strategies prescribed by each configuration, as depicted in Fig. 6, reveals stark differences in phosphorus release decisions.

Specifically, the choice of RBF significantly influences the release decisions. Some configurations lead to decreasing phosphorus discharges at much lower lake phosphorus concentrations, strategically reducing pollution to preserve lake health. In contrast, other configurations may increase phosphorus releases under similar conditions to maximize further economic benefits. Yet, there are also configurations that maintain a constant phosphorus release, regardless of the lake's condition.

This scenario exemplifies the concept of equifinality. In this context, different pollution release policies can achieve similar trade-offs as observed in Fig. 5 despite their distinct approaches to managing phosphorus releases in Fig. 6.

4.2.1. Performance metrics for the Shallow Lake Problem

Fig. 8 presents performance metrics for each of the RBFs, following the same setup as Fig. 4. In terms of generational distance, there is swift convergence observed across all RBFs. Conversely, the additive epsilon indicator exhibits a somewhat erratic pattern across the RBFs, indicating gaps within the Pareto fronts of individual seeds relative to the reference set.

For the hypervolume indicator, the modified squared exponential, inverse multiquadric, exponential, and Matern 5/2 RBFs display a distinct outlier seed with relatively low hypervolume, indicating stagnation in their search. The inverse quadratic and Matern 3/2 RBFs also exhibit a similar outlier seed, but this seed starts to make progress again after approximately 50 k nfe and eventually converges to a similar hypervolume as the other seeds.

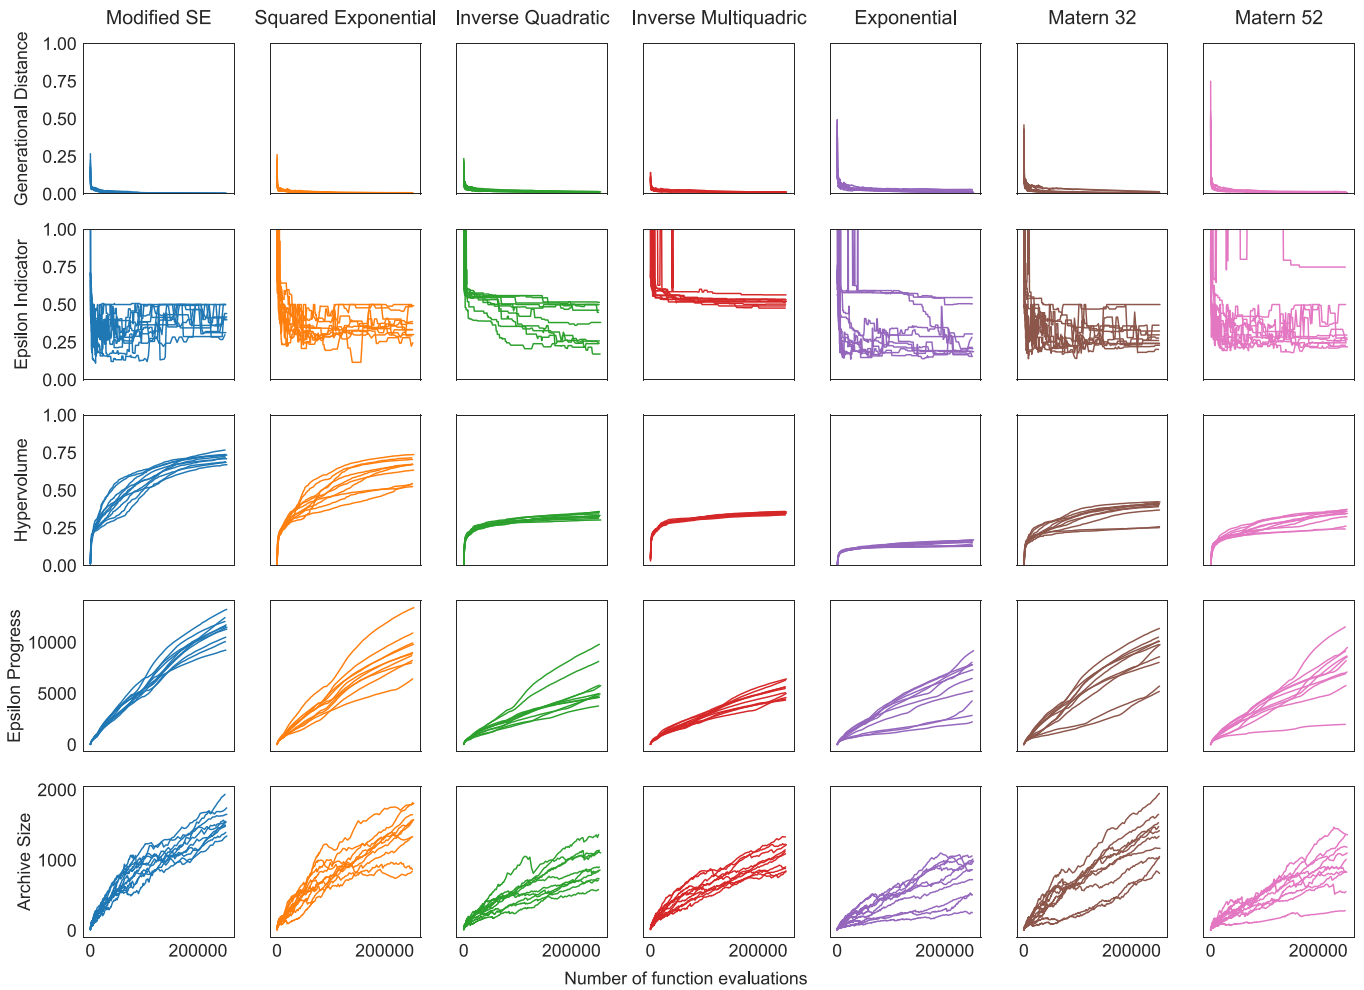


Fig. 4. Comparison of performance metrics for the Conowingo Reservoir system. This figure displays the search dynamics throughout the duration of the run. Each column, distinguished by color, represents a unique RBF configuration, while the rows depict the metrics used for evaluation. A smaller generational distance and epsilon indicator are preferable, as they signify proximity to the reference set and minimal translation distance, respectively. Conversely, a high hypervolume is desirable, as it indicates the volume dominated by the approximation set relative to the reference set—it entails a comprehensive measure of both diversity and convergence. A consistent epsilon progress, represented by a steep slope, is desired. The archive size is closely related to the diversity of the solution set, and this relationship is particularly noteworthy when the hypervolume exhibits high performance.

On the other hand, the remaining RBFs all converge to a nearly identical hypervolume, which is also relatively close to the hypervolume of the reference set. The epsilon progress indicator suggests a phase of steady and then stabilizing progress, while the archive size indicates outlier seeds with larger archives. Similarly, the inverse quadratic and Matern 3/2 RBFs display a similar outlier, which, at around 50k nfe, begins to approach the average size across the seeds, indicating a connection between larger archives and stalled search phases.

Table 3 gives an overview of some summary statistics across the different RBFs for the different performance metrics. Again it is notable how similar the results are across RBFs. The archive size is basically the same across all RBFs. There are some small differences for the other indicators. The squared exponential RBF, followed by the modified squared exponential are the main contributors to the reference set. Conversely, the inverse multiquadric and the Matern 5/2 do not contribute to the reference set, despite having similar average metric values relative to the other RBFs.

4.3. Kruskal–Wallis test to detect statistical differences for the objectives values


A Kruskal–Wallis test was performed to assess whether the objective values obtained for each RBF are statistically different. This

non-parametric statistical test compares more than two independent groups when the data is not normally distributed. The Kruskal–Wallis test is therefore suitable when the data does not meet the normality or equal variance assumptions of parametric tests like one-way ANOVA. It is robust to non-normality, does not assume equal variances, and can handle ordinal or continuous data, widely used in other diagnostics studies (see Hadka and Reed (2012), Gupta et al. (2020)). For the Kruskal–Wallis test, each data point in the samples is ranked, and the test measures whether the observed differences in these ranks are statistically significant. The hypotheses for the Kruskal–Wallis test are defined as follows:

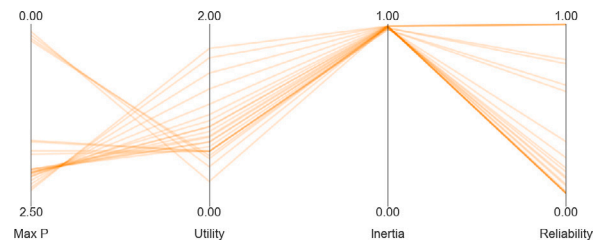
- Null Hypothesis (H_0): There is no difference between the groups (i.e., the medians are equal across all groups). The null hypothesis is accepted if the calculated χ^2 is less than 3.84146.
- Alternative Hypothesis (H_1): There is a difference between the groups (i.e., at least one group median is different from the others). The alternative hypothesis is accepted if the calculated χ^2 is greater than 3.84146.

For this test:


Significance Level (α) = 0.05



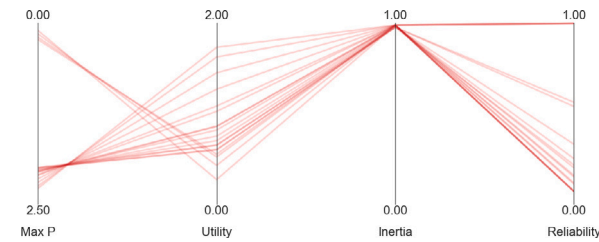
(a) Modified Squared Exponential RBF




(b) Squared Exponential RBF




(c) Inverse Quadratic RBF




(d) Inverse Multiquadric RBF



(e) Exponential RBF



(f) Matern 3/2 RBF



(g) Matern 5/2 RBF

Fig. 5. Trade-offs for The Shallow Lake Problem attained by each RBF configuration. Each axis represents a distinct objective, and each colored line corresponds to a unique policy. The direction of preference is upwards, indicating that a more desirable policy would be represented as a straight line situated at the top of the axes.

Table 3
Average performance metrics and reference set contribution for the shallow lake problem.

Activation function	Non-dominated solutions	Mean generational distance	Mean epsilon indicator	Mean hypervolume	Set contribution
Modified SE	21	0.018994	0.242975	0.468240	0.250
Squared Exponential	23	0.020191	0.253429	0.471526	0.542
Inverse Quadratic	19	0.022929	0.271876	0.464444	–
Inverse Multiquadric	19	0.025915	0.259099	0.463475	0.042
Exponential	23	0.021459	0.285110	0.465515	0.083
Matern 3/2	20	0.015791	0.318811	0.467675	0.083
Matern 5/2	22	0.020485	0.268411	0.466675	–

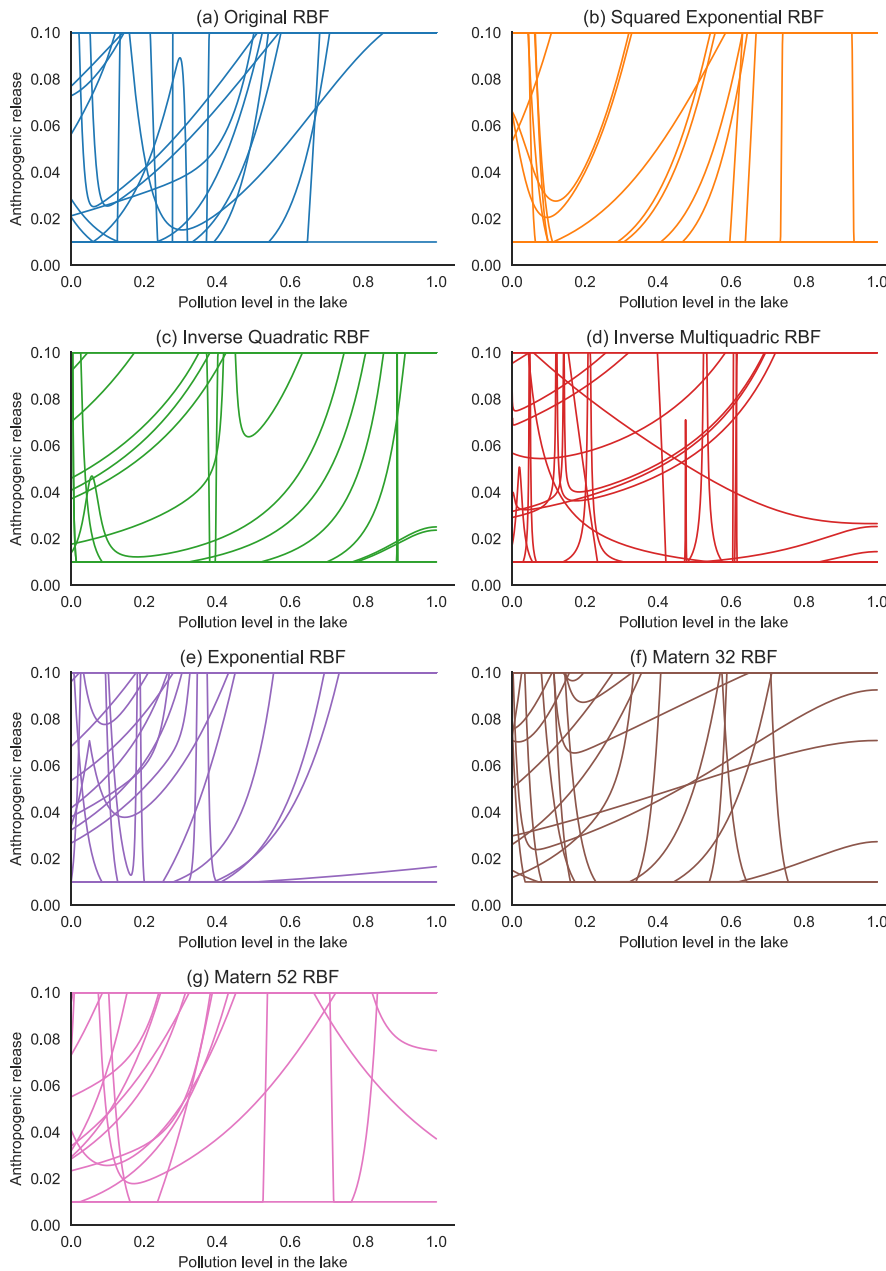


Fig. 6. Phosphorus emission policies for the Shallow Lake Problem: This figure illustrates the various policy approaches prescribed by each RBF tested for managing phosphorus emissions. The intuition is that each policy begins with a default maximum phosphorus release. As the phosphorus concentration in the lake rises, the policies generally dictate a reduction in anthropogenic emissions. However, upon crossing a critical tipping point once the lake’s health can no longer recover, several policies advocate for a sharp increase in emissions to maximize economic benefits. The depicted rules vary in how they adjust anthropogenic releases based on the state, with distinct patterns on how they increase emissions.

Degrees of Freedom = 1 (in this case, 2 samples tested against each other)

$$\chi^2 \text{ (Critical Value)} = 3.84146$$

For each score with a χ^2 value greater than 3.84146, the null hypothesis is rejected. This indicates that the sets are statistically different at a significance level of $\alpha = 0.05$. If the p -value is less than or equal to this significance level, we conclude that not all the group medians are equal.

Kruskal–Wallis test for the Conowingo Reservoir System

Fig. 9 visually represents these pairwise comparisons for different objectives in the case of the Conowingo Reservoir System. In this figure, the colorbar indicates the p -value, with dark purple grids signifying

cases where the sets are statistically different. In the Susquehanna case, we observe that the majority of the pairwise comparisons exhibit p -values below the 0.05 threshold, indicating significant differences, this is visually denoted by dark purple grids. This statistically supports our findings from the main study that the choice of RBF significantly impacts the objective values attained for the Conowingo Reservoir System.

Kruskal–Wallis pairwise comparison of objective values across RBF configurations for the Shallow Lake Problem

A similar test was conducted for the Lake Problem, with results presented in Fig. 10. In this figure, the colorbar represents the p -value,

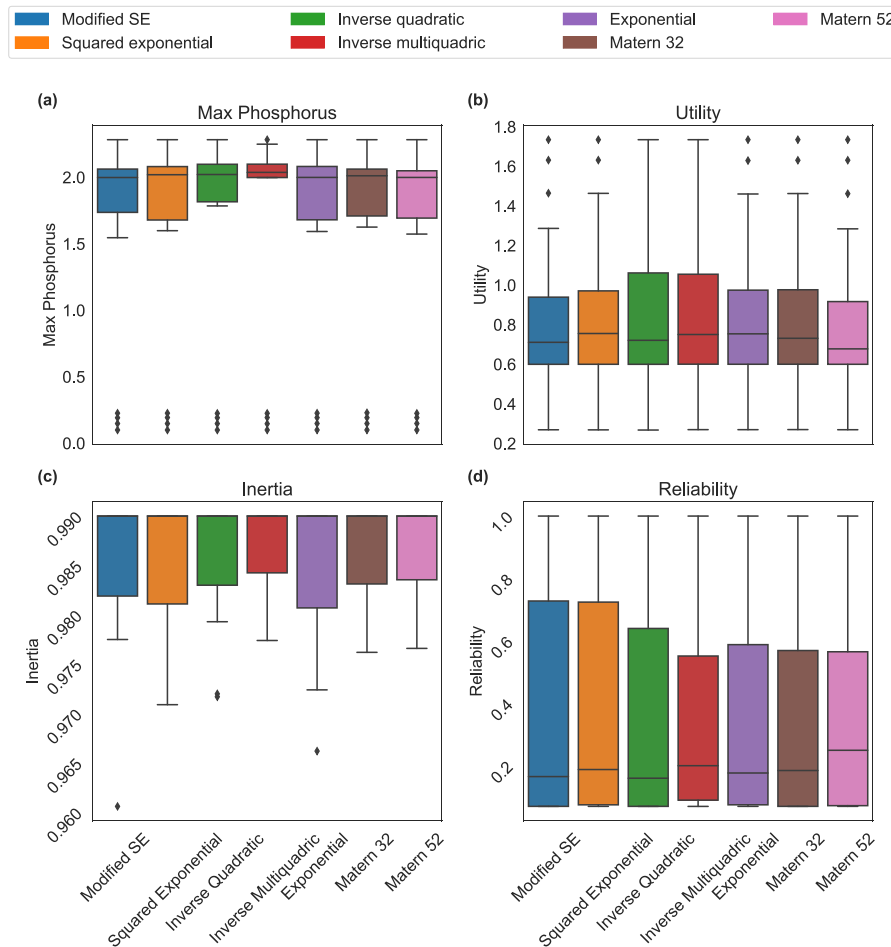


Fig. 7. Comparison of objective values for the Shallow Lake Problem per RBF configuration. The figure illustrates the distribution of objective values for various RBF configurations. Each subplot represents a distinct objective: (b) through (d) feature upward preferences due to the maximization objectives, whereas panel (a) features a downward preference as minimizing the maximum phosphorous level is desired. Within each subplot, individual box plots indicate the objectives' statistics obtained using different RBF configurations.

used to assess the statistical significance of the differences between objectives attained by each RBF configuration. Contrary to the Conowingo test results, the pairwise comparisons in the Lake Problem reveal no statistical significance; all the p-values exceed the 0.05 threshold. This outcome indicates that there are no notable differences between the objectives for the various RBF configurations examined in this case. This finding corroborates the conclusions drawn from both visual trade-off analysis and objective performance evaluations: while the policies prescribed by different RBF configurations exhibit varied behaviors, they do not result in statistically significant differences in the objectives they achieve.

5. Discussion

Does the choice of RBF within Evolutionary Multi-Objective Direct Policy Search influence both the search process and the resulting set of solutions? In the context of the Conowingo Reservoir System, the choice of RBF clearly has a substantial impact. Specifically, the modified squared exponential and the squared exponential RBFs yielded high-quality and diverse solutions. Our analysis highlights the considerable effect that the choice of RBF has on the objective space, the attributes of the release strategies, and the intra-annual variability within these policies. We observed that top-performing RBFs consistently align with and satisfy water demands over the year. In contrast, lower-performing RBFs exhibit a significant divergence between water demands and

actual releases particularly later in the year. From a reservoir management perspective, these fluctuations can have extensive implications for objectives sensitive to the timing of water releases. For instance, deviations in releases during critical periods may jeopardize environmental goals or the reliable supply of water to urban centers. On the other hand, for the Shallow Lake Problem, the choice of RBF resulted in only minor differences. Across multiple seeds, all RBFs converged to very similar results. Nevertheless, a detailed inspection of the emission strategies prescribed by each configuration reveals variations in phosphorus release decisions. In this context, disparate pollution release policies, despite their different approaches to managing phosphorus releases, are able to achieve nearly identical trade-offs. Furthermore, some RBFs had exhibited an outlier seed indicative of a stalled search.

What factors determine whether the choice of the RBF matters? A first partial explanation might be found in the nature of the control problem. For the Conowingo Reservoir System, the control setup is high dimensional. There are time-varying demands from the atomic power plant, and, to a lesser extent, from Chester. Baltimore's demand is stable over time. Hydropower revenue is also time-dependent. The net result is that the ideal control rule has to approximate as best as possible these different patterns over the year. In contrast, The Shallow Lake Problem requires a relatively simple control rule: start with a default maximum release, step the anthropogenic release down with an increasing pollution level, then increase the release again sharply once the tipping point in the lake has been crossed. If we compare the control rules for the modified squared exponential (best-performing)

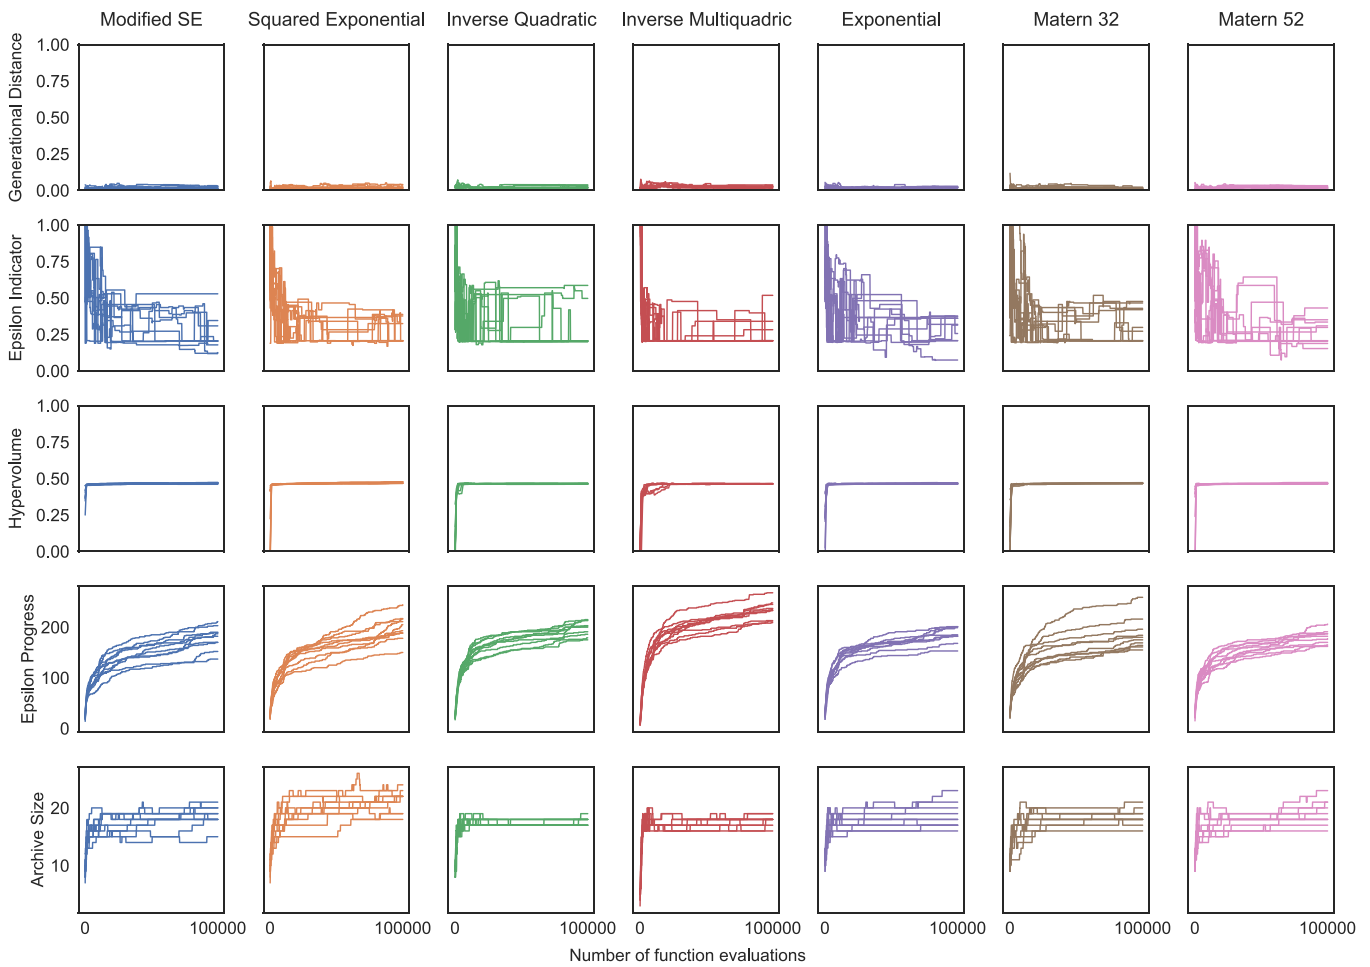


Fig. 8. Performance metrics for the Shallow Lake Problem. In this figure, each color and column represents a unique RBF configuration, while the rows illustrate the performance metrics. Lower values for the generational distance and epsilon indicator are favored, indicating closer proximity to the reference set and minimal translation distance, respectively. In contrast, a higher hypervolume is preferred, reflecting a larger volume dominated by the approximation set relative to the reference set; serving as a suitable measure of solution diversity and convergence. A steep slope in the epsilon progress is indicative of consistent advancement by adding new solutions to the archive. The relationship between archive size and solution diversity becomes especially relevant when the hypervolume is also high.

Modified squared exponential	hypower	6.1e-23	1e-224	5.8e-133	0	8.4e-288	3.7e-192
	atomicpowerplant	4.3e-05	2.9e-07	2.4e-08	1.3e-151	5.5e-37	7.4e-30
	baltimore	6.3e-05	2.4e-28	0.0066	2.7e-84	2.3e-25	4.3e-07
	chester	0.0012	3.5e-32	0.0037	5e-236	2.2e-123	2.5e-118
	environment recreation	0.15 5e-05	2.8e-49 1.9e-107	5.8e-249 2.1e-42	0.93 6.5e-71	2.1e-10 6e-18	0.63 8e-14
Squared exponential	hypower		8e-197	2.7e-114	0	1.3e-227	1.5e-145
	atomicpowerplant		0.44	0.1	1.8e-78	1.9e-13	1.6e-09
	baltimore		1e-50	2.8e-10	3.4e-98	1.1e-48	1.4e-21
	chester		2.3e-51	0.72	6.1e-258	9.2e-159	1.2e-152
	environment recreation		1.3e-57 1.3e-67	1.5e-222 6.5e-21	0.073 1.6e-40	8.3e-06 1.5e-05	0.27 0.00059
Inverse quadratic	hypower			0.3	2e-08	1.1e-138	2e-130
	atomicpowerplant			0.33	2.5e-74	5.9e-10	6.9e-07
	baltimore			7.1e-13	2.7e-07	0.024	7.7e-17
	chester			2.8e-34	7e-122	9.2e-32	2e-34
	environment recreation			7.7e-141 1.1e-13	6.8e-56 8.3e-05	4.5e-104 1.5e-42	1.2e-53 2.3e-40
Inverse Multiquadric	hypower				1.2e-09	6.7e-83	4.9e-78
	atomicpowerplant				8.6e-54	1.7e-06	0.00035
	baltimore				1.2e-35	9.3e-11	0.062
	chester				1.2e-182	1.4e-100	4.1e-98
	environment recreation				3.4e-309 0.00033	0 4.5e-08	1.5e-271 1e-08
Exponential	hypower					6.9e-248	1.8e-234
	atomicpowerplant					8.6e-48	1.1e-61
	baltimore					3.7e-21	9.5e-66
	chester					1.1e-89	9.9e-84
	environment recreation					1.5e-10 3.1e-21	0.7 3.9e-21 0.0035
Matern 32	hypower						0.39
	atomicpowerplant						3.9e-08
	baltimore						0.016
	chester						5e-09
	environment recreation						0.51
		Squared exponential	Inverse quadratic	Inverse Multiquadric	Exponential	Matern 32	Matern 52

Fig. 9. Kruskal–Wallis pairwise comparison of objective values across RBF configurations for the Conowingo Reservoir System. In this matrix, both rows and columns represent the pairwise comparisons between different RBF configurations, while the inner rows specify the objectives being compared. The colorbar encodes the *p*-value for statistical significance, with dark purple grids indicating high levels of significance and yellow grids denoting low statistical significance. In the Conowingo Reservoir case, it is evident that the majority of the comparisons yield statistically distinct outcomes.




Fig. 10. Kruskal-Wallis pairwise comparison of objective values across RBF configurations for the Shallow Lake Problem. In this matrix, both rows and columns represent the pairwise comparisons between different RBF configurations, while the inner rows specify the objectives being compared. The colorbar encodes the p -value for statistical significance, with dark purple grids indicating high levels of significance and yellow grids denoting low statistical significance. In the Shallow Lake Problem, it is evident that the majority of the comparisons do not yield statistically distinct outcomes.

and the exponential (worst-performing) for the Conowingo Reservoir System, we see that the exponential struggles much more to match the demand patterns. The RBFs' shape plays a role here, with the modified square exponential (more concave), outperforms the exponential (more convex) in matching the demand profiles for the atomic power plant, Baltimore and Chester (see also Table 1). This difference in shape is not relevant in the case of The Shallow Lake Problem.

A second partial explanation can be gleaned from Fig. 6. Despite quite different looking control rules, we see that the performance in the objective space is quite similar. That is, the Shallow Lake Problem shows a form of equifinality that is absent in the Conowingo Reservoir system. This might be due to the simple nature of The Shallow Lake Problem with a single input, single output, no competing demands, and strong tipping point behavior, so the system can only exist in two distinct states. In contrast, the Conowingo Reservoir System, is markedly more complicated, with multiple inputs and outputs, and no clear tipping point behavior.

There are two possible threats to the validity of these results. First, in our work, we fixed the number of RBFs following the guidance from the literature (Giuliani et al., 2016). It is conceivable that by increasing the number of RBFs, the observed differences for the Conowingo case might be reduced. However, this would increase the complexity of the optimization formulation and thus, in turn, affect the search behavior of the MOEA. Still, future research is needed to explore the impact of the number of RBFs. Second, we constrained the decision variables within the bounded normalized input space for the centroids of the RBF (i.e., c_i) and specified that the radius had to be strictly positive (i.e., $b_i \geq (0, 1]$). It is conceivable that allowing wider radii or allowing centroids outside the normalized input space might partially alleviate the observed search defects.

6. Conclusions

In recent years, there has been a growing interest in employing Evolutionary Multi-Objective Direct Policy Search (EMODPS) for the optimization of control rules in environmental systems, including reservoir control and pollution emissions (Quinn et al., 2017; Marangoni et al., 2021; Giuliani et al., 2014; Arnold et al., 2023). EMODPS relies on many-objective evolutionary algorithms to optimize control

rules, typically constructed using radial basis functions (RBFs), which serve as universal approximators. While many studies have followed the guidance outlined in the seminal work by Giuliani et al. (2016), which includes the use of modified square exponential RBFs, the literature also presents various other RBF variants with subtle shape differences. However, it has remained unclear whether and to what extent the choice of a specific RBF impacts the effectiveness of the multi-objective evolutionary algorithm (MOEA) and the quality of the resulting solutions. In this paper, we address this gap by investigating the performance of seven different RBFs in the context of optimizing control rules for the Conowingo Reservoir system and the Shallow Lake Problem.

For the Conowingo Reservoir problem, there is a clear and distinct difference in performance across the different RBFs. Looking at the overall attained hypervolume as a proxy for the overall quality of the set of identified solutions, the modified squared exponential, as suggested by Giuliani et al. (2016), performed the best, followed closely by the original squared exponential. The other five shapes all performed poorly on hypervolume. Likewise, other indicators for analyzing the search behavior of the MOEA also show the difficulty of finding high-quality solutions across multiple seeds for the MOEA. This performance disparity was clearly reflected in a visual inspection of the trade-offs and release strategies. High-performing RBFs produced release strategies that were consistently aligned with daily annual demands, while low-performing RBFs manifested poorer trade-offs and release strategies marked by substantial gaps between demands and actual releases. Conversely, for The Shallow Lake Problem, the choice of RBF had minimal impact. No discernible differences in search behavior or quality of solutions were observed between the various RBFs. Nevertheless, despite yielding nearly identical performance in the objective space, there were substantial differences in the emission strategies prescribed.

A possible explanation for the observed difference in the performance of the RBFs for the Conowingo Reservoir System is the difference in shape between e.g., the modified squared exponential and the exponential, with the former being more concave while the later is more convex. This difference enabled the (modified) squared exponential to better approximate the time-varying demands. In turn, this suggests that careful consideration of the control problem can guide the

selection of appropriate RBFs, although we still suggest experimenting with multiple shapes (e.g, modified squared exponential, and either a Matern or exponential).

In summary, our study highlights the substantial influence of Radial Basis Functions (RBFs) on the performance of Evolutionary Multi-Objective Direct Policy Search, particularly in the context of managing complex environmental systems like the Conowingo Reservoir System. This impact is evident through the quality of trade-offs achieved and the alignment of water release strategies with multiple demands.

However, it is worth noting that in simpler problems, such as the shallow lake problem, the choice of RBF has a relatively minor impact on solution quality. Nevertheless, it does result in variations in emission strategies.

Looking ahead, future research should prioritize investigating the effects of altering the number of RBFs and introducing constraints on RBF parameters. We propose the use of a diagnostic assessment framework, which holds the potential to extend its applicability to a wider range of cases that use EMODPS to find optimal control policies for multiple objectives.

Moreover, future research should also encompass the extension of this diagnostic assessment to problems featuring a higher number of objectives, increased uncertainties, and a larger number of states. In doing so, we can delve deeper into understanding the influence of global approximators on solution quality.

Furthermore, addressing another critical challenge lies in the interpretability of RBF control parameters, which play a central role in EMODPS. The lack of intuitive meaning in these parameters could hinder the real-world adoption and usability of EMODPS, where clear and understandable rationale behind strategies is crucial. Therefore, future research efforts should also prioritize enhancing policy interpretability.

Table 4
Susquehanna problem objective value's statistics across radial basis functions (RBFs).

Radial basis function	Hydropower	Atomic power plant	Baltimore	Chester	Environment	Recreation
Modified SE	Percentile 25: 57.09	Percentile 25: 0.68	Percentile 25: 0.34	Percentile 25: 0.55	Percentile 25: 0.05	Percentile 25: 1.0
	Percentile 75: 65.55	Percentile 75: 0.94	Percentile 75: 0.73	Percentile 75: 0.88	Percentile 75: 0.08	Percentile 75: 1.0
	Median: 59.70	Median: 0.85	Median: 0.53	Median: 0.76	Median: 0.06	Median: 1.0
	Min: 35.43	Min: 0.05	Min: 0.01	Min: 0.02	Min: 0.03	Min: 0.89
Squared exponential	Max: 81.10	Max: 0.99	Max: 0.92	Max: 0.95	Max: 0.11	Max: 1.0
	Percentile 25: 56.13	Percentile 25: 0.60	Percentile 25: 0.37	Percentile 25: 0.59	Percentile 25: 0.05	Percentile 25: 1.0
	Percentile 75: 60.50	Percentile 75: 0.94	Percentile 75: 0.75	Percentile 75: 0.88	Percentile 75: 0.08	Percentile 75: 1.0
	Median: 58.51	Median: 0.82	Median: 0.56	Median: 0.79	Median: 0.06	Median: 1.0
Inverse quadratic	Min: 36.05	Min: 0.05	Min: 0.01	Min: 0.04	Min: 0.03	Min: 0.93
	Max: 81.12	Max: 0.99	Max: 0.92	Max: 0.95	Max: 0.11	Max: 1.0
	Percentile 25: 37.54	Percentile 25: 0.60	Percentile 25: 0.29	Percentile 25: 0.47	Percentile 25: 0.06	Percentile 25: 0.96
	Percentile 75: 53.73	Percentile 75: 0.92	Percentile 75: 0.60	Percentile 75: 0.81	Percentile 75: 0.08	Percentile 75: 1.00
Inverse multiquadratic	Median: 44.91	Median: 0.81	Median: 0.43	Median: 0.68	Median: 0.07	Median: 1.00
	Min: 23.58	Min: 0.12	Min: 0.00	Min: 0.02	Min: 0.04	Min: 0.86
	Max: 79.90	Max: 0.99	Max: 0.87	Max: 0.91	Max: 0.10	Max: 1.00
	Percentile 25: 39.17	Percentile 25: 0.57	Percentile 25: 0.30	Percentile 25: 0.55	Percentile 25: 0.07	Percentile 25: 0.96
Exponential	Percentile 75: 68.77	Percentile 75: 0.93	Percentile 75: 0.70	Percentile 75: 0.90	Percentile 75: 0.08	Percentile 75: 1.00
	Median: 43.63	Median: 0.79	Median: 0.51	Median: 0.77	Median: 0.08	Median: 1.00
	Min: 21.26	Min: 0.02	Min: 0.01	Min: 0.00	Min: 0.06	Min: 0.86
	Max: 81.84	Max: 0.99	Max: 0.92	Max: 0.96	Max: 0.09	Max: 1.00
Matern 32	Percentile 25: 43.58	Percentile 25: 0.54	Percentile 25: 0.30	Percentile 25: 0.37	Percentile 25: 0.05	Percentile 25: 0.96
	Percentile 75: 50.11	Percentile 75: 0.78	Percentile 75: 0.52	Percentile 75: 0.60	Percentile 75: 0.07	Percentile 75: 1.0
	Median: 47.03	Median: 0.68	Median: 0.41	Median: 0.51	Median: 0.06	Median: 1.0
	Min: 24.33	Min: 0.06	Min: 0.02	Min: 0.06	Min: 0.03	Min: 0.82
Matern 52	Max: 69.69	Max: 0.86	Max: 0.62	Max: 0.66	Max: 0.09	Max: 1.0
	Percentile 25: 52.44	Percentile 25: 0.57	Percentile 25: 0.30	Percentile 25: 0.47	Percentile 25: 0.04	Percentile 25: 1.0
	Percentile 75: 56.38	Percentile 75: 0.90	Percentile 75: 0.61	Percentile 75: 0.72	Percentile 75: 0.07	Percentile 75: 1.0
	Median: 54.08	Median: 0.77	Median: 0.46	Median: 0.62	Median: 0.06	Median: 1.0
Matern 32	Min: 32.02	Min: 0.01	Min: 0.00	Min: 0.01	Min: 0.03	Min: 0.89
	Max: 76.15	Max: 0.98	Max: 0.76	Max: 0.81	Max: 0.10	Max: 1.0
	Percentile 25: 52.30	Percentile 25: 0.62	Percentile 25: 0.38	Percentile 25: 0.48	Percentile 25: 0.05	Percentile 25: 1.0
	Percentile 75: 56.93	Percentile 75: 0.89	Percentile 75: 0.61	Percentile 75: 0.70	Percentile 75: 0.07	Percentile 75: 1.0
Matern 52	Median: 54.52	Median: 0.78	Median: 0.51	Median: 0.61	Median: 0.06	Median: 1.0
	Min: 35.92	Min: 0.07	Min: 0.00	Min: 0.01	Min: 0.03	Min: 0.86
	Max: 75.13	Max: 0.97	Max: 0.74	Max: 0.78	Max: 0.09	Max: 1.0
	Percentile 25: 52.30	Percentile 25: 0.62	Percentile 25: 0.38	Percentile 25: 0.48	Percentile 25: 0.05	Percentile 25: 1.0

This entails striking a balance between the flexibility of mapping states to actions in systems with multiple states and objectives, while ensuring that these mappings are intuitively meaningful in decision-making contexts.

Reproducibility statement

For those interested in reproducing our work, we have provided the necessary software and guidelines in the following repository (Zatarain Salazar, 2023). These resources serve as a comprehensive reference for replicating our research findings.

CRedit authorship contribution statement

Jazmin Zatarain Salazar: Conceptualization of this study, Methodology, Software, Original draft preparation. **Jan H. Kwakkel:** Writing, Software. **Mark Witvliet:** Execution of the Conowingo study, Software.

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.

Data availability

I have shared the data and code information, also cited within the manuscript.

Appendix

Table 5
Lake Problem objective value's statistics across radial basis functions (RBFs).

Radial basis function	Max P	Inertia	Utility	Reliability
Modified SE	Percentile 25: 1.74 Percentile 75: 2.06 Median: 2.00 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.96 Max: 0.99	Percentile 25: 0.60 Percentile 75: 0.94 Median: 0.71 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.73 Median: 0.17 Min: 0.07 Max: 1.00
Squared exponential	Percentile 25: 1.68 Percentile 75: 2.08 Median: 2.02 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 0.97 Median: 0.76 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.72 Median: 0.19 Min: 0.07 Max: 1.00
Inverse quadratic	Percentile 25: 1.82 Percentile 75: 2.10 Median: 2.02 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 1.06 Median: 0.72 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.64 Median: 0.16 Min: 0.07 Max: 1.00
Inverse multiquadratic	Percentile 25: 1.99 Percentile 75: 2.10 Median: 2.04 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 1.06 Median: 0.75 Min: 0.27 Max: 1.73	Percentile 25: 0.09 Percentile 75: 0.55 Median: 0.20 Min: 0.07 Max: 1.00
Exponential	Percentile 25: 1.68 Percentile 75: 2.08 Median: 2.00 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 0.97 Median: 0.75 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.59 Median: 0.18 Min: 0.07 Max: 1.00
Matern 32	Percentile 25: 1.71 Percentile 75: 2.06 Median: 2.01 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 0.98 Median: 0.73 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.57 Median: 0.18 Min: 0.07 Max: 1.00
Matern 52	Percentile 25: 1.69 Percentile 75: 2.05 Median: 2.00 Min: 0.10 Max: 2.28	Percentile 25: 0.98 Percentile 75: 0.99 Median: 0.99 Min: 0.97 Max: 0.99	Percentile 25: 0.60 Percentile 75: 0.92 Median: 0.68 Min: 0.27 Max: 1.73	Percentile 25: 0.07 Percentile 75: 0.57 Median: 0.25 Min: 0.07 Max: 1.00

References

Arnold, W., Salazar, J.Z., Carlino, A., Giuliani, M., Castelletti, A., 2023. Operations eclipse sequencing in multipurpose dam planning. *Earth's Future* 11 (4), e2022EF003186.

Askari, M., Adibi, H., 2015. Meshless method for the numerical solution of the Fokker–Planck equation. *Ain Shams Eng. J.* 6 (4), 1211–1216.

Busoni, L., Ernst, D., De Schutter, B., Babuska, R., 2009. Policy search with cross-entropy optimization of basis functions. In: 2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning. IEEE, pp. 153–160.

Carpenter, S.R., Ludwig, D., Brock, W.A., 1999. Management of Eutrophication for lakes subject to potentially irreversible change. *Ecol. Appl.* 9 (3), 751–771.

Doering, K., Quinn, J., Reed, P.M., Steinschneider, S., 2021. Diagnosing the time-varying value of forecasts in multiobjective reservoir control. *J. Water Resour. Plan. Manag.* 147 (7), 04021031.

Fasshauer, G.E., 2007. Meshfree Approximation Methods with MATLAB, Vol. 6. World Scientific.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., Reed, P.M., 2016. Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations. *J. Water Resour. Plan. Manag.* 142 (2), 04015050.

Giuliani, M., Herman, J.D., 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.

Gupta, R.S., Hamilton, A.L., Reed, P.M., Characklis, G.W., 2020. Can modern multi-objective evolutionary algorithms discover high-dimensional financial risk portfolio tradeoffs for snow-dominated water-energy systems? *Adv. Water Resour.* 145, 103718.

Hadka, D., Herman, J.D., Reed, P.M., Keller, K., 2015. OpenMORDM: An open source framework for many objective robust decision making. *Environ. Model. Softw.* 74, 114–129. <http://dx.doi.org/10.1016/j.envsoft.2015.07.014>.

Hadka, D., Reed, P., 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evol. Comput.* 20 (3), 423–452.

Hashimoto, T., Stedinger, J.R., Loucks, D.P., 1982. Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. *Water Resour. Res.* 18 (1), 14–20.

Kollat, J.B., Reed, P.M., 2005. The value of online adaptive search: A performance comparison of NSGAI, ϵ -NSGAI and ϵ moea. In: International Conference on Evolutionary Multi-Criterion Optimization. Springer, pp. 386–398.

Kollat, J.B., Reed, P.M., 2006. Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. *Adv. Water Resour.* 29 (6), 792–807.

Marangoni, G., Lamontagne, J.R., Quinn, J.D., Reed, P.M., Keller, K., 2021. Adaptive mitigation strategies hedge against extreme climate futures. *Clim. Change* 166 (3), 1–17.

McInerney, D., Lempert, R., Keller, K., 2012. What are robust strategies in the face of uncertain climate threshold responses. *Climate Change* 112 (3), 547–568.

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. Softw.* 92, 125–141. <http://dx.doi.org/10.1016/j.envsoft.2017.02.017>.

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.

Schaback, R., 2007. A practical guide to radial basis functions. *Electron. Resource* 11, 1–12.

Singh, R., Reed, P.M., Keller, K., 2015. Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response. *Ecol. Soc.* 20 (3), <http://dx.doi.org/10.5751/ES-07687-200312>.

Ward, V.L., Singh, R., Reed, P.M., Keller, K., 2015. Confronting tipping points: Can multi-objective evolutionary algorithms discover pollution control tradeoffs given environmental thresholds? *Environ. Model. Softw.* 73 (1), 27–43.

Williams, C.K., Rasmussen, C.E., 2006. Gaussian Processes for Machine Learning, Vol. 2. MIT press Cambridge, MA.

Zatarain Salazar, J., 2023. MultiObjPolicy-RBF analysis kit: Exploring radial basis functions for multiobjective policy optimization. <http://dx.doi.org/10.5281/zenodo.10073327>.

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.

Zhang, H., Chen, Y., Guo, C., Fu, Z., 2014. Application of radial basis function method for solving nonlinear integral equations. *J. Appl. Math.* 2014.

Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., Da Fonseca, V.G., 2003. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Trans. Evol. Comput.* 7 (2), 117–132.