Exploring global approximators for multiobjective reservoir control

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

DOI
10.1016/j.ifacol.2022.11.006

Publication date
2022

Document Version
Final published version

Published in
IFAC-PapersOnline

Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 55-33 (2022) 34–41

Exploring global approximators for multiobjective reservoir control

Jazmin Zatarain Salazar * **Jan Kwakkel** * **Mark Witvliet** *

* *Faculty of Technology, Policy and Management, Delft University of Technology, Delft, The Netherlands (e-mail:J.ZatarainSalazar@tudelft.nl, J.H.Kwakkel@tudelft.nl)*

Abstract: Efficient multi-purpose reservoir control policies are crucial in the face of frequent and severe floods and droughts, and to balance water allocation across conflicting demands. Evolutionary Multi-Objective Direct Policy Search (EMODPS) is a popular approach to design control policies for multi-purpose reservoir systems. EMODPS, however, relies on experimental choices within the key components of the framework particularly when coupling multi-objective evolutionary optimization with nonlinear approximation networks. This study explores a suite of radial basis functions (RBFs) used to map the system's states to control actions in a flexible manner as time-varying, non-linear relationships. We provide a systematic assessment of different RBF functions to explore their suitability to obtain Pareto efficient control policies. We use the Susquehanna river basin case study in which competing water demands for hydropower, environment, urban water supply, atomic power plant cooling and recreation need to be met. Our findings suggest that the choice of RBF functions have a large impact on the model outcomes and the search behavior of the optimization algorithm.

Keywords: Optimal operation of water resources systems,direct policy search, global approximators

1. INTRODUCTION

Effective operations in existing reservoir systems are essential to meet competing water demands and to cope with challenging hydro-climatic conditions. Evolutionary Multi-Objective Direct Policy Search (EMODPS) has proven to be a flexible and generalizable approach to design 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 approximators, with multi-objective evolutionary algorithms (MOEAs). 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). The success of the EMODPS, is nonetheless, highly dependent on the function which maps states to actions. A flexible structure that is capable of capturing nonlinear relationships is required and so far, radial basis functions 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).

Busoniu et al. (2009) recommends the use of a Gaussian distribution for continuous RBF parameters due to its unbounded support. Giuliani et al. (2014) later used Gaussian RBFs for direct policy search within the EMODPS framework. However, the performance of alternative universal approximators requires further exploration. We aim to understand to what extent nonlinear approximation within the EMODPS framework can affect the recommendations that will influence efficient reservoir management strategies. We use the Conowingo reservoir in the Lower Susquehanna River Basin (LSRB), USA, to study the impact of nonlinear approximation networks over multi-objective reservoir control described in Section 1.1.

1.1 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 demands for hydropower, 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 the simulation time horizon of one year. A yearly simulation horizon has been chosen 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 (to be maximized) is defined as the economic revenue gained from hydropower production at the Conowingo hydropower dam in US$/MWh defined

⋆

10.1016/j.ifacol.2022.11.006

in Eq. 1. The energy prices are defined by the seven hour moving average of the energy price trajectory in the Pennsylvania, New Jersey–Maryland (PJM) energy market (Exelon, 2010). The hourly energy production (MWh) is defined by Eq.2 , where η is the turbine efficiency, g is the gravitational acceleration ($9.81m/s^2$), γ_w is the water density ($1000kg/m^3$), $\bar{h_t}$ is the net hydraulic water level difference (head) in meters and q_t^{Turb} is the turbine flow in m^3/s.

$$J^{hyd} = \sum_{t=1}^{H} (HP_t \cdot \rho_t) \tag{1}$$

$$HP_t = \eta g \gamma_w \bar{h_t} q_t^{Turb} \cdot 10^{-6} \tag{2}$$

Water supply reliability to the Atomic Power Plant, Chester and Baltimore (to be maximized) The daily average volumetric reliability defined as:

$$J^{VR,i} = \frac{1}{H} \sum_{t=1}^{H} (Y_t^i / D_t^i) \tag{3}$$

where $Y_t^i(m^3)$ is the daily delivery, $D_t^i(m^3)$ is the corresponding demand, and i represents the water supply to Baltimore, Chester or the Atomic Power Plant. Figure **??** illustrates the demand in cubic feet per second for each objective.

Environmental Shortage (to be minimized) specified as the daily average shortage index in regard to the FERC minimum flow requirements, defined as follows:

$$J^{SI} = \frac{1}{H} \sum_{t=1}^{H} \left(\frac{max(Z_t - Y_t, 0)}{Z_t} \right)^2 \tag{4}$$

where Y_t (m3) is the daily release and Z_t (m3) is the corresponding FERC flow requirement. The quadratic formulation (Eq. 4) is intended to penalize substantial deficits in a single time step while allowing for more frequent, minor shortages Hashimoto et al. (1982). Figure **??** depicts the monthly water supply demands as well as the FERC minimum flow criteria.

Recreation (to be maximized) defined in Eq. 5 as the storage reliability throughout the tourist season's weekends. Given by the relationship between the number of weekend days in the peak season that are less than the intended level (n_F) and the total number of weekends in the tourist season (N_{we}). To ensure recreational activities, the target level is 32.5 m (106.5 ft).

$$J^{SR} = 1 - \frac{n_F}{2N_{we}} \tag{5}$$

2. METHODS

An initial step towards offering guidance for flexible reservoir control using EMODPS, requires understanding the influence that the family of functions shaping the control policy have over the tradeoffs of interest. The main components of the framework are (1) the use of nonlinear approximators to parameterize the operating policy for multiple objectives, and (2) the identification of parameterizations that yield Pareto approximate reservoir control policies using multi-objective evolutionary search (Zatarain Salazar et al., 2016; Giuliani et al., 2016). We focus on the first component of the framework, the non-linear approximators, in particular on RBFs. The main differences between an RBF network and other ANN architectures is that it always consists of an input layer, a hidden layer with a nonlinear activation function and a linear output layer. Each neuron in the hidden layer has a kernel function, also known as an activation function, distinguished by a center x_i and radius σ. The idea is to apply locality-sensitive transformations of the data points into high-dimensional spaces, allowing the altered points to be linearly separable (Aggarwal, 2018). Correlation is high when the points are close together and is low when they are far apart. The output layer consists of a neuron for each output, giving the predicted value by multiplying the sum of the weights with the euclidean distance for each RBF. This gives an RBF network the ability to handle nonlinear problems.

The control policy is described in Equation 6, where the vector u_t is a sum of basis functions defined by the input vector x, and the output of the k^{th} node in the output layer (with $k = 1, ..., N_u$):

$$u^k = \sum_{i=1}^{n} w_i^k \phi_i(x_t) \tag{6}$$

where $\mathbf{n}$ is the number of RBFs and $\mathbf{w_i}$ 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 to one (*i.e.*, $\sum_{i=1}^{n} w_i = 1$). The function is only a valid covariance function when its variance is non-negative for every possible choice of $\mathbf{n}$. The function models the joint variability of the functions' random decision variables. It returns the modelled covariance between each pair in x_a and x_b. $\phi_i(x_t)$ is an activation function. The initial activation function used in the Susquehanna case is:

$$\phi_i(x) = \exp \left[- \sum_{j=1}^{m} \frac{(x_j - c_{j,i})^2}{b_{j,i}^2} \right] \tag{7}$$

where $\mathbf{m}$ is the number of input variables $\mathbf{x}$, and $\mathbf{c}_i$, $\mathbf{b}_i$ are the m-dimensional center and radius vectors of the ith RBF, respectively. The centers of the RBF must lie within the bounded input space and the radii must strictly be positive (*i.e.*, using normalized variables, $\mathbf{c_i} \in [-1, 1]$ and $\mathbf{b}_i \geq (0, 1]$). The parameter vector θ is therefore defined as $\theta = [c_{i,j}, b_{i-,j}, w_i^k]$, with $i = 1, ..., n, j = 1, ..., m$, and $k = 1, ..., N_u$ (Giuliani et al., 2016).

2.1 Activation functions (Kernels)

An activation function qualifies as a radial basis function when it is semi-positive definite and isotropic (Williams and Rasmussen, 2006). Positive definite functions have the property that the function of x is greater or equal to zero for all x. Activation functions that are isotropic only depend on the difference between $x - x'$ (*i.e.* the

euclidean distance). We explored popular isotropic and positive definite functions, as well as the multiquadric function shown in Table 1 (Fasshauer, 2007; Schaback, 2007; Askari and Adibi, 2015; Zhang et al., 2014).

It is worth noting that all RBF kernels are implemented with the Euclidean distance. With the modified squared exponential RBF Kernel, the working range of the RBFs improved when the Euclidean distance was implemented, this was particularly true for the Matern 3/2 and the Matern 5/2 kernel.

2.2 Experimental design

Our goal is to understand how the shape of the RBF affects the performance of the tradeoffs when using EMODPS. The number of RBF functions was kept constant at $n = 4$. The centers, radii and weights where 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 where preserved across all the tested activation functions in Table 1. Each configuration was run for 10 random seed trials to account for variability on the initial MOEA population, and each experiment was run for 100,000 function evaluations.

2.3 Assessment Metrics

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

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

Additive Epsilon-Indicator The additive ϵ indicator (Zitzler et al., 2003) assesses consistency of the Pareto approximate set, that is, the ability to capture all the regions of the tradeoff space. The metric is calculated as the largest distance that an approximation set must move in order to dominate the reference set, making it extremely sensitive to gaps in trade-offs (Zatarain Salazar et al., 2016). In other words, it measures the worst case distance needed in the approximation set in order to dominate its nearest neighbor in the reference set (Reed et al., 2013; Hadka and Reed, 2012).

Hypervolume 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. Where each objective adds another dimension to the multidimensional volume. It calculates the difference in hypervolume between the reference set and the Pareto approximation set (Reed et al., 2013).

Epsilon Progress ϵ-progress is a computationally efficient indication of search progress and stagnation. When the current solution sits in a different ϵ-box that dominates the previous solution, then ϵ-progress occurs. 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.

Archive Size 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 are used to compute the performance metrics. A bigger archive size can give a more complete image of the trade-off space. But this can only be argued when both convergence and diversity are also high. The archive size can also give information about micro evolution on different parts of the Pareto front.

3. RESULTS

We assessed the performance of each of the activation functions specified in Table 1 via a visual analysis of the tradeoffs and through multiobjective performance metrics that indicate their ability to converge and diversify.

3.1 Comparison of trade-offs and performance metrics

Figure 1 shows the trade-offs attained by the different activation functions. Each RBF configuration is depicted by a different color, 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 tradeoff was discovered.

Inspecting the tradeoffs for each RBF configuration, we note that the modified squared exponential RBF (1a), the squared exponential RBF (Figure 1b) and the inverse multiquadric RBF (Figure 1d) attain a diverse solution space with high performance across each objective, reflected by lines hitting the top of each objective axis. The inverse quadratic RBF in Figure 1c performs slightly worse, however, it finds a diverse set of solutions. In contrast, the Exponential RBF (Figure 1e), Matern 3/2 kernel (Figure 1f), and Matern 5/2 kernel (Figure 1g), find a narrow trade-off space. This may indicate that the exponential RBF, Matern 3/2 kernel and Matern 5/2 kernel have difficulty traversing the fitness landscape, causing the approximation set to be less diverse.

3.2 Performance across the Lower Susquehanna River basin objectives

In Figure 2 we inspect in detail the performance across each of the LSRB objectives. Each pane shows the objective distributions attained by different activation functions. The boxplots show consistent results with the parallel coordinate plots in Figure 1, with the benefit of detecting detailed differences across each objective.

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

Activation Function	$\phi(r)$	Reference
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)

We observe that the the modified squared exponential and the squared exponential RBFs attain the highest median values for the hydropower objective (panel a). The interquartile range for these two RBFs reflects a difference of approximately 10 M USD/ year, with several low performance outliers. In contrast, the inverse quadratic and the inverse multiquadratic, show a median value of approximately 45 M USD/year, with a difference of approximately 30 M USD/year between the lower and upper quartiles. The whiskers reflect, nonetheless, that these functions find the entire range of hydropower values, this is particularly true for the inverse multiquadratic RBF. The distribution of the exponential and Matern 3/2 RBFs result in a short box with values clustered around 45-50 and 50-55 M USD/year respectively. The Matern 3/2 and Matern 5/2 achieve similar median values, however, the outcomes from the Matern 5/2 RBF result in a higher upper quartile. Moving on to panel b, all the activation functions achieved a median value higher than 75% reliability for the atomic power plant, with the exception of the exponential RBF. It is worth noting that the squared exponential variants, the inverse quadratic and inverse multiquadratic RBFs found solutions close to 100% reliability for this objective. The lowest performance for the atomic power plant was linked to the exponential RBF, with the 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, in this case, the interquartile range is broader for the modified squared exponential. The inverse multiquadratic and the exponential RBFs attain the lowest median reliability for Baltimore at roughly 40%, however, 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, whose boxplot is short and its lowest and highest performance is bounded between 0.15 and 0.62. The median values attained by Matern 3/2 and Matern 5/2 are around 0.5, with both whiskers skewed towards lower reliabilities. We follow next the reliability results for Chester in panel d, we note that the modified squared exponential and the squared exponential RBFs again have the highest performance reflected by the median values, upper quartile ranges and their highest attained value, these two configurations are followed by the inverse multiquadratic RBF with a slightly lower median value and lower 25 % score values. Finally, the exponential RBF again is outperformed by all the other tested RBFs for water supply reliability. The modified squared exponential and the squared exponential RBFs find the largest range for the environmental objective in panel e. Interestingly, the Matern 3/2 and Matern 5/2 kernels find a good range of solutions for this objective, where the Matern 3/2 outperforms the squared exponential variants which attained consistently high performance across the other LSRB objectives. This could be explained by the stark trade-off between the environmental objective and the other objectives highlighted in section 3.1.

Finally, 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.

3.3 Performance metrics compared

Figure 3 shows the performance metrics for each activation function. The metrics where 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), set in this case to a maximum of 100,000. Further, each row shows a different metric value, and each column shows a distinct RBF. Finally, the colored plots depict the runtime dynamics for each configuration.

Generational distance 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 our configurations.

Additive epsilon indicator 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 sample of the MOEA 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 fail to find epsilon values close to the global reference set.

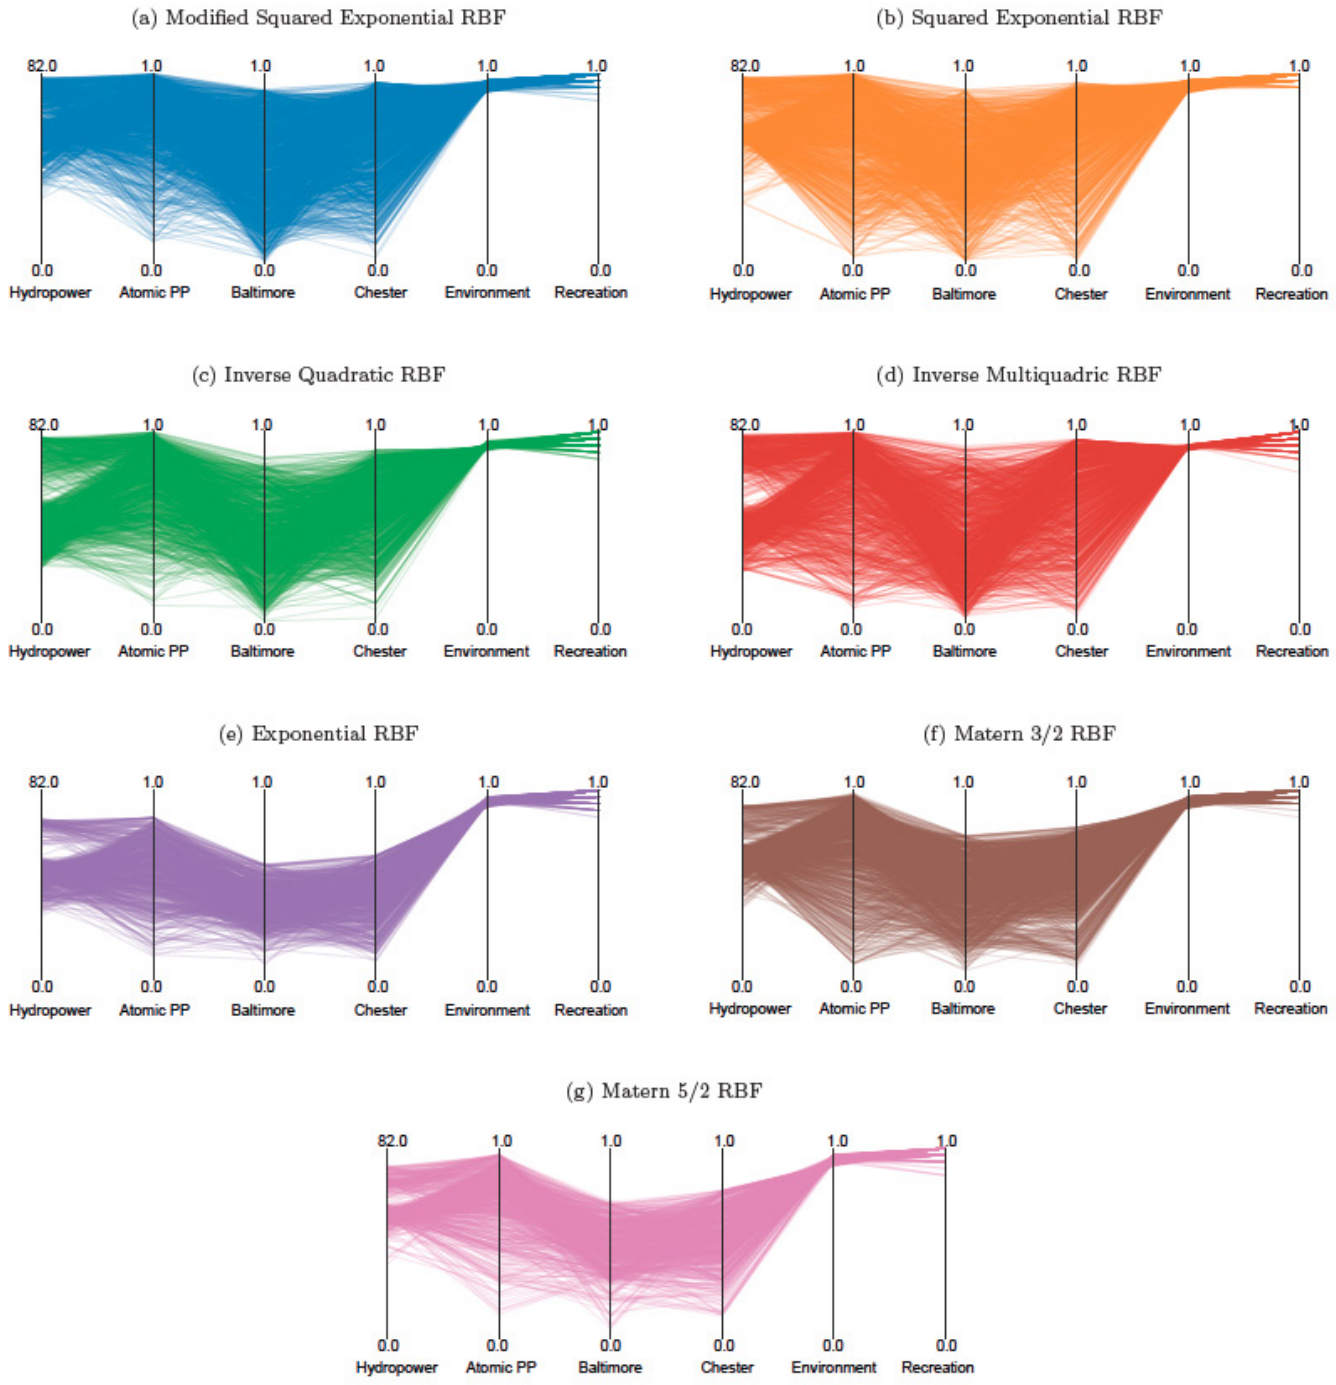

Fig. 1. Parallel coordinates compared

Hypervolume Hypervolume provides the most thorough measure of convergence and diversity making it the hardest metric to meet. Namely, 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 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 nfes. 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.

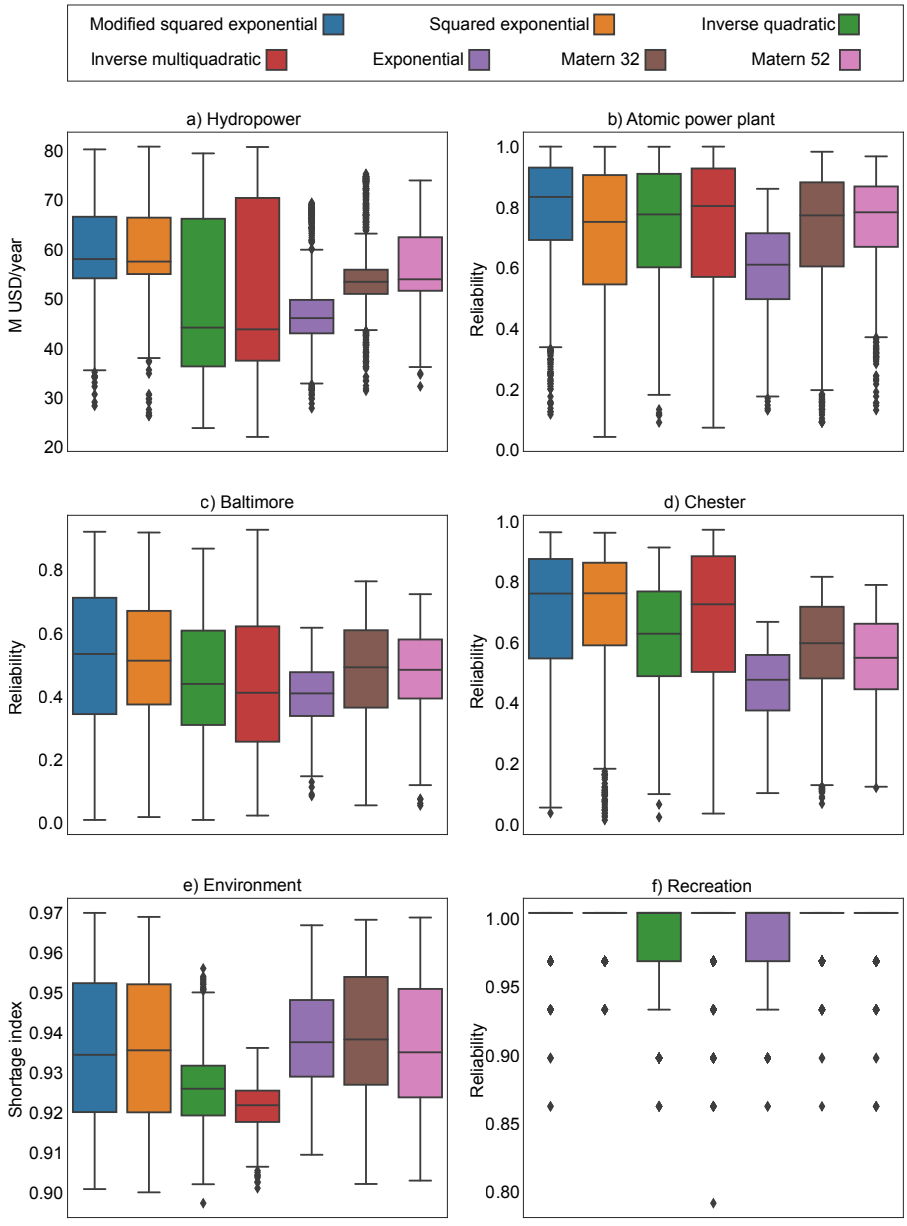

Fig. 2. Comparing RBFs per objective

Epsilon Progress 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 shows milder progress throughout the search and a wider performance range across random seed trials. This is also true for the exponential, and for the Matern 3/2 and Matern 5/2, 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. 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 and others get stuck in local optima.

Archive size 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 the 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 100k nfes. Evidently, the archive sizing is highly

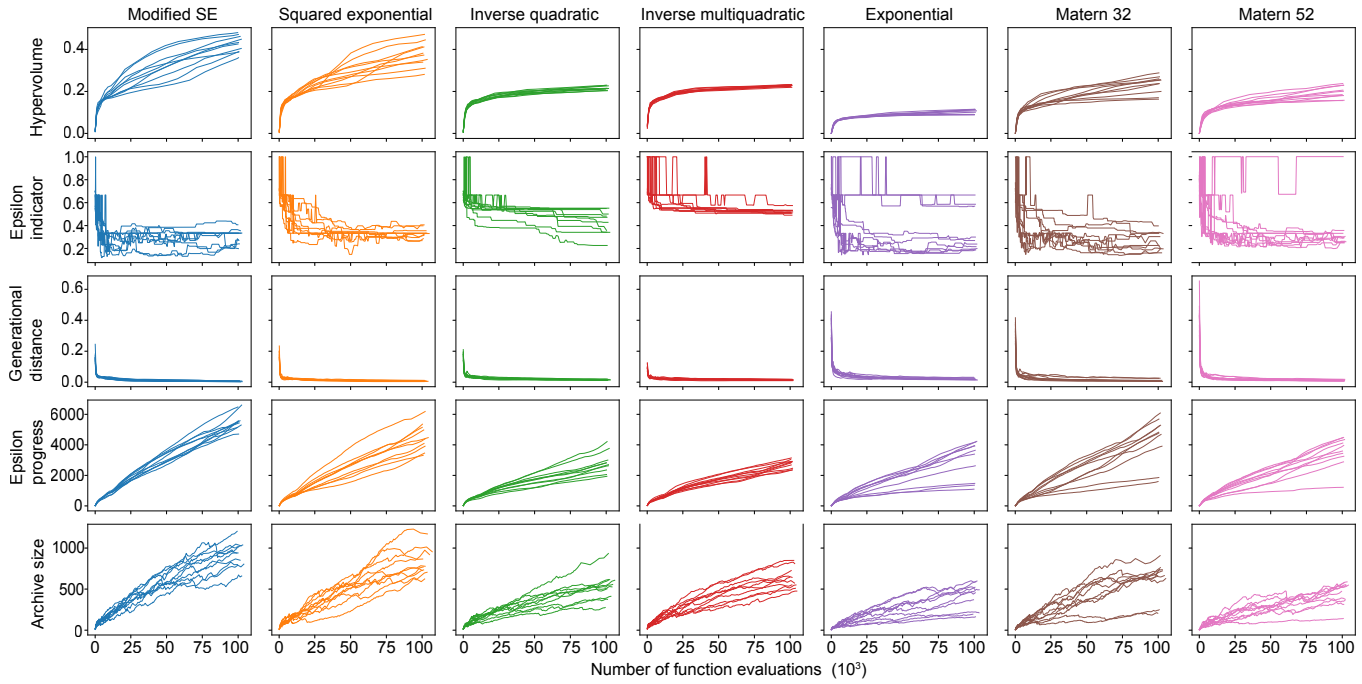

Fig. 3. Comparison of hypervolume, Additive epsilon indicator and generational distance of different RBFs

dependent on the starting populations reflected by a large random seed variability across all the activation functions.

Table 2 summarizes the statistical differences across the assessed activation functions. It shows the average across seeds of the maximum and minimum attained values of the metrics shown in Figure 3. The contribution of each activation function to the global reference set is also presented. This represents the amount of non-dominated solutions which are present in the global reference set which were generated from the model with that activation function. As seen in the table, the modified squared exponential RBF had the largest contribution (56%) to the global reference set, followed by the squared exponential RBF (28%), whereas the lowest contribution is attributed to the exponential RBF (0.1%).

4. CONCLUSION

The aim of this study was to explore how the choice of nonlinear approximation networks within the EMODPS framework impact the tradeoffs of multi-sector water allocations. We used the Lower Susquehanna River Basin model as a test case as it embodies common challenges in reservoir control, including multiple competing demands for hydropower, urban water supply, environment, atomic power plant cooling and recreation, making it a suitable candidate to test non-linear approximators for EMODPS. Seven popular activation functions were investigated within this study; the modified squared exponential, the squared exponential , the inverse quadratic, the inverse multiquadric, the exponential, the Matern(3/2), and the Matern(5/2). The activation functions where coupled with the *epsilon*-NSGAII MOEA due to the algorithm's ability to balance convergence and speed.

Our analysis shows that the shape of the activation functions significantly influence the performance of EMODPS models. The question is: What characteristics of shape of the activation function are causing them to behave more desirably? The differences in performance can be explained by changes in the fitness landscape traversed by the evolutionary algorithm resulting in distinct search dynamics. Key results show that the modified squared exponential RBF and the squared exponential RBF achieved a diverse set of tradeoffs and consistently attained high performance metric values. This could indicate that the optimization algorithm was able to create a smoother fitness landscape with the squared exponential variants compared to the other activation functions. Overall, this resulted in high hypervolume performance and epsilon progress, reflecting the ability to deal better with local optima and quickly find a diverse set of tradeoffs. The analysis suggests that both the modified squared exponential and the squared exponential activation functions are suitable candidates to be coupled with EMODPS. Lastly, the selection of an appropriate family of functions is increasingly relevant when finding tradeoffs through generative methods, in which a compromise between flexibility and accuracy is inevitable. Even if the goal is not to find accurate but good enough tradeoffs for decision support, clear diagnostic assessments of methodological choices that impact the results obtained via parameterization-optimization-simulation are required to increase their chances of success, to improve their reliability and trust, and potentially move towards narrowing the gap between the research and the intended applications in reservoir control.

Table 2. Average best attained performance metrics per activation function

Activation Function	Non-dominated solutions	Mean Generational distance	Mean Epsilon indicator	Mean Hypervolume	Set Contribution
Modified SE	1514	0.005	0.308	0.426	56%
Squared Exponential	1357	0.007	0.337	0.378	28%
Adapted Inverse Quadratic	1266	0.016	0.425	0.216	4%
Adapted Inverse Multiquadric	1157	0.014	0.522	0.229	10%
Exponential	946	0.019	0.343	0.102	0.1%
Matern 3/2	1294	0.012	0.269	0.234	1%
Matern 5/2	857	0.011	0.362	0.196	2%

REFERENCES

Aggarwal, C.C. (2018). Neural networks and deep learning. *Springer*, 10, 978–3.

Askari, M. and Adibi, H. (2015). Meshless method for the numerical solution of the fokker–planck equation. *Ain Shams Engineering Journal*, 6(4), 1211–1216.

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

Doering, K., Quinn, J., Reed, P.M., and Steinschneider, S. (2021). Diagnosing the time-varying value of forecasts in multiobjective reservoir control. *Journal of Water Resources Planning and Management*, 147(7), 04021031.

Fasshauer, G.E. (2007). *Meshfree approximation methods with MATLAB*, volume 6. World Scientific.

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

Giuliani, M., Herman, J.D., Castelletti, A., and Reed, P. (2014). Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management. *Water resources research*, 50(4), 3355–3377.

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

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

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

Reed, P.M., Hadka, D., Herman, J.D., Kasprzyk, J.R., and Kollat, J.B. (2013). Evolutionary multiobjective optimization in water resources: The past, present, and future. *Advances in water resources*, 51, 438–456.

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

Williams, C.K. and Rasmussen, C.E. (2006). *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.

Zatarain Salazar, J., Reed, P.M., Herman, J.D., Giuliani, M., and Castelletti, A. (2016). A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances in water resources*, 92, 172–185.

Zhang, H., Chen, Y., Guo, C., and Fu, Z. (2014). Application of radial basis function method for solving nonlinear integral equations. *Journal of Applied Mathematics*, 2014.

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