An Evolutionary Computation Approach for Optimizing Support Vector Regression in Predicting Pollen Levels

Tiberiu Alexandru MIHAI[1,2], Smaranda BELCIUG[3]*

[1] School of Advanced Studies of the Romanian Academy, 125 Calea Victoriei, Bucharest, 010071, Romania

[2] "Simion Stoilow" Institute of Mathematics of the Romanian Academy, 21 Calea Grivitei, Bucharest, 010702, Romania
tiberiu.mihai@gmail.com

[3] Department of Computer Science, University of Craiova, 13 A. I. Cuza St., Craiova, 200585, Romania
sbelciug@inf.ucv.ro (**Corresponding author*)

Abstract: Support Vector Regression (SVR) is a machine learning technique widely used for modeling complex and non-linear data. Its performance is highly sensitive to the choice of the kernel and hyperparameters. This study explores the use of an evolutionary computation approach for optimizing the SVR parameters, namely a systematic framework that balances predictive accuracy, computational efficiency, and generalization. The proposed methodology optimizes the key SVR parameters - including the regularization constant (C), the kernel coefficient (γ), and the ε-insensitive loss parameter - while simultaneously exploring different kernel families. Unlike the conventional grid search optimization, the evolutionary algorithm dynamically refines the candidate solutions through selection, crossover, and mutation. In order to evaluate this method, it is applied to forecasting airborne pollen counts, a task which is complicated by variability, seasonality, and the non-linear dynamics of aeroallergen data. The experimental results demonstrate that kernel adaptation through evolutionary computation leads to a consistently improved accuracy and efficiency. Overall, this approach enables the construction of SVR models that are both more accurate and computationally efficient than those produced by using traditional optimization methods.

Keywords: Support vector regression, Evolutionary computation, Statistical analysis, Airborne pollen counts.

1. Introduction

Worldwide, between 10 and 30% of the population suffers from allergic rhinitis, also known as hay fever. In Europe, around 40% of the people are sensitized to different pollen allergens, while in the United States, pollen affects more than 60 million people each year (25.7% of the adults and 18.9% of children), (U.S. Centers for Disease Control and Prevention, 2023; European Climate and Health Observatory, 2025). It is clear that airborne pollen is no longer a seasonal inconvenience, it represents a trigger for symptoms that can range from mild distress to life-threatening respiratory distress. Besides the physical burden, the economic consequences are considerable: the United States healthcare costs related to pollen allergies account for $ 18 billion each year, while the European Union healthcare costs exceed 282 billion/year (Luengo-Fernandez et al., 2023).

It is crucial to be able to forecast the pollen levels. High pollen counts worsen the allergies and asthma symptoms leading to an increased use of medication, and spikes in the emergency care units. If the pollen levels could be forecast then different preventive measures could be undertaken, such as reducing the duration of outdoor activities, using air filters, or starting antihistamine medication in advance. This way, while parents would be able to limit or adjust their children's outdoor playtime, the healthcare providers would be able to make room in their consultation schedule for allergy patients. In this manner, the surges in demand can be forecast in due time, helping hospitals manage their resources.

In the last decade, there have been many scientific attempts to forecast pollen levels. For instance, in (Kitinoja et al., 2020), the authors provided a systematic review and meta-analysis in which they demonstrated that even small increases in airborne pollen concentrations lead to statistically significant increases in allergy and asthma symptoms, with predominantly respiratory and ocular symptoms. The study shows that a 10-grain/m^3 increase in pollen concentration is associated with a 2% increase in any allergy or asthma symptoms, a 7% increase in upper respiratory tract symptoms, and a 11% increase in ocular symptoms. Oh (2022) showed that rising temperatures and elevated CO_2 levels play a crucial role in climate change, which leads to longer pollen seasons and consequently to a higher allergenicity making accurate monitoring and prediction vital for anticipating health crises and risks. Thus, forecasting pollen levels represents not only a medical issue but also a societal and environmental one. Landesberger et al. (2023) tested a mobile application that forecasts pollen levels. The app was found to be a valuable tool for managing allergic rhinitis and asthma. The role

https://doi.org/10.24846/v34i3y202507

of community-level data networks in managing allergic conditions is emphasized in (Dwaraknath et al., 2024).

It is a fact that pollen level forecasting is not a luxury, but a necessity. As the climate changes, forecasting will be even more important, since it bridges ecological monitoring and human health. Using Artificial Intelligence (AI) enhances the experts' ability to forecast pollen concentrations, making a true difference in public health and allergy prevention. Different AI methods have proven to be effective in recognizing patterns in environmental data, offering short-term pollen predictions. For instance, Lops et al. (2020) used a convolutional neural network to forecast pollen counts one to seven days in advance. Depending on the pollen type, the model achieved high correlation coefficients, up to 0.90, demonstrating its potential in empowering people with allergies. Makra et al. (2024) used CatBoost and deep learning models to predict the pollen concentration up to 14 days ahead. Through their study they showed that the temperature and soil moisture are the most important predictors of pollen levels. In (Astray et al., 2025), the authors focused on a specific allergenic genus, Parietaria pollen, using different artificial intelligence models such as Random Forests, Support Vector Machines, and Artificial Neural Networks. The employed models achieved high correlations ranging between 0.713 and 0.859, with a low root squared error (5.55-7.66 grains/m^3) for a forecast 1-3 days ahead. In the work of Corić et al. (2023), a recurrent neural network called PollenNet was designed for forecasting daily airborne pollen (ragweed, grass, and birch pollen) using historical data and meteorological inputs. AIpollen, a deep-learning tool based on ResNet34 was introduced by Yu et al. (2024). The model is able to achieve a 97% accuracy in identifying pollen form image data.

Recent advances in applied Artificial Intelligence show the growing importance of robustness, optimization, and explainability across different application domains. One study extended the use of federated learning systems by combining them with explainability tools and showing this integration not only improved prediction accuracy, but also ensured transparency and data privacy (Ribeiro et. al., 2024). Another contribution focused on forensic deception detection with Support Vector Machine, Neural Networks, and Random Forest (Rad et al., 2024).

The aim of this study is to automatically optimize the SVR hyperparameters (e.g. $C, \gamma, \varepsilon,$ and the kernel) for predicting the pollen levels. For benchmarking purposes, the novel model's results were compared with the results obtained by the stand-alone SVR model that uses Grid Search optimization.

The remainder of this paper is organized as follows. Section 2 presents the design and implementation of the proposed approach, while section 3 discusses the dataset, and the design of experiments. Section 4 presents the results and their discussion. Finally, Section 5 concludes this paper and outlines possible future research directions.

2. Model Design

2.1 Support Vector Regression

Support vector machines (SVMs) are a class of supervised artificial intelligence models that can be used for classification or regression purposes. For classification tasks they identify the optimal hyperplane that separates the data points in different classes. The data points that surround the hyperplane are the support vectors that define the decision boundary. For regression tasks, the SVR is used, which extends the SVMs principles from classification to regression. SVRs do not search for a hyperplane to separate classes; they build a regression function that approximates the ground truth values within a tolerance margin. The regression model is kept as simple as possible.

The tolerance margin or zone is called the $\varepsilon-$ insensitive tube. If the prediction is within the interval $(-\varepsilon, \varepsilon)$ then it does not count as an error in the forecast. However, if the prediction is outside this interval, then the model has made an error in the prediction, and it is penalized by using the slack variables (ξ, ξ^*) that measure the distance of the error from the $\varepsilon-$ insensitive tube.

From a mathematical point of view a SVR function is defined as follows. Let $X \subset \mathbb{R}^n$ be a set. Given a training dataset $(x_i, y_i) \in X \times \mathbb{R}$, $i = 1, 2, \ldots, m,$ a SVR function $f: X \to \mathbb{R}$ is considered:

$$f(x) = w^T \cdot \phi(x) + b \tag{1}$$

Here $x_i \in X$ denotes the input vectors, y_i denotes the scalar outputs, and m denotes the size of the dataset. $\phi: X \to \mathbb{R}^d$ is a nonlinear function called the kernel function, and the vector $w \in \mathbb{R}^d$ denotes the weight vector.

The goal is to find the optimum SVR function that deviates from the targets y_i by at most ε.

In general, with regression tasks, the tendency is to minimize the absolute error. In the case of SVR, the $\varepsilon-$ insensitive loss function is minimized.

$$L_\varepsilon\left(y, f\left(x\right)\right)=\begin{cases}0, if \mid y-f\left(x\right)\mid \leq \varepsilon \\ \mid y-f\left(x\right)\mid -\varepsilon, otherwise\end{cases} \quad (2)$$

The following optimization problem arises:

$$\min_{w,b,\xi,\xi^*}\frac{1}{x}\|w\|^2+C\sum_{i=1}^{m}\left(\xi_i+\xi_i^*\right) \quad (3)$$

which is subject to:

$$\begin{aligned} &y_i-\left(w\cdot\phi\left(x_i\right)+b\leq\varepsilon+\xi_i\right), \\ &\left(w\cdot\phi\left(x_i\right)+b\right)-y_i\leq\varepsilon+\xi_i^*, \\ &\quad \xi_i,\xi_i^*\geq 0, i=1,...,m \end{aligned} \quad (4)$$

where:

- $\phi(x)$ is the kernel function;
- ε is the tolerance margin;
- $C>0$ is the regularization parameter that controls the trade-off between the model's tolerance for errors and flatness;
- ξ_i, ξ_i^* are the slack variables for under- or overestimation.

By applying the standard Lagrangian duality approach to the optimization problem (equations (3) and (4)), as described in (Cortes & Vapnik, 1995), the SVR prediction function can be expressed in its dual form as:

$$f\left(x\right)=\sum_{i=1}^{m}\left(\alpha_i-\alpha_i^*\right)K\left(x_i,x\right)+b \quad (5)$$

where $K(x_i, x)$ is the kernel function that implicitly computes the inner product in the transformed feature space.

2.2 Evolutionary Computation Approach for SVR

Evolutionary computation (EC) is an optimization algorithm inspired by the natural evolution. It uses a population of candidate solutions that evolves toward a better performance. The evolution process uses biologically inspired operators such as selection, crossover, and mutation.

The aim of this study is to use evolutionary computation for optimizing the SVR method. Equation (5) represents the standard prediction function in SVR derived from the dual optimization problem. In this study, the goal is not to re-derive in the SVR formulation, but to optimize its hyperparameters using evolutionary computation. With this purpose in mind, a genome g with mixed genes, that is categorical and continuous ones was encoded:

$$g=[C,\gamma,\varepsilon,kernel] \quad (6)$$

Depending on the kernel function, some genes become irrelevant. For instance, for a linear kernel, γ is ignored. The evolutionary algorithm searches this mixed representation to identify the parameter settings that maximize predictive performance, with equation (5) serving as the regression function evaluated during fitness computation.

The search space range for each gene and the optimal range found are presented in Table 1.

Table 1. Search space range for genes

Gene	Search Range	Optimal range	Evaluations
C	[0.001, 1000]	20-80	250 000
γ	[1e-6, 10]	0.001-0.01	100
ε	[0.001, 1.0]	0.08-0.2	
kernel	Linear, polynomial, RBF, sigmoid	RBF (67%)	

At first, an initial population of N candidates: $P^{(0)}=(g_1^{(0)}, g_2^{(0)}, ..., g_N^{(0)})$ is randomly generated from the search domain. Each candidate is afterwards evaluated using the fitness function F, which is the SVR function.

A new generation $P^{(t+1)}$ is produced from $P^{(t)}$ at each generation t, using the variation operators. In this study, the following operators were used:

- *Selection*: the candidates that have lower fitness function values are favored and are selected (tournament selection);
- *Crossover* (total arithmetic): if there are two parents g^a, g^b, then two offsprings can be created by exchanging parts of their genome:

$$g^1=\alpha g^a+\left(1-\alpha\right)g^b, \alpha\in\left[0,1\right] \quad (7)$$

$$g^2=\alpha g^b+\left(1-\alpha\right)g^a, \alpha\in\left[0,1\right] \quad (8)$$

- *Mutation*: diversity in the population is created by randomly perturbing genes. If a

gene is categorical then one switches from 0 to 1, and vice versa. If the gene is continuous then a random value is added or subtracted:

$$g_j^{'} = g_j \pm \sigma \quad (9)$$

EC uses elitism, which means that the top candidates E are copied directly to the next generation. The algorithm stops when reaching a predefined number of generations or the convergence.

To ensure that the evolutionary search explored only valid configurations, constraints were applied to the genome. For instance, γ was considered only when the kernel was a RBF or polynomial kernel and ignored when it was linear. Polynomial kernels are restricted to valid degree values, and all continuous parameters (C,γ,ε) are kept within predefined numeric bounds (Table 1). These constraints prevented the generation of unrealistic parameter combinations and improved convergence stability.

The SVR/EC algorithm can be summarized as follows:

Algorithm. SVR/EC

1. Encode the candidate solutions;
2. Initialize a population of N individuals by randomly generating it within the predefined ranges;
3. For each candidate solution, train a SVR model with those hyperparameters (candidate solutions);
4. Measure predictive performance using 10-fold cross-validation;
5. Compute error metrics such as RMSE, MAE, and R^2;
6. Select parents for reproduction. Choose the individuals that have better fitness values;
7. Apply crossover;
8. Apply mutation;
9. Form the next generation;
10. Repeat for a fixed number of generations G, or until no improvement in fitness is observed for several iterations;
11. Return the best solution;
12. Retrain the SVR model using those hyperparameters and record the performance metrics together with the support vectors.

3. Dataset and Design of Experiments

3.1 Dataset Description

In 2021, the National Atmospheric Deposition Program (NADP) in collaboration with the U.S. Geological Survey, University of Wisconsin, PollenSense Inc. Aerobiology Research Laboratories, and the National Allergy Bureau, started a pollen-monitoring project. The project was carried out in three core monitoring locations: Duke Forest, North Carolina, the University of Wisconsin, Wisconsin, and Utah State University Agricultural Research Station, Utah.

The data was gathered using traditional microscopy (samples were identified and counted under microscopes), automated real-time monitoring (sensors collected continuously 1-minute images of airborne particles which were further on classified by AI models), wet deposition sampling (composite precipitation samples were filtered and stained and pollen grains were identified microscopically), and high-volume air samplers (filter captures airborne particles) (Wetherbee, 2022).

The data was gathered between February 2 and November 2, 2021. The following types of data were collected: total and taxa-specific pollen counts (tree, grasses, weeds, Acer, Quercus, Poaceae, Ambrosia etc.), fungal spores (Cladosporium, Alternaria, Aspergillus, etc.), non-biological particles (smoke, ash, dust, microplastic particles), precipitation data linked to pollen deposition (depth, volume, catch efficiency). Due to the fact that the data regarding pollen was the most consistent, this study focused on determining the pollen count.

3.2 Analysis Aim

This study's aim was to find whether the performance of a traditional SVR optimized with Grid Search Optimization can be improved by changing the optimization method with Evolutionary Computation. An important part of this study was the benchmarking process that compared the performances of the two models on the same dataset. As a validation method the 10-fold cross-validation was used. Each model was executed in 100 independent computer runs of a complete cross-validation. The number of computer runs was chosen after performing a power analysis, which revealed that for a statistical

power greater than 95% type I error $\alpha = 0.05$, the algorithms should be run 100 times. For each model the following were computed:

- *Root Mean Square Error* (RMSE), $\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}$, which penalizes large error heavily, and is more sensitive to outliers;
- *Mean Absolute Error* (MAE), which is defined as:

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|.$$

It measures the average magnitude of errors between the predicted values ($\hat{y}_i$) and the actual values (y_i), but it does not take into account the direction of the errors.

- *Coefficient of Determination*

$$R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}{\sum_{i=1}^{n}\left(y_i - \overline{y_i}\right)^2}, \overline{y_i} = \sum_{i=1}^{n} y_i$$

is the mean of the observed values. If its value is close to 1, then the model fits better. It computes how well the model explains the dependent variable's variance.

- *Support vectors (SVs)* represent the subset of training data that lie closest to the decision boundary. The SVs determine the position and orientation of the regression tube. The lower their number, the better the generalization and efficiency of the model.

4. Results

The first part of the experiment focused on determining the best SVR+EC model. For this, the performance of each candidate solution model was evaluated in terms of RMSE, MAE, R^2, and SV. The performance summary statistics in terms of mean, standard deviation (SD), minimum and maximum values, is displayed in Table 2.

Table 2. Summary Performance Statistics

Metric	Mean	SD	Min	Max
RMSE	33.456	3.234	28.123	42.789
MAE	25.234	2.567	21.456	32.123
R^2	0.804	0.042	0.7234	0.8789
SV	81.3	6.7	68.2	97.4

From Table 2, it can be seen that the model achieves a good accuracy (a low MAE and RMSE, with a high R^2). The errors seem to be generally small relative to the scale of pollen concentrations measured. The model uses approximately 80 SVs, which suggests it balances complexity and generalization.

Next, the statistical behaviour of the evolved parameters (C,γ,ε) is analysed (Table 3).

From Table 3, it can be seen that C shows the highest variability, which ultimately reflects the difficulty in balancing margin width and error penalty. γ and ε are more stable, which indicates that the kernel's shape and tolerance settings are consistent and robust. Figure 1 illustrates the 3D Hyperparameter Space versus the R^2 performance. The visualization shows that the optimal region in the hyperparameter space is not arbitrary but clustered, which indicates that a good performance appears when the three parameters are in balance. If one deviates too much from these ranges, there will be a sharp decline in the value of R^2, as shown by the darker points. EC was able to find these "sweet spots" where the accuracy peaked.

Table 3. Evolved Hyperparameter Statistics

	Mean	SD	Min	Max	Optimal range
C	54.237	47.89	0.234	234.567	20-80
γ	0.004	0.007	0.0001	0.034	0.001-0.01
ε	0.145	0.067	0.034	0.287	0.08-0.2

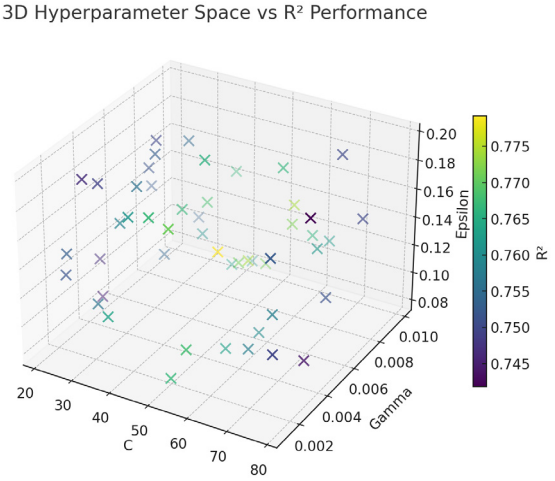

Figure 1. 3D hyperparameter space vs. R^2 performance

Regarding the kernel parameters, the experiments showed that the RBF kernel was by far the most effective, being selected in 67% of the computer runs. This shows that the employed data has a

Table 4. Efficiency and Generalization Analysis

	Description	Best value	Run ID	Parameters
Highest efficiency	Best R^2/SV ratio	0.012	66	C=47.234, $\gamma = 0.003$, $\varepsilon = 0.157$, kernel=RBF
Best generalization	Lowest SV count	68.2	66	C=47.234, $\gamma = 0.003$, $\varepsilon = 0.157$, kernel= RBF
Fastest convergence	Fewest fitness evaluations	1.890	23	C=35.678, $\gamma = 0.005$, $\varepsilon = 0.089$, kernel= RBF
Most stable	Lowest R^2 and SD	0.0234	45	C=41.789, $\gamma = 0.004$, $\varepsilon = 0.167$, kernel= RBF

strong non-linear nature. The linear kernel was selected in 18% of the cases, the sigmoid one accounted for 12% of the runs, and the polynomial kernel was selected in only 3% of the cases.

Next, the aim was to establish which of the SVR+EC candidate solutions featured the highest efficiency, the best generalization, the fastest convergence, or which was the most stable (Table 4). The highest efficiency was achieved in Run 66, with the best R^2-to-support vector ratio of 0.012, which indicates an excellent balance between the simplicity and the predictive accuracy of the model. Interestingly, the same model achieved the best generalization, using only 68.2 support vectors, which indicates that its configuration is the optimal trade-off between complexity and generalization.

The variation observed between different runs (e.g. Run ID 66) is explained by the stochastic nature of the evolutionary algorithm: since each of the 100 independent runs starts from a different random initialization, certain runs converged more quickly to favourable regions of the search space, resulting in a slightly stronger performance.

For the benchmarking analysis the SVR+EC model that had the best R^2 was selected. Its performance was compared with that of the traditional SVR+Grid Search optimization. The results revealed a *t*-statistic of 8.45 with a *p*-level < 0.001 and a large effect size (Cohen's *d*) of 1.67. This result shows that the EC provides meaningful insights into the systematic search strategies. Furthermore, Table 5 provides a baseline comparison of the default SVR, Grid Search SVR, and SVR+EC models in terms of R^2, RMSE, MAE, SVs, training time and improvement in terms of R^2. The comparative analysis demonstrates that indeed using EC for optimizing the SVR parameters outputs a superior model for pollen count forecasting. It can be seen that the Grid Search SVR is better than the baseline configuration, achieving measurable gains in predictive accuracy, increasing the value of R^2 by more than 14%, while simultaneously reducing the error rate and lowering the number of SVs needed. However, the improvements came with a higher computational cost in terms of training time. The SVR+EC model showed the best performance across all metrics. If one compares it to the default SVR, the RMSE decreased by 25.6%, the MAE by 24.8%, while the value of R^2 increased by 23.3%. This outcome reflects the model's precision and explanatory power. Moreover, the number of SVs needed declined by 14.7%, which indicates that the SVR+EC approach is not only more accurate but it is also more compact and can generalize well. The training time also increased.

Table 5. Baseline Comparison of the employed models

	Default SVR	Grid Search SVR	SVR+EC
R^2	0.6524	0.7456	0.805
RMSE	45.234	38.789	33.456
MAE	33.567	28.234	25.234
SVs	95.3	89.7	81.3
Training time	3.2 sec	45.6 sec	67.3 sec
Improvement of R^2	baseline	+14.3% R^2	+23.3% R^2

It is important to note that SVR+EC provided a higher accuracy than Grid Search SVR, but it required more computational resources, since multiple independent runs and generations were necessary to achieve convergence. In the context of pollen forecasting, this trade-off is acceptable if the training is conducted offline and the models are updated periodically, as the final prediction step is computationally inexpensive. However, for strict real-time forecasting where computational resources are limited, hybrid approaches may be more practical.

For a better visualization, this information was plotted in Figure 2.

The kernel coefficient γ is correlated with R^2 (0.267), which indicates that a careful tuning of its value will enhance the prediction accuracy. This fact is reinforced by the negative correlation with

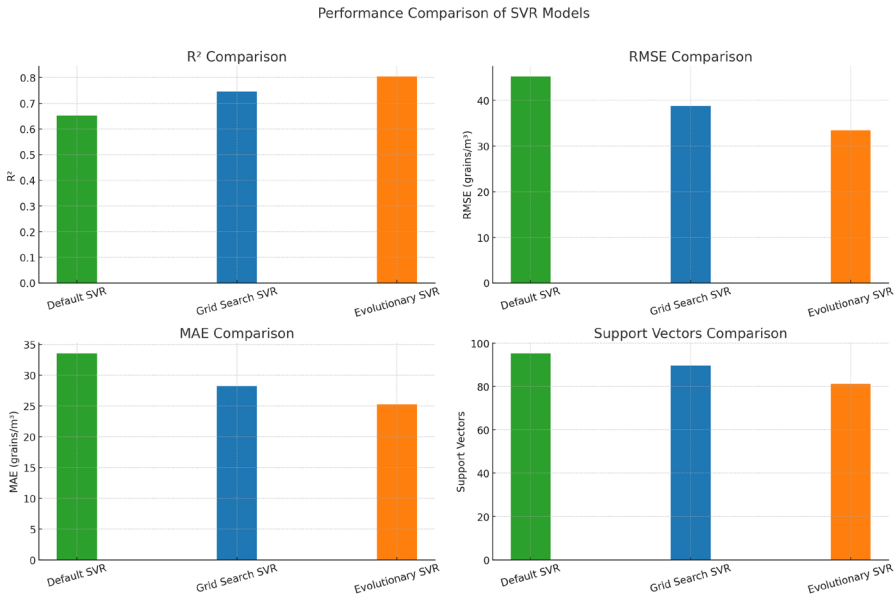

Figure 2. Performance comparison for the SVR models

RMSE (-0.245) and MAE (-0.223). An interesting finding is the fact that γ has a slightly negative correlation with the number of SVs (-0.156) which suggests that higher values of γ might lead to more compact models. In contrast with the other two parameters, ε features weaker and less consistent correlations. It has a negative relationship with R^2 (-0.189), and positive correlations with RMSE (0.167) and MAE (0.145), which indicates that a large value might lead to a decrease in accuracy. In terms of performance metrics, the findings confirm the expected inverse relationship between R^2 and RMSE and MAE. Finally, the number of SVs is negatively correlated with R^2 (-0.456) which indicates that a higher accuracy is achieved with more compact models. This finding aligns with the efficiency gains which were observed under evolutionary optimization (Figure 3).

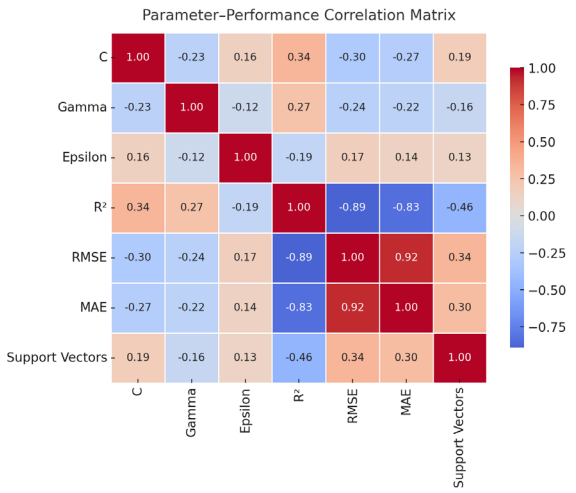

Figure 3. Correlation Heatmap

5. Conclusions

In this study, EC was applied for optimizing the SVR hyperparameters and its performance was compared with that of traditional Grid Search Optimization. The findings show that EC is a viable solution for improving performance and generalization. Nevertheless, some of its limitations should be acknowledged. First, the dataset used in this study is relatively small and restricted to specific monitoring sites, which ultimately limits the generalizability of the results to other regions that feature different pollen distributions or environmental conditions. Another drawback lies in the fact that for achieving a strong performance, the EC's computational power is more demanding than in the case of classical methods, due to the repeated runs required to account for randomness. Future work should investigate efficiency improvements such as parallelization and hybrid search strategies. When EC is applied on larger datasets, the expectation is that it remains effective. Nonetheless, scalability depends on computational effort and algorithmic efficiency. Since a smaller dataset was used in this study, the risk of overfitting which appears in cases when models are tuned too aggressively to local patterns should be considered. This risk has been mitigated using cross-validation.

Finally, the proposed EC-based model incorporated constraints that were able to avoid unrealistic and incompatible parameter combinations. Each hyperparameter was bounded

by predefined ranges and restricted to valid kernel types, which ensured that the evolutionary search was confined to meaningful solution spaces. In this manner, convergence stability was maintained, while still allowing the exploration of the parameter landscape.

Overall, the obtained results suggest that EC has the potential of optimizing SVR in prediction tasks. At the same time, further research is needed to assess its scalability, to minimize computational costs, and to confirm its robustness across different datasets and application domains.

REFERENCES

Astray, G., Amigo Fernández, R., Fernández-González, M. et al. (2025) Machine Learning to Forecast Airborne *Parietaria* Pollen in Northwest Spain Using Random Forest, SVM, and ANN Models. *Sustainability*. 17(4), art. ID 1528. https://doi.org/10.3390/su17041528.

Corić, R., Matijević, D. & Marković, D. (2023) PollenNet – a deep learning approach to predicting airborne pollen concentrations. *Croatian Operational Research Review*. 14(1), 1–13. https://doi.org/10.17535/crorr.2023.0001.

Cortes, C & Vapnik, V. (1995) Support-vector networks. *Machine Learning*. 20, 273-297. https://doi.org/10.1007/BF00994018.

Dwarakanath, D., Milic, A., Beggs, P. J. et al. (2024) A global survey addressing sustainability of pollen monitoring. *World Allergy Organization Journal*. 17(12), art. ID 100997. https://doi.org/10.1016/j.waojou.2024.100997.

European Climate and Health Observatory. (2025) *Pollen.* https://climate-adapt.eea.europa.eu/en/observatory/evidence/health-effects/aeroallergens [Accessed 24th August 2025]

Kitinoja, M.A., Hugg, T. T. & Siddika, N. et al. (2020) Short-term exposure to pollen and the risk of allergic and asthmatic manifestations: a systematic review and meta-analysis. *BMJ Open.* 10(1), art. ID e029069. https://doi.org/10.1136/bmjopen-2019-029069.

Landesberger, V., Grenzebach, K., Schreiber, F. et al. (2023) Conception and pilot testing of a self-management health application for patients with pollen-related allergic rhinitis and allergic asthma - the APOLLO app. *Scientific Reports*. 13(1), art. ID 21568. https://doi.org/10.1038/s41598-023-48540-4.

Lops, Y., Choi, Y., Eslami, E. et al. (2020) Real-time 7-day forecast of pollen counts using a deep convolutional neural network. *Neural Computing and Applications*. 32(15), 11827–11836. https://doi.org/10.1007/s00521-019-04665-0.

Luengo-Fernandez, R., Walli-Attaei, M., Gray, A. et al. (2023) Economic burden of cardiovascular diseases in the European Union: a population-based cost study. *European Heart Journal.* 44(45), 4752-4767. https://doi.org/10.1093/eurheartj/ehad583.

Makra, L., Coviello, L., Gobbi, A. et al. (2024) Forecasting daily total pollen concentrations on a global scale. *Allergy*. 79(8), 2173-2185. https://doi.org/10.1111/all.16227.

Oh, J.-W. (2022) Pollen Allergy in a Changing Planetary Environment. *Allergy Asthma Immunology Research*. 14(2), 168–181. https://doi.org/10.4168/aair.2022.14.2.168.

Rad, D., Kiss, C., Paraschiv, N. et al. (2024) Automated Comparative Predictive Analysis of Deception Detection in Convicted Offenders Using Polygraph with Random Forest, Support Vector Machine, and Artificial Neural Network Models. *Studies in Informatics and Control*. 33(3), 39–48. https://doi.org/10.24846/v33i3y202404.

Ribeiro, J., Santos, R., Analide, C. et al. (2024) Implementing Federated Learning and Explainability Techniques in Regression Models to Increase Transparency and Reliability. *Studies in Informatics and Control*. 33(4), 15–24. https://doi.org/10.24846/v33i4y202402.

U.S. Centers for Disease Control and Prevention (CDC). (2023) *Allergies*. https://www.cdc.gov/nchs/fastats/allergies.htm [Accessed 24th August 2025].

Wetherbee, G.A., (2022) *National Atmospheric Deposition Program Pollen Study Data for 2021 Pollen Season*. https://www.sciencebase.gov/catalog/item/62aa03b2d34ec53d277114de [Accessed 4th August 2025]

Yu, X., Zhao, J., Xu, Z. et al. (2024) AIpollen: An Analytic Website for Pollen Identification Through Convolutional Neural Networks. *Plants*. 13(22), art. ID 3118. https://doi.org/ 10.3390/plants13223118.