Chapter 6: EVALUATING DECISION VARIABLE IMPORTANCE IN LARGE SCALE GROUNDWATER OPTIMIZATION

6.1 Introduction

Decision-making in GW multi-objective optimization involves balancing the trade-off(s) between an obtained set of non-dominated solutions. While the studies yield valuable insights, uncertainties from simulation models' parameters could potentially affect the effectiveness and feasibility of proposed solutions (Raso et al., 2014). Decision-making typically occurs within the objective space of optimization problems, often neglecting decision variables that offer deeper insights into pumping well behavior. Thus, identifying critical decision variables and producing more robust solutions, previously unexplored in multi-objective GW management problems aiming to maximize extraction while considering R-A exchanges, holds promise for improving outcomes.

The pareto front is often analyzed for choosing a single solution from a set of non-dominated solutions to manage GW extraction and other objectives. Three main methods are used: A-priori, interactive, and A-posteriori. The abovementioned methods have been employed independently in the literature, and their comparison in the final output has also been critically examined. Thus far, different tools have been used to increase the efficiency of S-O models to obtain pareto fronts, e.g., novel metaheuristic algorithms were introduced that can handle large-scale optimization, artificial intelligence-based surrogate modeling such as physics-informed neural networks were used to speed up the simulation process.

However, A common observation in previous research is that choosing an optimal solution is mainly based on objective space. The decision space in multi-objective optimization of groundwater problems is relatively unexplored. Whether the solution of an S-O model implicitly depicts the decision space performance can be debated. For example, the solution of a GW pumping optimization (pareto front) would imply a combination of decision variables (pumping well rates). However, other instances may exist where a slight variation in this DV space may not find its place in the objective space. In such cases, finding an optimal solution based on stakeholders' preferences may not be in the Pareto set. However, it can be found in the decision space. Thus, exploration of DV space in GW optimization problems can be important, e.g., providing an understanding of objective functions and pumping wells, revealing patterns and relationships between a combination of wells and their effect on the final outputs, decision making in terms of proposing new zones for recharge and sustainability or GW potential zones.

Simulation optimization has been mostly studied under the assumption of a known input model without accounting for input parameter uncertainty. The input models are often constructed from data; hence, input uncertainty arises due to the finiteness of data. Therefore, when uncertainty in the model outputs is considered, a management plan deemed optimal while utilizing a deterministic S-O model may no longer be the best course of action.

To account for this uncertainty, previous studies have suggested using stochastic simulations, chanced constraint programming approach, multiple realization approach, and Monte Carlo simulations (Rajabi and Ketabchi, 2017). These studies are useful in addressing the uncertainty and error accounting in simulation models. However, the optimization model’s uncertainty is rarely discussed. It is widely recognized that various

optimizers provide varying outcomes in groundwater management through SO. This work raises at least two inquiries regarding simulation optimization in input uncertainty.

Firstly, how can we measure the extent to which the uncertainty in the input affects the outcomes of our optimization process for the management problem? Emphasis is placed on the decision variables utilized during the design and implementation stages. It is evident that each distinct decision variable, even from the same unknown distribution, yields a different optimal solution. Therefore, it is crucial to understand how to assess the outcome of this optimization and the extent to which it statistically differs from one another.

A second question related to the first is how to make second-order judgments or optimize the system by leveraging insights from the variability of these crucial decision variables.

There is limited literature on groundwater management through the S-O approach in the Rhône-Méditerranée region of France which has shown successively decline in GW levels and GW exchanges to rivers. The government has adopted strategies like Schéma Directeur d'Aménagement et de Gestion des Eaux (SDAGE) for holistic surface, ground, and coastal water management for socio-economic development. However, it is difficult to assess future challenges like groundwater depletion, insufficient supply of GW for irrigation, and river aquifer interaction for ecological development.

In this regard, the present study attempts to analyze the pattern and distribution of decision variables obtained from multi-objective optimization using statistical, explainable artificial intelligence (EAI) and visualization techniques and aid in the final decision-making of GW management problem. The study's targeted objectives include: (i) to explain the variability and distribution of DVs obtained from solving multi-objective optimization by varying (a.) optimization algorithms, (b.) GW model boundary, and (c.) DV range; (ii.) to visualize and analyze the factors in (a.) EAI(METHODS), (b.) statistical analysis and (c.) visualization responsible for DV pattern change; (iii.) to assess the spatial distribution and critical DVs

in the lower Ain basin, France and (iv.) multi-step decision making to identify optimal and new regions of sustainability and pumping by considering pareto front and its decision variables.

6.2 Methods

A total of 12 S-O models were developed in the PyGWMO platform. Three optimization algorithms, NSGA-II, MOPSO, and MOEA/D were used to check the efficacy and variation of decision variables on the optimization models. These algorithms were chosen because they are widely accepted for handling non-linear, complex functions in optimization. Based on the area-specific policies and accounting for aquifer parameter uncertainty, the groundwater extraction from pumping well zones was set to vary between 500 to 1500 and 500 to 2500 m^3/d. The study considered two groundwater model geometry domains to see the effect of boundary conditions on the optimal results.

The groundwater model was developed in MODFLOW, while changes in its boundary conditions and coupling with optimization models were achieved in Flopy (Hughes et al., 2024). Detailed discussion about GW model development has been discussed in chapter 3. The decision variables corresponding to each pareto set of nondominated solutions were studied with statistical analysis, including correlation analysis, Principal component analysis (PCA), self-organizing maps (SOM), and explainable artificial intelligence techniques (xAI). Critical variables were determined, which greatly impacted the optimization result. Finally, a spatial distribution of the decision variables was analyzed with the help of critical decision variables to identify the regions for environment protection and development zones.

6.2.1 Data preprocessing

All the pareto front data generated by solving multi-objective S-O were exported to individual CSV files, which were later combined to form a comprehensive dataset covering

various data points. There were 12 pareto fronts comprising (30 to 220) pareto solutions. The change in the number of solutions could be attributed to the different optimizers used. Moreover, every pareto point corresponded to 31 decision variables. Therefore, analyzing the multivariate relationship between the decision variables and output pareto optimal solutions becomes essential. To compensate for the low number of pareto solutions, 200 synthetic data points were created using kernel density estimation (KDE) by simulating observations for both predictor variables (pumping well zones) and target variables (objective functions). Epanechnikov kernel was chosen after experimenting with different kernel types. Synthetic data creation was limited to pareto fronts with less than 200 solutions. This was done to broaden the datasets and allow for more rigorous analysis. Statistical tests, such as the Anderson-Darling test, were performed to ensure that the synthetic data properly mirrored the statistical properties of the real dataset. Multicollinearity among the decision variables can cause unstable estimates of influence. To overcome this, each decision variable's variance inflation factor (VIF) was performed. The VIF values were utilized to detect any significant correlations or redundancies among the decision variables. The decision variables obtained by NSGA-II did not show much variation in their values nor any relation among themselves or with objective functions. Thus, they were omitted from the further analysis. The data was cleaned and standardized for further study. 1445 nondominated solutions were found with corresponding decision variables at the end of preprocessing.

6.2.2 Statistical analysis

6.2.2.1 Correlation analysis

The correlation analysis of decision variables provides essential insights into their interdependencies and relationships within the simulation-optimization framework. Calculating the pairwise correlation coefficients makes it possible to identify how strongly

or weakly decision variables influence one another. High correlations between a decision variable and an objective function suggest that the variable significantly influences achieving that objective, whereas low correlations indicate limited impact. A positive correlation between two decision variables demonstrates that they increase or decrease together, suggesting a degree of redundancy or potential for joint optimization. Conversely, negative correlations signify inverse relationships, where an increase in one variable leads to a decrease in the other. It is particularly relevant for dimensionality reduction techniques, as highly correlated decision variables can often be consolidated to streamline the optimization process without significant loss of information. Using Python's seaborn library, the correlation matrix of the decision variables was computed, representing the pairwise linear relationships among all decision variables. The matrix was aggregated to compute average correlations, ensuring the representation captures a robust and interpretable structure.

6.2.2.2 Principal Component Analysis

Principal Component Analysis (PCA) was employed to analyze and reduce the dimensionality of the decision variable space while retaining the maximum variance of the data. PCA is a widely used statistical technique for transforming correlated decision variables into a smaller set of uncorrelated variables known as principal components. The first step involves standardizing the decision variable dataset (X) to ensure that all variables are on the same scale, achieved using

$$Z = \frac{X - \mu}{\sigma} \tag{6.1}$$

where (Z) is the standardized data matrix, (μ) is the mean, and (σ) is the standard deviation. Subsequently, the covariance matrix Σ of the standardized data is computed using:

$$\Sigma = \frac{1}{n-1} Z^T Z \tag{6.2}$$

where n is the number of observations, the covariance matrix provides insights into how the decision variables vary concerning one another. Eigenvalues and corresponding eigenvectors of the covariance matrix are then calculated to identify the directions of maximum variance in the data. The eigenvectors represent the principal components, while the eigenvalues quantify the variance each component explains. The principal components are ranked in descending order of their eigenvalues, and the data is projected onto the first k-principal components to reduce dimensionality, where k is determined based on the cumulative explained variance (e.g., retaining 90% to 95% of the total variance). Here, the threshold of 85% was taken for all 12 pareto fronts. The transformation of the data allows for identifying the most influential decision variables driving the variation in the dataset. The application of PCA in this study provides a robust methodology to identify critical decision variables that significantly contribute to enhancing river-aquifer exchanges while minimizing redundancy and computational complexity in the subsequent optimization process.

6.2.3 Explainable Artificial Intelligence Techniques

xAI provides a thorough understanding of the hydrogeological parameters and critical decision variables in analyzing DV behaviour in the S-O framework. This work includes data generation, data generation/preprocessing, ensemble model development, and xAI techniques.

6.2.3.1 Ensemble model development

The xAI model aims to identify the S-O result affected by different well zones and explore the relationship between groundwater pumping and pareto front solutions. We utilized a bagging ensemble strategy to enhance the predictive accuracy and robustness of our pareto front prediction multi-output regression model. Bagging, another term for bootstrap

aggregating, creates several separate models using bootstrap samples from the original dataset. Every model in the ensemble was trained on a distinct bootstrap sample, which added variety to the training procedure. This diversity is crucial for minimizing overfitting and improving the model's generalization capacity. During the prediction phase, the bagging model aggregates the predictions of all different models by averaging their outputs to provide a robust and dependable forecast. Grid search was used to optimize the hyperparameters, and the best regressor (XgBoost) was chosen based on cross-validation score. XgBoost model was evaluated on a test dataset, producing accurate results. A bagging ensemble was utilized to limit the influence of outliers, decrease model variance, and enhance the accuracy and stability of our objective function value forecasts.

6.2.3.2 SHAPely and LIME analysis

In this study, we utilized xAI methods, particularly SHAP (Lundberg and Lee, 2017) and LIME analysis (Núñez et al., 2023), to understand the connection between discharge from each pumping well zone and the two objectives of maximizing R-A exchanges and total discharge. We utilized SHAP analysis to determine the significance and influence of each variable on the pareto optimal solutions. Based on cooperative game theory, SHAP values measure each decision variable's impact on the predictions of an objective function. They indicate the incremental effect of DV on the expected value of the objective function. The SHAP value of a feature xi for a prediction f(x) is given as:

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} [f(S \cup \{i\}) - f(S)] \qquad (6.3)$$

Where, ϕ_i Is the SHAP value for the feature i.

F is the set of all features.

S is any subset of F excluding feature i.

$f(S)$ is the model output (prediction) using the subset S.

$f(S \cup \{i\})$ is the model output, including the feature i. $|S|!$ and $(|F| - |S| - 1)!$ are factorial terms ensuring fair attribution.

The SHAP summary map offers a thorough overview of the variable contributions and aids in quantifying the sensitivity and importance of each predictor, showing non-linearities within the decision variable ranges.

LIME, conversely, creates localized explanations of model predictions by simplifying large models with more interpretable ones focused on instances of decision variables. It identifies the key decision variables that have the biggest impact on forecasting the value of the objective function at a certain location. For an instance x, LIME solves the following optimization problem:

$$\text{argmin}_{g \in G} \mathcal{L}(f, g, \pi_x) + \Omega(g) \quad (6.4)$$

Where, g is the local surrogate model (e.g., a linear model).

G is the class of simple models.

$\mathcal{L}(f, g, \pi_x)$ measures how well g approximates the complex model f in the local region defined by π_x.

π_x is a locality kernel that assigns weights to samples based on their proximity to x.

$\Omega(g)$ penalizes the complexity of the surrogate model g.

The pareto front-to-decision variable dataset was trained based on an XGB regressor with 10000 trees in the ensemble model and a learning rate 0.01. Early stopping was used to determine the optimal number of estimators dynamically. LIME generates synthetic data by perturbing the input instance, evaluates the model's predictions on these perturbed samples, and fits a simple interpretable model (in this case, a linear regression) to approximate the model's behavior locally around the instance. The weights of this local model indicate the importance of each feature for the prediction. xAI approaches were used to improve the interpretability and robustness of the study, helping comprehend the critical decision variables that influence the objective functions.

6.2.4 Generalized Additive Modeling

Generalized Additive Models (GAMs) are statistical tools that model complex, non-linear relationships between dependent and independent variables. GAMs extend traditional linear models by replacing linear terms with smooth functions, enabling them to capture intricate patterns in data. The functional form of a GAM is:

$$y = \beta_0 + f_1(x_1) + f_2(x_2) + \cdots + f_p(x_p) + \epsilon \quad (6.5)$$

Where:

- y is the dependent variable,
- β_0 is the intercept term,
- $f_i(x_i)$ represents smooth, non-linear functions of the predictor variables x_i,
- ϵ is the error term.

GAMs are particularly useful for identifying the non-linear effects of predictors and their partial contributions to the outcome variable without assuming a specific functional form. In optimization studies, GAMs are valuable tools for analyzing and interpreting the influence of decision variables on objective functions, especially in contexts involving multi-objective optimization.

In this study, GAMs were applied to examine the relationships between groundwater pumping decisions, represented by 31 decision variables across spatial zones, and the two primary objective functions. Two independent GAMs were constructed to model the relationships; GAM 1 captures the effect of decision variables on objective function 1, while GAM 2 captures the impact of decision variables on objective function 2.

The smooth functions for each decision variable and $f_i(x_i)$were fitted using spline-based smoothing techniques implemented with the pygam library. Cross-validation optimized the smoothing parameters, ensuring model generalizability and preventing overfitting. Partial Dependence Plots (PDPs) were generated for each decision variable to isolate and visualize

their contributions to the objective functions. PDPs allow for the quantification of the average effect of each decision variable while holding all other variables constant. This step provided insights into the sensitivity of the objective functions to specific decision variables and revealed potential non-linear patterns.

Finally, a spatial hotspot analysis was done using the Local Getis-Ord Gi* statistic to identify regions with significant clustering of values for four criteria: SHAP values (SHAP_abs), LIME values (LIME_abs), variable loadings from PCA, and net partial dependency scores (dependen_1). The analysis is conducted on a geospatial dataset loaded from a shapefile, ensuring valid geometries and reprojecting to a suitable CRS for spatial analysis. A k-nearest neighbors spatial weights matrix (with k=3) is created to capture spatial relationships between features. The Local Getis-Ord Gi* statistic is computed for each criterion, generating Z-scores and p-values to identify statistically significant clusters (hotspots and cold spots).

6.3 Results and discussions

6.3.1 Correlation Analysis and Variable Loading

The resulting violin plot in Figure 6.1 highlights the distribution and variability of correlations for each of the 31 decision variables. The y-axis reveals a range of correlation values, with most variables showing moderate variability around the zero-correlation baseline. A few decision variables exhibit a wider spread, such as variables 3, 10, 19, and 25, indicating higher variability and potential outliers within their correlation structures. This spread suggests that these variables have a stronger interaction or influence on the system, which could reflect their significance in groundwater exchange processes. Some decision variables 4, 12, and 22 notably show tightly clustered distributions around zero, signifying weaker interdependencies. Such insights are crucial for identifying decision

variables that contribute most to the variability in groundwater management scenarios and for prioritizing them in further optimization processes.

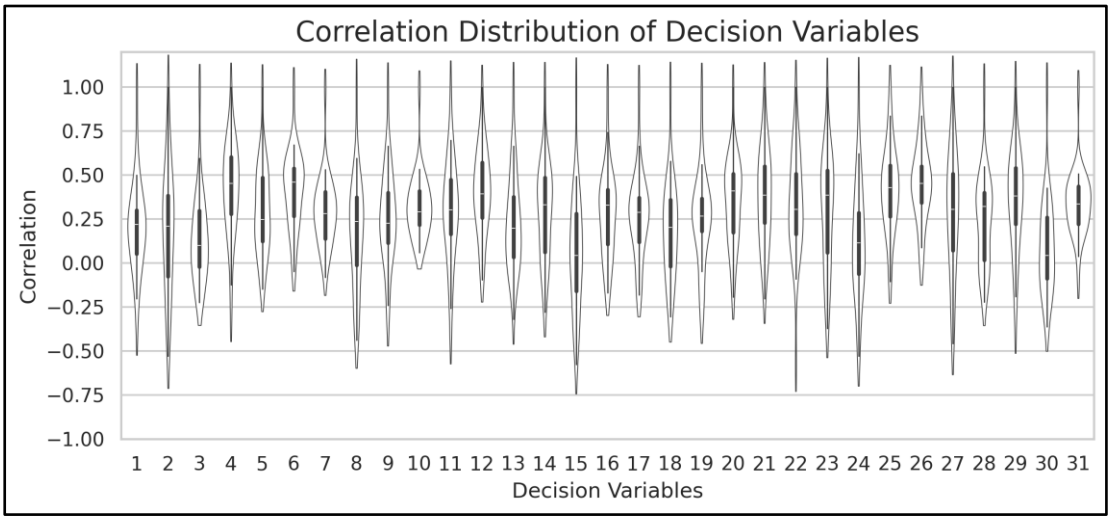

Figure 6.1 Correlation distribution among decision variables

Positive correlations with Objective 1, can be seen predominantly in variables 3, 7, 14, 19, 25, and 30, suggest their significant role in enhancing river-aquifer exchanges (Figure 6.2). Conversely, strong negative correlations with Objective 2, observed in variables 4, 9, 13, and 22, indicate their critical impact on groundwater discharge.

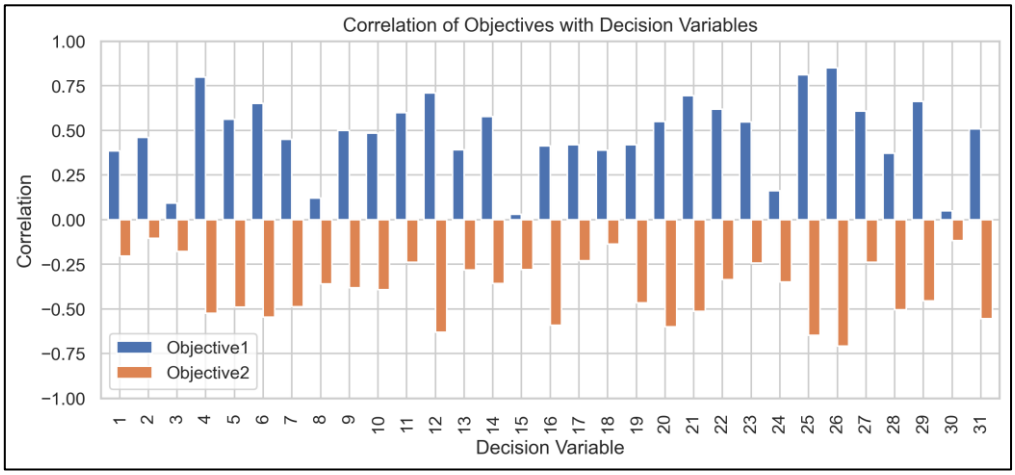

Figure 6.2 Correlation between decision variables and objective functions

The dendrograms in clustered heatmap of correlations in Figure 6.3 represent how the rows and columns have been clustered hierarchically. Notably, there are distinct clusters forming

indicating different groups of correlations between decision variables and objective functions.

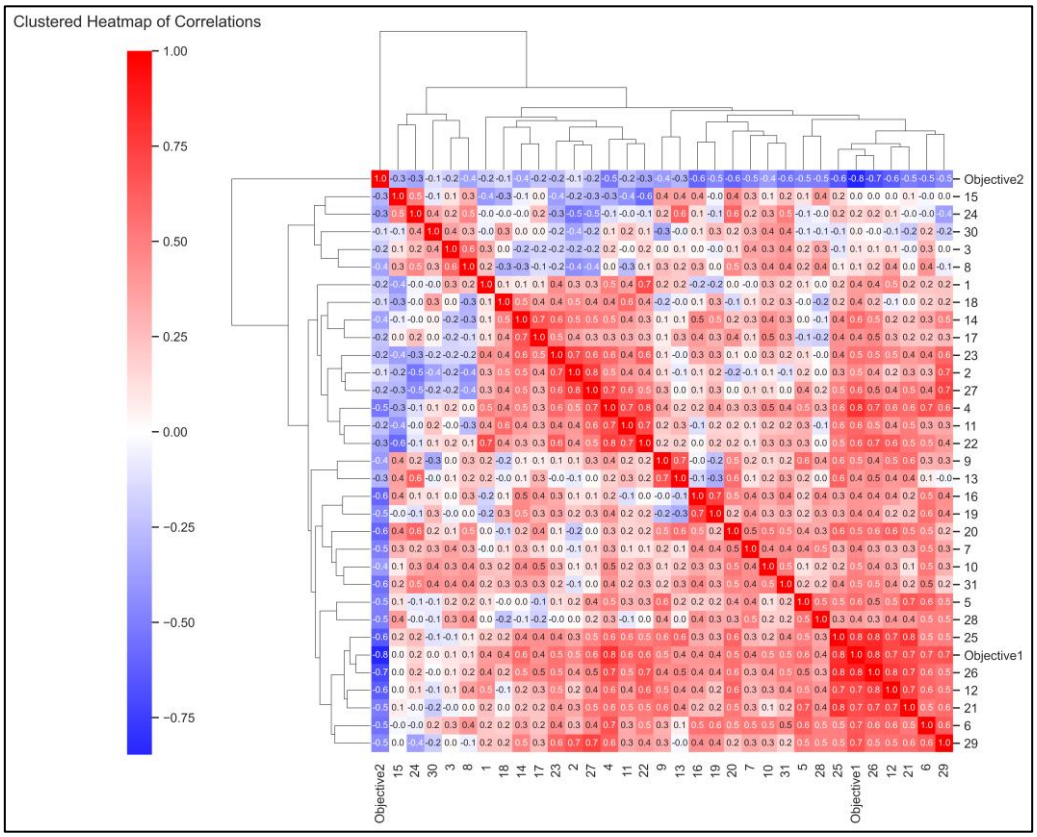

Figure 6.3 Clustered heatmap of correlations

The variable loading plot in Figure 6.4 from the PCA analysis illustrates the contribution of each decision variable to the first two principal components (PC1 and PC2). The length and direction of the arrows indicate the magnitude and influence of the variables on the principal components. Variables such as 15, 8, 25, and 3 exhibit strong positive loadings on PC2, suggesting they are key contributors to variations captured by this component. On the other hand, variables like 4, 22, and 11 show strong negative loadings along PC1, indicating their significant influence in explaining the variance in that direction. Variables closer to the origin, such as 18 and 1, contribute minimally to both PCs, reflecting their weaker impact on the dataset's variance.

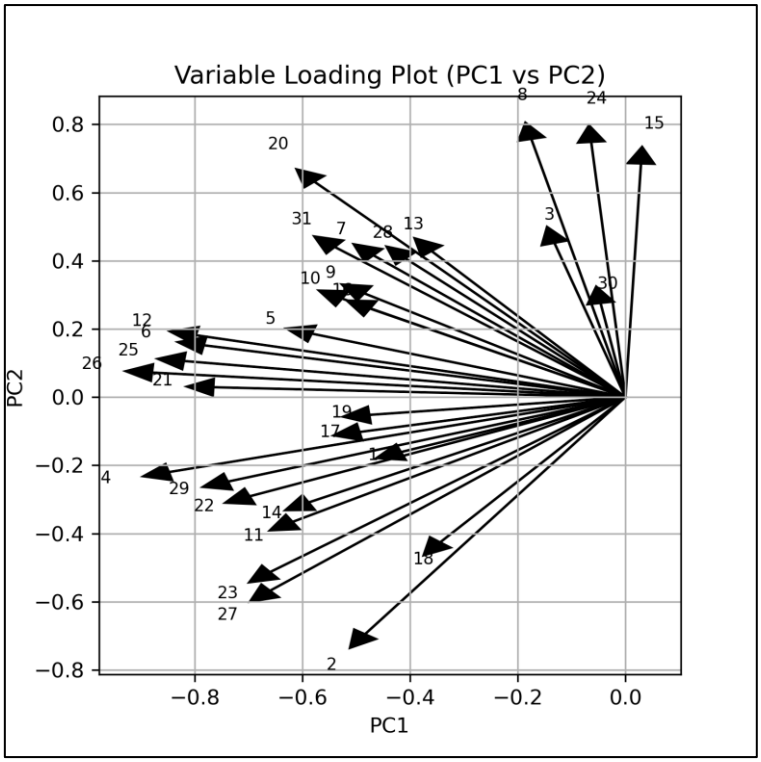

Figure 6.4 Variable loading plot using PCA of decision variables

6.3.2 Variable importance

The LIME analysis provided evaluation of the relationship between the 31 decision variables and the two objectives. The training and validation loss is greatly minimized representing a relation between input and output. The model loss (RMSE) reached to 0.01 at the end of 400 epochs. The loss is represented in Figure 6.5(A). The part B of visualization highlights the mean absolute weights of the decision variables, revealing the most influential ones in shaping the outcomes. Variables such as zones 4, 25, and 26 exhibit significantly higher mean absolute weights, underscoring their critical roles in achieving the objectives.

The part C of visualization delves into the LIME weight distribution, illustrating the contribution of each variable across various samples. The polarity and magnitude of these weights provide insights into whether the variable promotes or restricts the objectives. Notably, zones 4, 25, 26, and 9 consistently demonstrate dominant influences, making them

pivotal decision points. These critical zones should be prioritized in management strategies for achieving optimal groundwater extraction and R-A exchanges.

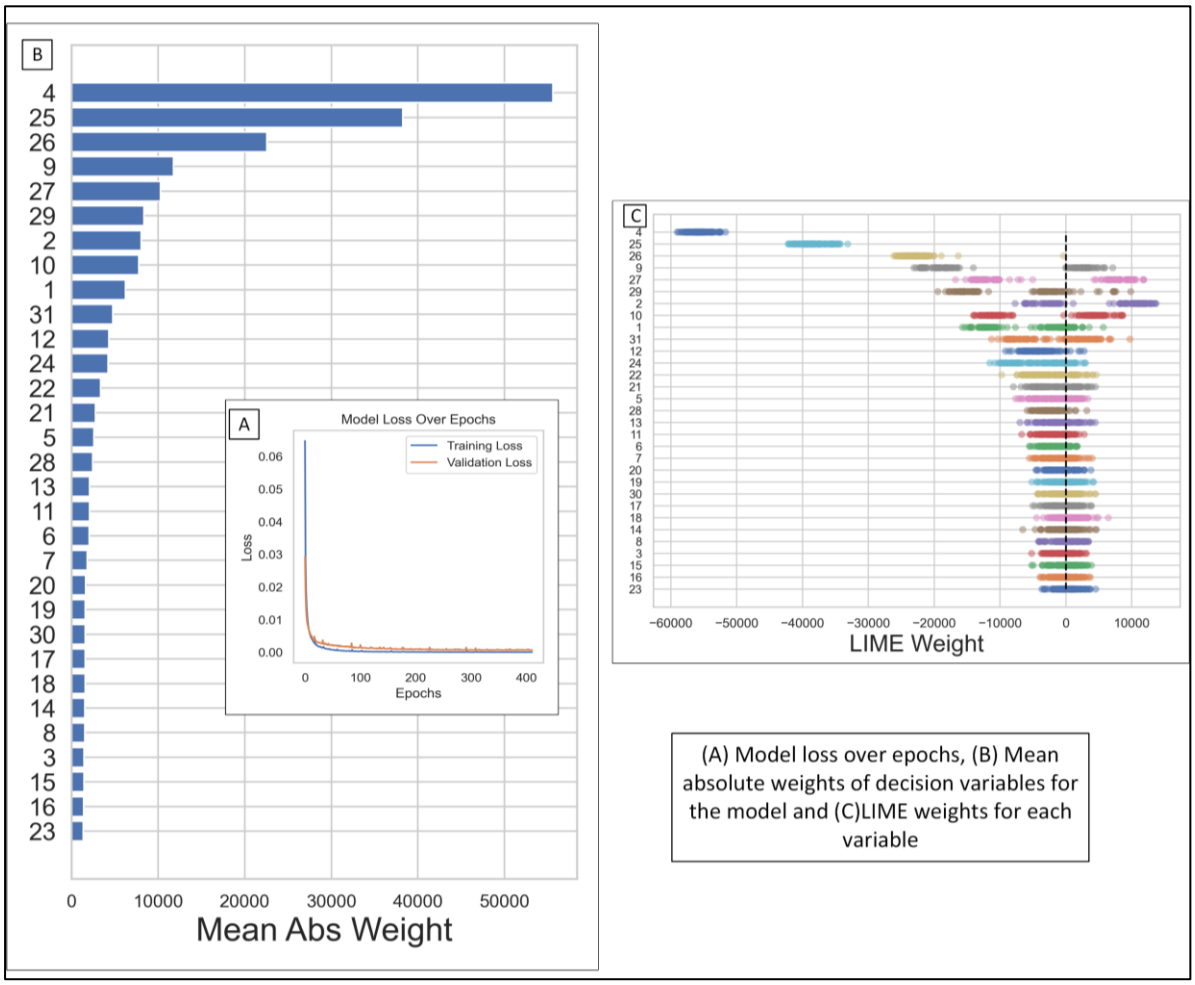

Figure 6.5 (A) Training and validation loss over the epochs (B) mean absolute weights for all the variables in the algorithm and (C) LIME weights

SHAP is a game-theoretic approach to explaining the output of machine learning models by attributing the prediction to the individual contribution of each feature. The attached SHAP summary plot in Figure 6.6 showcases the influence of the decision variables (representing groundwater pumping zones) on the objective function of maximizing recharge-to-aquifer exchanges and groundwater extraction. In the visualization, each dot represents a data point, and its position along the x-axis reflects the SHAP value, which indicates the magnitude and direction of the variable's impact on the model output. Positive SHAP values signify a positive contribution toward achieving the objective, while negative

values denote a negative effect. The color gradient, ranging from blue (low feature value) to red (high feature value), highlights the feature values' variation.

The spread of dots indicates variability in the influence of a variable across different samples, capturing both the interaction effects and the feature importance. Notably, decision variables such as zones 12, 4, and 25 exhibit substantial variation in SHAP values for objective 1, emphasizing their critical role for objective two variables such as 4, 12, 25, and 16 exhibit a significant influence, as evident from their wide spread of SHAP values. The variation within each variable demonstrates the heterogeneity of its impact across the dataset. In contrast, features like 30, 28, and 3 show a comparatively minor influence, as indicated by their smaller SHAP value ranges. This analysis highlights the importance of prioritizing key variables 4, 12, 25, and 16 for effectively managing groundwater extraction strategies.

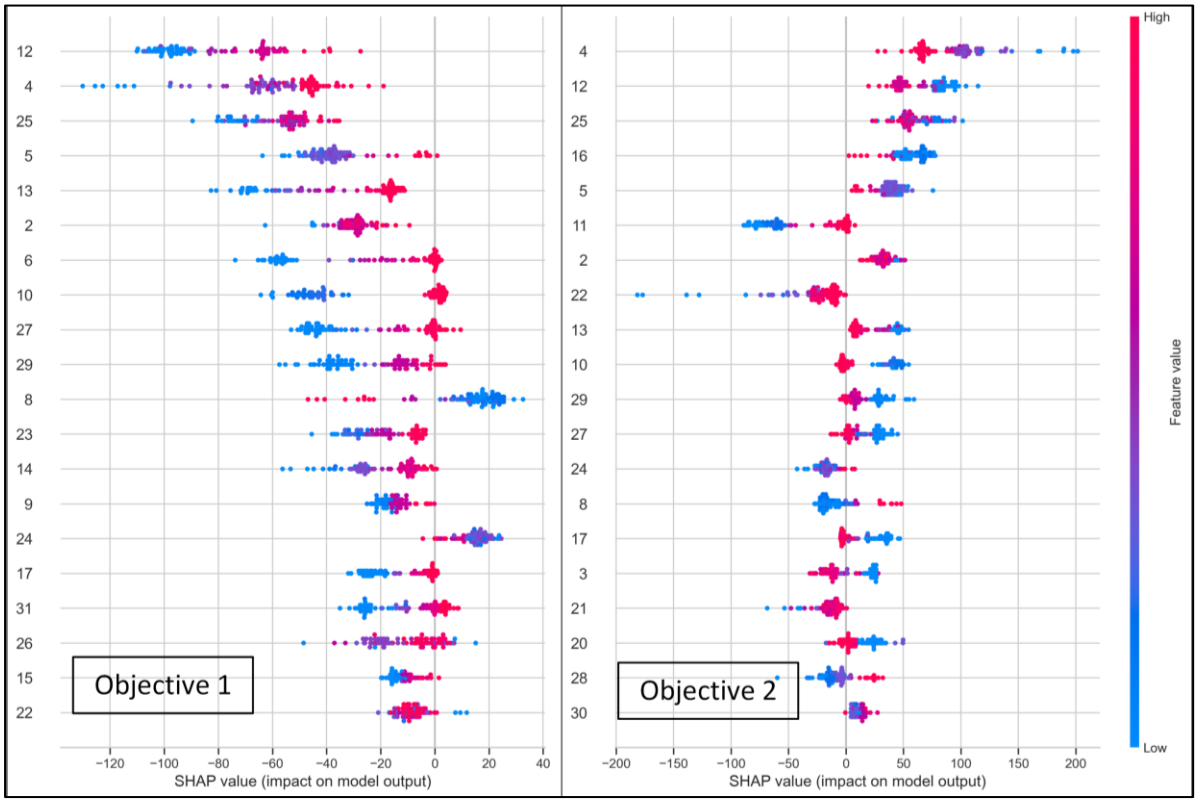

Figure 6.6 SHAP value plot for each decision variable in two objectives

6.3.3 Variable dependency

The partial dependence plot in Figure 6.7 generated using GAM reveals the average impact of decision variables (zones of groundwater pumping) on the two objectives: maximizing R-A exchanges and maximizing groundwater extraction. The oscillating patterns across the variables suggest non-linear relationships between decision variables and objective functions, underscoring the appropriateness of GAMs for capturing these dynamics. The graph reveals significant variability in the partial dependence across decision variables. Certain variables 4, 10, and 28 positively or negatively impact one or both objective functions, indicating their critical influence on the optimization process. The impact on the two objective functions is opposing for some decision variables. For example, 10 exhibits a strong positive partial dependence on objective two but is neutral or slightly negative for objective 1. This highlights potential trade-offs in decision-making, as optimizing one objective might adversely affect the other. Variables 4, 10, and 28 emerge as influential, suggesting that targeted interventions in these zones could yield significant improvements in R-A exchanges or groundwater discharge.

Conversely, other decision variables, such as 12 and 16, have relatively negligible partial dependence values, implying minimal influence on the objectives. The analysis provides actionable insights for prioritizing certain zones (decision variables) to achieve desired outcomes.

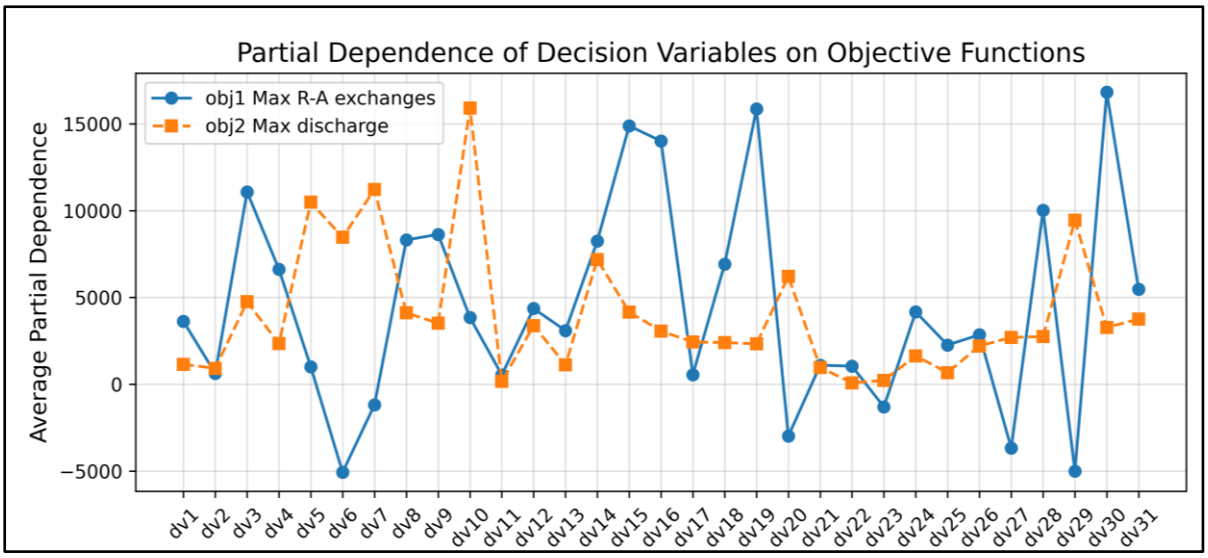

Figure 6.7 Partial dependency plot of decision variables for each objective.

6.3.4 Environmental and Development Zones

Figure 6.8 represents the hotspot analysis using the Getis-Ord spatial clustering method. It identifies regions within the basin that are significant in terms of SHAP absolute values, representing critical decision variables influencing stream-aquifer interactions. The map highlights areas with high (red) and low (blue) spatial clustering of SHAP values, denoting strong and weak influence zones, respectively. Environmental zones, characterized by low SHAP values, align with blue clusters, suggesting limited groundwater extraction to preserve river-aquifer exchanges. In contrast, development zones, marked by high SHAP values (red clusters), are associated with significant groundwater pumping activities. The central and southern regions demonstrate a mixture of hotspot clusters, indicating areas requiring balanced management strategies to mitigate environmental impacts while meeting groundwater demands. The environmental zones are clustered together. While, in LIME based analysis, the environmental zones are scattered all around the model. In variable loading-based clustering, the middle region near the confluence of Albarine, to the western edge comes under environmental zones, while the faraway wells represent the development zones.

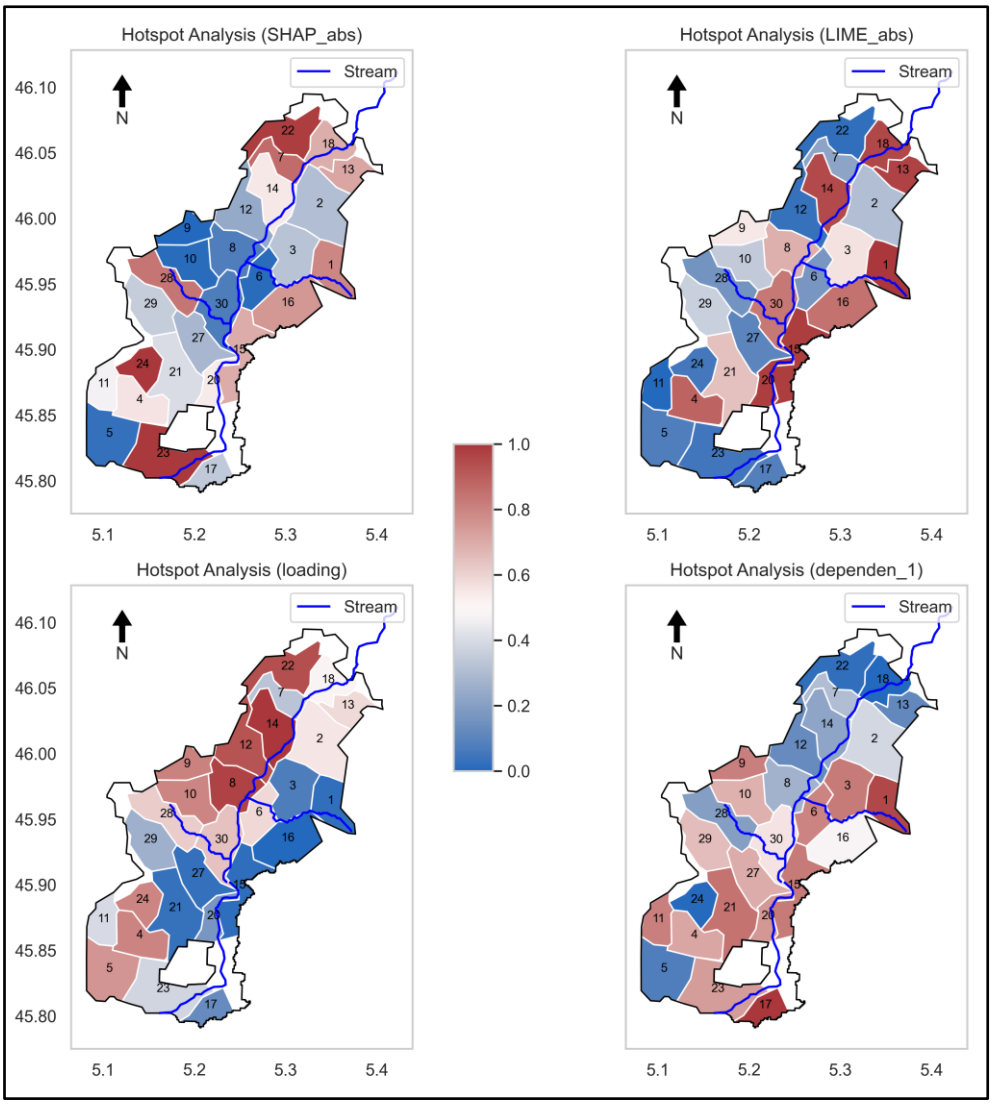

Figure 6.8 Hotspot analysis based on SHAP, LIME, variable loading, and partial dependency

6.4 Summary

This chapter delves into the exploration of decision variable spaces within groundwater simulation-optimization (GW S-O) models and evaluates their impact on the optimal results. The integrated S-O framework produced 12 Pareto sets for various combinations of algorithms, spatial regions, and variable bounds. Initially, correlation analysis and principal component analysis were employed to identify the most influential decision variables. Subsequently, explainable artificial intelligence techniques, such as SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations), were utilized to highlight critical decision variables, with variables 4, 25,

26, and 9 consistently demonstrating significant influence on the objective functions. To refine the selection of critical decision variables, Generalized Additive Modeling was applied, revealing additional influential variables, such as 10, 31, 19, and 6, as evident in the partial dependency plots. Integrating insights from these analyses, a hotspot analysis was conducted to delineate environmental zones (requiring groundwater extraction restrictions to maintain river-aquifer exchanges) and development zones (characterized by high groundwater pumping). This comprehensive approach provides a robust framework for managing groundwater resources sustainably.