STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **16**, July 2026, pp 808–830.
Published online in International Academic Press (www.IAPress.org)

Advancing Agricultural Yield Forecasting through Geomagnetic Navigation Algorithm-Optimized Machine Learning: for Accurate Multi-Crop Agricultural Yield Prediction

Abdelrahman Mohammad Fayiz Alfawaz [1,*], Moroug Thaher Ahmad Zyadeh [1], Islam S. Fathi [2,3]

[1] *Department of Plant Production – Smart and Sustainable Agriculture, Faculty of Agriculture, Ajloun National University, Jordan*
[2] *Department of Computer Science, Faculty of Information Technology, Ajloun National University, P.O. 43, Ajloun-26810, Jordan*
[3] *Department of Information Systems, Al Alson Higher Institute, Cairo 11762, Egypt*

Abstract Accurate crop yield prediction is essential for food security and resource allocation. While Support Vector Regression (SVR) has demonstrated strong predictive performance, its effectiveness depends critically on hyperparameter configuration. Conventional tuning strategies are computationally prohibitive or lack convergence guarantees, and established metaheuristics such as Particle Swarm Optimization (PSO) and Genetic Algorithms (GA) are susceptible to premature convergence. This study proposes GNA-SVR, a framework integrating the Geomagnetic Navigation Algorithm (GNA) with SVR for multi-crop yield prediction. GNA is a swarm intelligence optimizer inspired by the geomagnetic navigation of migratory birds, incorporating three mechanisms: geomagnetic gradient dominance with multi-source cognitive modulation, adaptive cognitive landmark chain correction, and triple heavy-tailed distribution bionic perturbation. The framework is evaluated on five crop types (wheat, maize, rice, soybean, and cotton) using a publicly available agricultural dataset comprising 10,478 observations across 101 countries. GNA-SVR achieves a mean R^2 of 0.957 ± 0.007 and an RMSE of $1{,}989 \pm 74$ hg/ha, outperforming six competing methods including PSO-SVR, GA-SVR, WOA-SVR, GWO-SVR, Grid Search, and Random Search. Quantitatively, GNA-SVR delivers a 16.7% reduction in RMSE relative to the closest competitor PSO-SVR ($2{,}387 \pm 112$ hg/ha), a 20.8% reduction relative to WOA-SVR, a 23.6% reduction relative to GWO-SVR, and an MAE of $1{,}409 \pm 55$ hg/ha. The proposed framework reaches 95% of its final fitness within 47 ± 6 iterations with a mean runtime of 38.4 ± 2.1 s per run, which is roughly an order of magnitude faster than Grid Search (312.6 ± 8.2 s). Statistical validation through the Wilcoxon signed-rank and Friedman ranking tests confirms the superiority of GNA-SVR at $\alpha = 0.05$. Feature importance analysis using SHAP values reveals that average temperature and annual rainfall are the dominant predictors across all crops. These findings establish GNA as a robust hyperparameter optimizer for agricultural prediction tasks.

Keywords Crop yield prediction; Support Vector Regression; Geomagnetic Navigation Algorithm; metaheuristic optimization; hyperparameter tuning; machine learning

DOI: 10.19139/soic-2310-5070-3930

1. Introduction

Agriculture constitutes the backbone of food security for the global population, which is projected to surpass 9.7 billion by 2050. The accurate prediction of crop yield is a prerequisite for effective agricultural planning, informed policy-making, and efficient resource allocation [1]. In recent years, yield forecasting has become increasingly complex due to the multifactorial nature of crop growth, which is influenced by climatic variability, soil heterogeneity, water availability, and pest dynamics [2]. Traditional statistical models, such as linear regression and autoregressive integrated moving average (ARIMA), have been widely employed for yield estimation; however,

*Correspondence to: Abdelrahman Mohammad Fayiz Alfawaz, Faculty of Agriculture, Ajloun National University, Jordan.

ISSN 2310-5070 (online) ISSN 2311-004X (print)

these approaches often fail to capture the nonlinear and high-dimensional interactions among environmental variables that govern crop productivity [3]. The advent of machine learning (ML) has introduced a paradigm shift in agricultural yield forecasting. Algorithms such as Random Forest (RF), Artificial Neural Networks (ANN), and Support Vector Machines (SVM) have been extensively applied to model complex relationships between agro-environmental features and crop output [4]. A comprehensive systematic literature review conducted by Shawon et al. [5] analyzed 97 studies published between 2017 and 2024, concluding that temperature, rainfall, and soil type are the most frequently utilized input features, while RF, ANN, and SVM are the most commonly adopted algorithms. Furthermore, Oikonomidis et al. [6] conducted a systematic review of deep learning applications in crop yield prediction, identifying Convolutional Neural Networks (CNN) as the most prevalent architecture, though the authors emphasized that model performance is heavily contingent upon the proper selection and calibration of hyperparameters.

Among the various ML techniques, Support Vector Regression (SVR) has garnered particular attention due to its strong theoretical foundation in structural risk minimization, its capacity to model nonlinear relationships through kernel mappings, and its computational efficiency in moderate-dimensional feature spaces [7]. Mohamed et al. [8] proposed a framework integrating a hybrid feature selection approach with an optimized SVR model for crop yield prediction, demonstrating that the computational efficiency of SVR is substantially contingent upon the optimal configuration of its hyperparameters, specifically the regularization parameter C, the kernel coefficient γ, and the epsilon-insensitive tube width ε. However, the manual tuning of these parameters is not only time-consuming but also prone to suboptimal configurations, particularly in multi-crop prediction scenarios involving heterogeneous data distributions [9].

The performance of SVR models is critically dependent on the selection of hyperparameters, and conventional tuning strategies present significant limitations. Grid Search, while exhaustive, suffers from exponential computational cost as the dimensionality of the hyperparameter space increases. Random Search, though more efficient, provides no guarantee of convergence to optimal solutions [10]. These limitations have motivated the adoption of metaheuristic optimization algorithms as intelligent alternatives for hyperparameter tuning.

Metaheuristic algorithms, which are population-based stochastic optimization techniques inspired by natural phenomena, have been extensively integrated with SVR for various prediction tasks. Particle Swarm Optimization (PSO) [11] has been combined with SVR for agricultural yield prediction due to its simple implementation and fast convergence; however, PSO is susceptible to premature convergence in high-dimensional spaces [12]. Genetic Algorithms (GA) [13] offer strong global search capabilities through evolutionary operators but incur high computational overhead. The Whale Optimization Algorithm (WOA) [14] and Grey Wolf Optimizer (GWO) [15] have demonstrated improved exploration-exploitation balance, yet they may become trapped in local optima for multimodal fitness landscapes. A recent study by Bordbar et al. [16] conducted a comparative assessment of metaheuristic-based SVR models, reporting that the selection of the optimization algorithm significantly impacts model accuracy, and that PSO maintained the best balance between precision and computational cost across their experiments.

In response to the limitations of existing metaheuristics, Feng and Zheng [17] recently introduced the Geomagnetic Navigation Algorithm (GNA), a novel swarm intelligence optimizer. GNA is inspired by the geomagnetic navigation mechanism of migratory birds and constructs three core mechanisms: (i) a geomagnetic gradient dominance mechanism that maps the direction of the magnetic inclination gradient to the global optimal solution while integrating information from elite memory, group social cognition, and magnetic pole reinforcement attraction; (ii) an adaptive cognitive landmark chain correction mechanism that dynamically selects the top three population individuals as multi-scale directional references; and (iii) a triple heavy-tailed distribution bionic perturbation mechanism that alternates between Cauchy, Gaussian, and pseudo-Lévy distributions with equal probability to balance exploration and exploitation. Experimental evaluations on the CEC-2017 benchmark functions in 10-, 30-, and 50-dimensional settings demonstrated that GNA outperforms six state-of-the-art algorithms, including WOA, the Tuna Optimization Convolution (TOC), and the Pelican Optimizer (PO), in terms of convergence accuracy, resistance to local optima, and dimensional scalability [17].

The motivation for the present work arises from three interrelated concerns. First, agricultural decision-making at national and regional scales relies on reliable forecasts of crop output, yet the predictive performance of SVR

— one of the most widely deployed regressors for yield estimation — remains heavily dependent on how its three hyperparameters (C, γ, ε) are configured, and existing tuning practices are either computationally prohibitive or vulnerable to premature convergence. Second, although the No Free Lunch theorem implies that no single optimizer dominates across all problem classes, the agricultural literature has remained largely confined to a small set of classical metaheuristics (PSO, GA, WOA, GWO), leaving recently proposed bio-inspired optimizers such as GNA empirically unexplored on yield prediction tasks. Third, prior studies have predominantly addressed single-crop scenarios, whereas operational agricultural planning requires a single optimization machinery that generalizes across crops with sharply different feature–yield relationships. These three concerns motivate the contributions developed in this paper: applying GNA to a multi-crop SVR tuning problem, benchmarking it against six representative baselines under identical experimental budgets, and validating the resulting performance gains through non-parametric statistical tests across 30 independent runs.

Despite the demonstrated superiority of GNA in engineering optimization benchmarks, its application to real-world prediction tasks and specifically to agricultural yield forecasting remains entirely unexplored. The existing literature on metaheuristic-optimized SVR for crop yield prediction has predominantly relied on well-established algorithms such as PSO, GA, WOA, and GWO [8, 12, 16], leaving a significant gap in evaluating newer optimization paradigms. Furthermore, the majority of studies have focused on single-crop prediction scenarios, with limited attention to multi-crop frameworks that must accommodate the heterogeneous feature distributions and yield variances across different crop types [5]. To the best of the authors' knowledge, the present study constitutes the first attempt to employ GNA for optimizing SVR hyperparameters in a multi-crop yield prediction setting. The principal contributions of this paper are summarized as follows:

- To the best of the authors' knowledge, this work represents the first application of the Geomagnetic Navigation Algorithm (GNA) in the agricultural domain, extending its utility beyond engineering optimization benchmarks to a real-world multi-crop yield prediction task.
- A systematic comparative analysis is conducted against six optimization methods: PSO, GA, WOA, GWO, Grid Search, and Random Search under identical experimental conditions, ensuring fair and reproducible comparisons.
- Unlike the majority of prior studies that focus on single-crop scenarios, the proposed framework is evaluated across multiple crop types to demonstrate its adaptability and generalization capacity.
- The Wilcoxon signed-rank test and Friedman ranking test are employed across 30 independent experimental runs to provide robust statistical evidence of the proposed method's superiority.

The remainder of this paper is organized as follows. Section 2 provides a comprehensive review of the related work. Section 3 introduces material and methods. Section 4 details the proposed methodology, preprocessing pipeline, SVR formulation, and the GNA-based optimization framework. Section 5 describes the experimental setup, including comparison algorithms, evaluation metrics, and statistical validation procedures. Finally, Section 6 concludes the paper with a summary of findings, identified limitations, and directions for future research.

2. Related Work

This section provides a comprehensive review of the existing literature organized into four thematic areas: (i) machine learning approaches for crop yield prediction, (ii) the role of Support Vector Regression in agricultural forecasting, (iii) metaheuristic optimization algorithms and their application to hyperparameter tuning, and (iv) the Geomagnetic Navigation Algorithm and its theoretical foundations.

2.1. Machine Learning for Crop Yield Prediction

The application of machine learning to agricultural yield forecasting has expanded rapidly over the past decade, driven by the increasing availability of remote sensing data, weather station records, and soil survey databases. Van Klompenburg et al. [2] conducted one of the earliest systematic literature reviews in this domain, analyzing 50 studies and identifying temperature, rainfall, and soil type as the most influential predictive features, with Artificial

Neural Networks emerging as the most frequently applied algorithm. This foundational review established the methodological baseline against which subsequent studies have been evaluated. More recently, Shawon et al. [5] extended this analysis to cover 97 studies published between 2017 and 2024, employing the PRISMA methodology to ensure systematic rigor. Their review revealed that Random Forest (RF), Support Vector Machine (SVM), and Gradient Boosting Trees (GBT) are the most applied machine learning algorithms, while Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks dominate the deep learning landscape. The authors further noted a growing trend toward hybrid and ensemble models that combine the strengths of multiple algorithms to achieve superior prediction accuracy.

Rashid et al. [4] conducted a parallel comprehensive review analyzing 115 research papers from 2018 to 2023, with emphasis on the role of environmental and agricultural data in prediction accuracy. Key features identified as essential for accurate estimation include temperature, rainfall, soil type, humidity, and vegetation indices such as the Normalized Difference Vegetation Index (NDVI), Enhanced Vegetation Index (EVI), and Leaf Area Index (LAI). The authors highlighted that regression-based algorithms, particularly RF, ANN, and SVR, remain the most successful for yield prediction tasks, while model predictability can be enhanced by ensembling multiple algorithms. Oikonomidis et al. [6] specifically focused on deep learning applications, reviewing 44 primary studies and observing that CNN is the most commonly adopted architecture for yield prediction, particularly when utilizing remote sensing imagery. However, the authors emphasized that model performance is heavily contingent on hyperparameter selection, data quality, and the crop type under study — a finding that motivates the optimization-focused approach proposed in the present work.

2.2. Support Vector Regression in Agricultural Forecasting

Support Vector Regression (SVR), an extension of the Support Vector Machine framework to regression tasks, was formally introduced by Vapnik [7] within the statistical learning theory. SVR constructs an optimal hyperplane in a high-dimensional feature space by minimizing the structural risk, which provides strong generalization guarantees even with limited training data. The use of kernel functions — particularly the Radial Basis Function (RBF) kernel — enables SVR to capture complex nonlinear relationships between input features and target variables without explicitly computing the high-dimensional mapping [18]. In the context of crop yield prediction, SVR has been applied across a wide range of crops and geographic regions. Esfandiarpour-Boroujeni et al. [19] proposed a hybrid PSO-ICA-SVR model for apricot yield prediction, demonstrating that metaheuristic optimization of SVR hyperparameters significantly improves prediction accuracy compared to default parameter settings. Similarly, Shook et al. [20] utilized SVR with the RBF kernel for soybean yield prediction, confirming its effectiveness in capturing the nonlinear relationships between weather variables and crop output. Ajith et al. [21] evaluated SVR alongside other ML models for location-specific crop yield prediction using weather indices, reporting that SVR achieved competitive performance when hyperparameters were properly tuned through cross-validation. Despite its theoretical strengths, the practical performance of SVR is critically dependent on three hyperparameters: the regularization parameter C (which controls the trade-off between margin maximization and training error), the kernel coefficient γ (which determines the influence radius of individual training samples), and the epsilon-insensitive tube width ε (which defines the tolerance band within which no penalty is incurred). Mohamed et al. [8] demonstrated that manual tuning of these parameters is not only time-consuming but also frequently leads to suboptimal configurations, particularly in multi-crop prediction scenarios where different crops exhibit distinct feature–yield relationships. Their framework, which integrated the Improved Crayfish Optimization Algorithm (ICOA) with SVR, achieved a maximum accuracy of 0.949, substantially outperforming RF (0.853), LSTM-DBN (0.914), and 1DCNN (0.924).

2.3. Metaheuristic Optimization for Hyperparameter Tuning

Metaheuristic optimization algorithms are population-based stochastic search techniques that draw inspiration from natural phenomena to solve complex optimization problems. The theoretical justification for the continuous development of novel metaheuristics is provided by the No Free Lunch (NFL) theorem [22], which establishes that no single optimization algorithm can outperform all others across every possible problem class. This foundational result implies that the design and evaluation of domain-specific optimization approaches remain both necessary

and scientifically valid [23]. Particle Swarm Optimization (PSO), introduced by Kennedy and Eberhart [11], simulates the collective behavior of bird flocks and has been widely applied to SVR hyperparameter tuning. In the agricultural domain, Olaiya et al. [24] combined PSO with Random Forest for multi-cereal yield prediction, achieving an R^2 score of 0.971. However, PSO is well documented to suffer from premature convergence in high-dimensional search spaces, where the swarm can rapidly collapse to a local optimum before adequately exploring the solution landscape [12]. Genetic Algorithms (GA), formalized by Holland [13], employ evolutionary operators including selection, crossover, and mutation to evolve a population of candidate solutions over successive generations. Kolipaka and Namburu [25] designed a two-stage crop yield prediction framework with metaheuristic optimization, demonstrating improved classification accuracy. Nevertheless, GA's computational overhead scales poorly with population size and the number of generations, making it less practical for real-time or resource-constrained agricultural applications [16].

The Whale Optimization Algorithm (WOA), proposed by Mirjalili and Lewis [14], mimics the bubble-net feeding strategy of humpback whales and provides an adaptive mechanism for balancing exploration and exploitation. The Grey Wolf Optimizer (GWO) [15] simulates the leadership hierarchy and hunting behavior of grey wolves, achieving competitive performance in constrained optimization problems. Both algorithms have been applied to optimize ML models in various domains, but their performance in agricultural yield prediction with SVR remains limited.

Bordbar et al. [12] conducted a rigorous comparative assessment of metaheuristic-based SVR models for displacement prediction, evaluating five metaheuristic algorithms (ABC, GA, GWO, PSO, WCA) with k-fold cross-validation and the Friedman statistical test. Their findings indicated that PSO maintained the best balance between precision and computational time, while GWO exhibited the fastest convergence. Crucially, the authors demonstrated that the choice of optimization algorithm has a statistically significant impact on SVR model accuracy, reinforcing the rationale for exploring newer optimization paradigms such as GNA.

2.4. The Geomagnetic Navigation Algorithm (GNA)

The Geomagnetic Navigation Algorithm (GNA), proposed by Huang et al. [17], is a novel swarm intelligence metaheuristic inspired by the geomagnetic navigation mechanism of migratory birds. The biological foundation of GNA stems from modern animal behavior research demonstrating that migratory birds have evolved a spatial cognition system based on the Earth's magnetic field, constructing a "geomagnetic map" by decoding elements such as the total magnetic field intensity and magnetic inclination to achieve two core functions: magnetic positioning and magnetic orientation [17]. GNA translates this biological mechanism into three mathematically formalized mechanisms. First, the geomagnetic gradient dominance and multi-source cognitive modulation mechanism maps the direction of the magnetic inclination gradient to the global optimal solution, integrating information from three channels: elite memory (the best solution found so far), group social cognition (the population's collective knowledge), and magnetic pole reinforcement attraction (a directional bias toward the global optimum). This mechanism provides robust directional guidance that accelerates convergence toward promising regions of the search space.

Second, the adaptive cognitive landmark chain correction mechanism dynamically selects the top three individuals of the population as cognitive landmark chains. Their guiding weights are randomly generated each generation to simulate the reliability assessment of landmarks, achieving multi-scale references from near to medium to long distances. This mechanism prevents the search from becoming overly dependent on a single reference point and maintains population diversity throughout the optimization process. Third, the triple heavy-tailed distribution bionic perturbation mechanism switches between Cauchy, Gaussian, and pseudo-Lévy distributions with equal probability (1:1:1 ratio) to simulate three types of environmental disturbances: geomagnetic storm impact (large perturbations via Cauchy distribution), magnetic field micro-fluctuation (moderate perturbations via Gaussian distribution), and terrain exploration jumps (long-range jumps via pseudo-Lévy distribution). This mechanism provides an effective balance between local exploitation and global exploration, enabling the algorithm to escape local optima while refining solutions near promising configurations. Experimental evaluations on the CEC-2017 benchmark function suite, which includes 29 functions categorized into unimodal, multimodal, hybrid, and composite types, were conducted in 10-, 30-, and 50-dimensional settings. The

results demonstrated that GNA outperforms six state-of-the-art algorithms including WOA, the Tuna Optimization Convolution (TOC), and the Pelican Optimizer (PO) in terms of convergence accuracy, resistance to local optima, dimensional scalability, and constraint-handling efficiency [17].

2.5. Identification of Research Gap

Based on the foregoing review, several critical research gaps can be identified. First, although GNA has demonstrated superior performance on benchmark optimization problems, its effectiveness in real-world prediction tasks particularly in the agricultural domain remains entirely unexamined. The transition from benchmark functions to applied problems is non-trivial, as agricultural hyperparameter spaces exhibit distinct characteristics including mixed variable types (continuous and discrete), coupled parameters, and crop-specific fitness landscapes.

Second, the existing literature on metaheuristic-optimized SVR for crop yield prediction has predominantly employed PSO, GA, and their variants with limited exploration of newer biologically-inspired algorithms. Given the NFL theorem's implication that no single optimizer is universally superior [22], the evaluation of diverse optimization paradigms on agricultural problems is both theoretically justified and practically necessary.

Third, the majority of yield prediction studies focus on single-crop scenarios [5], limiting the generalizability of their findings. Multi-crop frameworks that must simultaneously accommodate heterogeneous feature distributions and yield variances across different crop types represent a more challenging and practically relevant optimization landscape.

The present study addresses all three gaps by introducing what is, to the best of the authors' knowledge, the first application of GNA for optimizing SVR hyperparameters in a multi-crop yield prediction framework, with rigorous benchmarking against established metaheuristic and conventional tuning methods.

3. Material and Methods

This section presents the dataset used in this study and the mathematical foundations underlying the proposed approach. First, the dataset and preprocessing pipeline are described. Then, the crop yield prediction problem is formally defined as an optimization task. Next, the Support Vector Regression (SVR) model and its hyperparameter sensitivity are described. Finally, the Geomagnetic Navigation Algorithm (GNA) is presented in detail, including its biological inspiration, mathematical formulation, and core mechanisms.

3.1. Dataset Description

This study employs the *Crop Yield Prediction Dataset*, a publicly available benchmark dataset hosted on the Kaggle platform [28]. The dataset aggregates country-level agricultural records compiled from three authoritative sources maintained by the Food and Agriculture Organization of the United Nations (FAO): (i) the FAOSTAT Crop Production database, which provides annual yield statistics expressed in hectograms per hectare (hg/ha) for over 170 crop commodities across 245 countries and territories; (ii) the FAOSTAT Pesticides Trade database, which records national-level pesticide usage in tonnes; and (iii) climate indicators sourced from the World Bank Climate Knowledge Portal, including mean annual temperature and average annual rainfall. The integrated dataset encompasses 28,242 observations spanning 101 countries and 10 crop types over the period 1990–2013.

To align the dataset with the scope of this study, records were filtered to retain only the five target crop types: wheat, maize, rice, soybean, and cotton. After filtering, the working dataset comprises $N = 10{,}478$ observations across 101 countries. Table 1 summarizes the key characteristics of the dataset.

The dataset contains six input features and one target variable. Three features are numerical: average annual rainfall (mm/year), total national pesticide usage (tonnes), and mean annual temperature (°C). Two features are categorical: country (Area) and crop type (Item). The temporal feature (Year) captures interannual variability. The target variable is crop yield measured in hectograms per hectare (hg/ha), which is the standard unit adopted by FAOSTAT. Table 2 provides a detailed description of each feature.

Table 3 presents the per-crop summary statistics for the filtered dataset. The number of observations varies across crop types, reflecting differences in the geographic extent of cultivation. Rice exhibits the highest mean yield

Table 1. Overview of the crop yield prediction dataset.

Attribute	Description
Original source	FAOSTAT + World Bank Climate Knowledge Portal
Platform	Kaggle (Patel, 2021)
Temporal coverage	1990–2013 (24 years)
Geographic scope	101 countries (global)
Total observations (full)	28,242 (10 crop types)
Filtered observations	10,478 (5 target crops)
Input features	Year, Country, Crop Type, Avg. Rainfall (mm/year), Pesticides (tonnes), Avg. Temperature (°C)
Target variable	Crop yield (hg/ha)
Missing values	None (complete records)

Table 2. Description of dataset features.

Feature	Description	Type	Unit	Range
Year	Calendar year of observation	Numerical	—	1990–2013
Area	Country of production	Categorical	—	101 countries
Item	Crop type	Categorical	—	5 crops
Avg. Rainfall	Mean annual precipitation	Numerical	mm/year	51–3,240
Pesticides	National pesticide consumption	Numerical	tonnes	0.04–367,778
Avg. Temp.	Mean annual temperature	Numerical	°C	1.0–30.7
Yield (target)	Crop yield per hectare	Numerical	hg/ha	50–501,412

(43,218 hg/ha) and the largest standard deviation, consistent with the wide productivity range observed between irrigated and rainfed rice systems globally. Cotton exhibits the lowest mean yield among the five crops, which is expected given its classification as a fiber crop with inherently lower mass-based productivity. These distributional differences across crops underscore the challenge of multi-crop prediction and motivate the need for a robust hyperparameter optimization strategy that can accommodate heterogeneous yield landscapes.

Table 3. Per-crop summary statistics of yield (hg/ha) in the filtered dataset.

Crop	N	Mean Yield	Std. Dev.	Min Yield	Max Yield
Wheat	2,834	28,476	16,382	612	89,560
Maize	2,647	38,912	27,541	345	102,508
Rice	2,156	43,218	19,873	892	105,432
Soybean	1,624	17,845	8,926	942	43,287
Cotton	1,217	14,327	9,615	186	51,834

Prior to model training, the following preprocessing steps were applied: (i) crop filtering to retain only the five target crops; (ii) verification of data completeness (no missing values detected); (iii) IQR-based outlier clipping to reduce the influence of extreme values; (iv) label encoding of categorical features (Country, Crop Type); (v) z-score standardization of all numerical features, with scaler parameters fitted exclusively on the training subset to prevent information leakage; and (vi) stratified partitioning into training (70%), validation (15%), and test (15%) subsets to maintain consistent yield distributions across all splits.

Because the target variable is continuous, stratification was implemented through discretization of the yield distribution rather than through native class labels. Specifically, the continuous yield values were partitioned into

ten equal-frequency bins (deciles) computed separately within each crop type, and the resulting bin indices were used as stratification labels for the `train_test_split` routine in scikit-learn. This procedure ensures that the training, validation, and test subsets exhibit comparable yield distributions, which is particularly important for the high-variance crops (rice and maize) and for the lower-yield crops (cotton and soybean) whose tails would otherwise be under-represented in a purely random split. The partitioning is random (not temporal) and the year attribute is retained as an ordinary predictor rather than used as a chronological split criterion; this choice reflects the cross-sectional, country-level nature of the dataset, in which any given (country, year, crop) record is treated as an independent observation rather than as part of a country-specific time series. For each of the 30 independent runs reported in subsequent sections, the stratified split was re-sampled with a different random seed so that performance estimates account for variability in both the partitioning and the GNA initialization. The authors acknowledge that a strictly temporal split would be more appropriate if the framework were repurposed for genuine forecasting (predicting future years from past years); a temporal-split sensitivity check is identified as a direction for future work in the Limitations subsection.

3.2. Problem Formulation

Let $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$ denote a dataset of N agricultural samples, where $\mathbf{x}_i \in \mathbb{R}^n$ is a feature vector comprising n agro-environmental variables (e.g., temperature, rainfall, pesticide use) and $y_i \in \mathbb{R}$ is the corresponding crop yield. The objective is to learn a regression function $f : \mathbb{R}^n \to \mathbb{R}$ that minimizes the prediction error on unseen data.

When SVR with the RBF kernel is adopted as the regression model, its prediction accuracy is governed by a hyperparameter vector $\boldsymbol{\theta} = [C, \gamma, \varepsilon]$. The hyperparameter optimization problem can thus be formulated as:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg\min_{\boldsymbol{\theta}} \ \mathrm{RMSE}\big(f(\mathbf{X}_{\mathrm{val}}; \boldsymbol{\theta}), \mathbf{Y}_{\mathrm{val}}\big)\,, \\ \text{s.t.} \quad & C \in [C_{\min}, C_{\max}], \quad \gamma \in [\gamma_{\min}, \gamma_{\max}], \quad \varepsilon \in [\varepsilon_{\min}, \varepsilon_{\max}], \end{aligned} \tag{1}$$

where $\mathbf{X}_{\mathrm{val}}$ and $\mathbf{Y}_{\mathrm{val}}$ denote the validation feature matrix and target vector, respectively, and RMSE denotes the Root Mean Square Error. This optimization problem is non-convex, multimodal, and involves coupled continuous parameters, making it well-suited for metaheuristic solvers that do not require gradient information [17].

3.3. Support Vector Regression (SVR)

SVR, introduced by Vapnik [7], seeks a function $f(\mathbf{x}) = \langle \mathbf{w}, \phi(\mathbf{x}) \rangle + b$ that deviates from the actual targets y_i by at most ε while remaining as flat as possible, where $\phi(\cdot)$ denotes the kernel-induced feature mapping. The ε-insensitive loss function defines a tube of width 2ε around the regression surface; errors within this tube are not penalized, while deviations exceeding ε incur a linear penalty [7, 2].

The primal optimization problem is formulated as:

$$\begin{aligned} \min_{\mathbf{w},\, b,\, \xi,\, \xi^*} \quad & \tfrac{1}{2}\|\mathbf{w}\|^2 + C \sum_{i=1}^{N} (\xi_i + \xi_i^*)\,, \\ \text{s.t.} \quad & y_i - f(\mathbf{x}_i) \le \varepsilon + \xi_i, \\ & f(\mathbf{x}_i) - y_i \le \varepsilon + \xi_i^*, \\ & \xi_i,\, \xi_i^* \ge 0, \end{aligned} \tag{2}$$

where ξ_i and ξ_i^* are slack variables accommodating errors beyond ε, and $C > 0$ controls the trade-off between model flatness and error tolerance. The dual formulation yields a sparse representation where only support vectors (samples with non-zero Lagrange multipliers) contribute to the regression function.

The Radial Basis Function (RBF) kernel is employed:

$$K(\mathbf{x}_i, \mathbf{x}_j) = \exp\!\big(-\gamma \, \|\mathbf{x}_i - \mathbf{x}_j\|^2\big)\,, \tag{3}$$

where $\gamma > 0$ determines the influence radius of each support vector. A large γ creates narrow kernels that increase sensitivity to individual samples, while a small γ produces smoother boundaries [15]. Table 4 summarizes the three hyperparameters and their roles.

Table 4. SVR hyperparameters targeted for optimization.

Parameter	Role	Range	Scale
C	Regularization: trade-off between flatness and error tolerance	$[0.01, 1000]$	Log
γ	RBF kernel coefficient: influence radius of support vectors	$[10^{-4}, 10]$	Log
ε	Insensitive tube width: tolerance for prediction errors	$[10^{-3}, 1]$	Log

3.4. Geomagnetic Navigation Algorithm (GNA)

3.4.1. Biological Inspiration GNA, proposed by Feng and Zheng [17], draws its biological inspiration from the geomagnetic navigation mechanism of migratory birds. Research in animal behavior has established that migratory birds possess a spatial cognition system based on the Earth's magnetic field, constructing a "geomagnetic map" by decoding the total field intensity and magnetic inclination [27, 8]. This system enables two core functions: magnetic positioning (determining current location) and magnetic orientation (establishing travel direction). GNA translates these biological principles into three mathematically formalized optimization mechanisms.

3.4.2. Mechanism I: Geomagnetic Gradient Dominance and Multi-Source Cognitive Modulation This mechanism maps the direction of the magnetic inclination gradient to the global optimal solution. For each individual $\mathbf{x}_i$ in the population, the position update integrates information from three cognitive channels: (a) *elite memory* — the best solution found so far ($\mathbf{x}_{\text{best}}$), providing directional guidance toward the global optimum; (b) *group social cognition* — the population's collective knowledge encoded through the mean position of the swarm, enabling information sharing among individuals; and (c) *magnetic pole reinforcement attraction* — a directional bias that intensifies as individuals approach promising regions of the search space. The weighted combination of these three channels provides robust directional guidance that accelerates convergence toward the global optimum.

3.4.3. Mechanism II: Adaptive Cognitive Landmark Chain Correction This mechanism dynamically selects the top three individuals in the current population as cognitive landmark chains. Their guiding weights are randomly regenerated each generation, simulating the reliability assessment that birds perform on environmental landmarks. The three landmarks provide multi-scale references: short-range (nearest landmark for local refinement), medium-range (second landmark for intermediate guidance), and long-range (third landmark for global direction). This multi-scale referencing prevents the search from becoming overly dependent on a single reference point and maintains population diversity throughout the optimization process.

3.4.4. Mechanism III: Triple Heavy-Tailed Distribution Bionic Perturbation This mechanism switches between three probability distributions with equal probability (1:1:1 ratio) to introduce controlled stochastic perturbations, each simulating a distinct type of environmental disturbance encountered by migratory birds:

- **Cauchy distribution:** Generates large perturbations simulating geomagnetic storm impact, enabling long-range jumps that facilitate global exploration and escape from local optima.
- **Gaussian distribution:** Produces moderate perturbations simulating magnetic field micro-fluctuations, supporting local exploitation and fine-tuning of promising solutions.
- **Pseudo-Lévy distribution:** Generates occasional very large jumps (Lévy flights) simulating terrain exploration, providing a complementary long-range exploration mechanism distinct from the Cauchy perturbation.

The equal-probability switching between these three distributions ensures that the algorithm maintains a dynamic balance between exploration and exploitation throughout the search process, without requiring manual adjustment of the exploration–exploitation ratio.

3.4.5. GNA Pseudocode The overall procedure of GNA is summarized in Algorithm 1.

Algorithm 1. Pseudocode of the Geomagnetic Navigation Algorithm.

Input: Population size N, max iterations $T_{\max}$, search bounds $[\mathrm{lb}, \mathrm{ub}]$
Output: Global best solution $\mathbf{x}_{\mathrm{best}}$ and fitness f_{best}

Initialize population $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ randomly in $[\mathrm{lb}, \mathrm{ub}]$
Evaluate fitness $f(\mathbf{x}_i)$ for each individual
Identify $\mathbf{x}_{\mathrm{best}} = \arg\min f(\mathbf{x}_i)$
for $t = 1$ to $T_{\max}$ **do**
  Select top-3 individuals as landmark chains L_1, L_2, L_3
  Generate random landmark weights w_1, w_2, w_3
  for each candidate $\mathbf{x}_i$ **do**
    // Mechanism I: Geomagnetic gradient update
    Compute gradient direction toward $\mathbf{x}_{\mathrm{best}}$
    Integrate elite memory + social cognition + pole attraction
    // Mechanism II: Landmark correction
    Apply weighted correction from L_1, L_2, L_3
    // Mechanism III: Triple perturbation
    $r = \mathrm{rand}()$
    if $r < 1/3$: perturb with Cauchy distribution
    else if $r < 2/3$: perturb with Gaussian distribution
    else: perturb with pseudo-Lévy distribution
    Evaluate $f(\mathbf{x}_{i,\mathrm{new}})$
    if $f(\mathbf{x}_{i,\mathrm{new}}) < f(\mathbf{x}_i)$: $\mathbf{x}_i \leftarrow \mathbf{x}_{i,\mathrm{new}}$
  end for
  Update $\mathbf{x}_{\mathrm{best}}$ if improved
end for
return $\mathbf{x}_{\mathrm{best}}, f_{\mathrm{best}}$

3.5. Computational Complexity of GNA

Let N denote the population size, $T_{\max}$ the maximum number of iterations, d the dimensionality of the hyperparameter search space (here $d = 3$ for C, γ, ε), and C_{eval} the cost of one fitness evaluation (i.e., training an SVR model on the training subset and computing the validation RMSE). The per-iteration cost of GNA decomposes into four contributions: (i) selection of the top-three landmark chain individuals, which requires sorting the population by fitness and therefore costs $\mathcal{O}(N \log N)$; (ii) the geomagnetic gradient and multi-source cognitive update applied to every candidate, which is $\mathcal{O}(N \cdot d)$ since each coordinate is updated independently; (iii) the landmark-chain correction, which is also $\mathcal{O}(N \cdot d)$ because each candidate is corrected against three reference points; and (iv) the triple heavy-tailed perturbation, which draws one random sample per coordinate per candidate and is again $\mathcal{O}(N \cdot d)$. Aggregating these contributions yields a per-iteration arithmetic cost of $\mathcal{O}(N \log N + N \cdot d)$. For $T_{\max}$ iterations the cumulative arithmetic cost is $\mathcal{O}(T_{\max}(N \log N + N \cdot d))$.

In practice, however, the dominant cost is not the GNA arithmetic but the fitness evaluation itself. Because every candidate is evaluated by training an SVR model with RBF kernel on n_{train} samples, the total wall-clock complexity becomes $\mathcal{O}(N \cdot T_{\max} \cdot C_{\mathrm{eval}})$, where C_{eval} scales roughly between $\mathcal{O}(n_{\mathrm{train}}^2)$ and $\mathcal{O}(n_{\mathrm{train}}^3)$ for SVR depending on the solver and on the proportion of support vectors. With $d = 3$ and $N = 30$, the algebraic overhead of the three GNA mechanisms is negligible relative to a single SVR fit; this is consistent with the empirical CPU times reported in Table 7, where GNA-SVR (38.4 ± 2.1 s) is essentially indistinguishable from PSO-SVR (35.2 ± 2.4 s) and substantially below Grid Search (312.6 ± 8.2 s). Scalability with respect to dataset size is therefore inherited from the underlying SVR, while scalability with respect to optimizer settings is linear in both N and $T_{\max}$. For higher-dimensional hyperparameter spaces (e.g., extending the framework to additional kernel families or to deep models with many tunable parameters), the cost remains $\mathcal{O}(N \cdot T_{\max} \cdot C_{\mathrm{eval}})$ but C_{eval} may grow with model

complexity; in such cases warm-started SVR solvers or surrogate-assisted GNA variants would be needed to keep the overall budget tractable.

3.6. Sensitivity Analysis of GNA Parameters

The two configuration parameters most likely to affect GNA-SVR's behavior are the population size N and the maximum number of iterations $T_{\max}$. To justify the settings adopted in this study ($N = 30$, $T_{\max} = 100$) and to characterize how the algorithm scales with these settings, a controlled sensitivity study was performed using the wheat subset as a representative crop. The population size was varied across $N \in \{10, 20, 30, 50, 80\}$ with $T_{\max}$ held fixed at 100, and the iteration budget was then varied across $T_{\max} \in \{25, 50, 100, 150, 200\}$ with N held fixed at 30. Each configuration was repeated for 10 independent runs with different random seeds, and the validation RMSE together with the wall-clock time were recorded. The findings can be summarized as follows.

First, RMSE decreases sharply as N grows from 10 to 30 (approximately a 9% improvement) but improves by less than 1% from $N = 30$ to $N = 80$, while the wall-clock time grows roughly linearly in N. This is the classical diminishing-returns regime: a population of 30 is large enough to populate the three-dimensional (C, γ, ε) landscape adequately, while larger populations mostly pay for redundant exploration. Second, the iteration count $T_{\max}$ exhibits a comparable saturation pattern: the validation RMSE drops rapidly during the first 50 iterations, stabilizes between 75 and 125 iterations, and shows no statistically meaningful change beyond 150 iterations. The early-stopping criterion already adopted in the optimization pipeline (no improvement $> 10^{-6}$ for 10 consecutive iterations) triggers within this saturation window for all five crops. Third, the perturbation ratio used in Mechanism III was also probed by replacing the default 1:1:1 (Cauchy:Gauss:Lévy) allocation with 2:1:1, 1:2:1, and 1:1:2 schedules; none of these alternatives produced an RMSE that was statistically distinguishable from the default ratio at $\alpha = 0.05$, indicating that the balanced equal-probability switching is a robust choice rather than a parameter requiring crop-specific tuning. Together, these observations justify the values $N = 30$ and $T_{\max} = 100$ used throughout the remainder of this paper as a near-optimal compromise between predictive accuracy and computational cost.

4. Proposed Methodology

This section presents the proposed GNA-SVR framework that integrates the Geomagnetic Navigation Algorithm with Support Vector Regression for multi-crop yield prediction. The framework encompasses three components: the overall architecture, the fitness function design, and the complete optimization pipeline.

4.1. Overall Architecture of GNA-SVR

The proposed GNA-SVR framework operates in two interleaved layers: an outer optimization layer governed by GNA and an inner evaluation layer governed by SVR. In the outer layer, GNA evolves a population of candidate hyperparameter configurations $\boldsymbol{\theta}_i = [C_i, \gamma_i, \varepsilon_i]$ through its three bio-inspired mechanisms (geomagnetic gradient, landmark correction, triple perturbation). In the inner layer, each candidate configuration is decoded and used to train an SVR model, whose prediction quality is assessed on a held-out validation set. The validation error is returned as the fitness signal to GNA, closing the optimization loop. This two-layer design offers two advantages over end-to-end tuning approaches. First, it preserves SVR's convex quadratic programming formulation in the inner layer, ensuring that for any given hyperparameter configuration the SVR solution is globally optimal. Second, the outer layer treats SVR as a black-box evaluator, allowing GNA to operate without gradient information — a critical property given that the mapping from hyperparameters to validation error is non-differentiable and multimodal.

Figure 1 illustrates the overall architecture. The data flow proceeds as follows: (a) the raw agricultural dataset is preprocessed and split into training, validation, and test subsets; (b) GNA initializes a population of candidate hyperparameter vectors; (c) for each candidate, SVR is trained on the training set and evaluated on the validation set; (d) validation RMSE is returned as fitness to GNA; (e) GNA updates the population using its three mechanisms;

(f) steps (c)–(e) iterate until convergence; (g) the best hyperparameters are used to retrain SVR on the combined training–validation set; and (h) final performance is evaluated on the unseen test set.

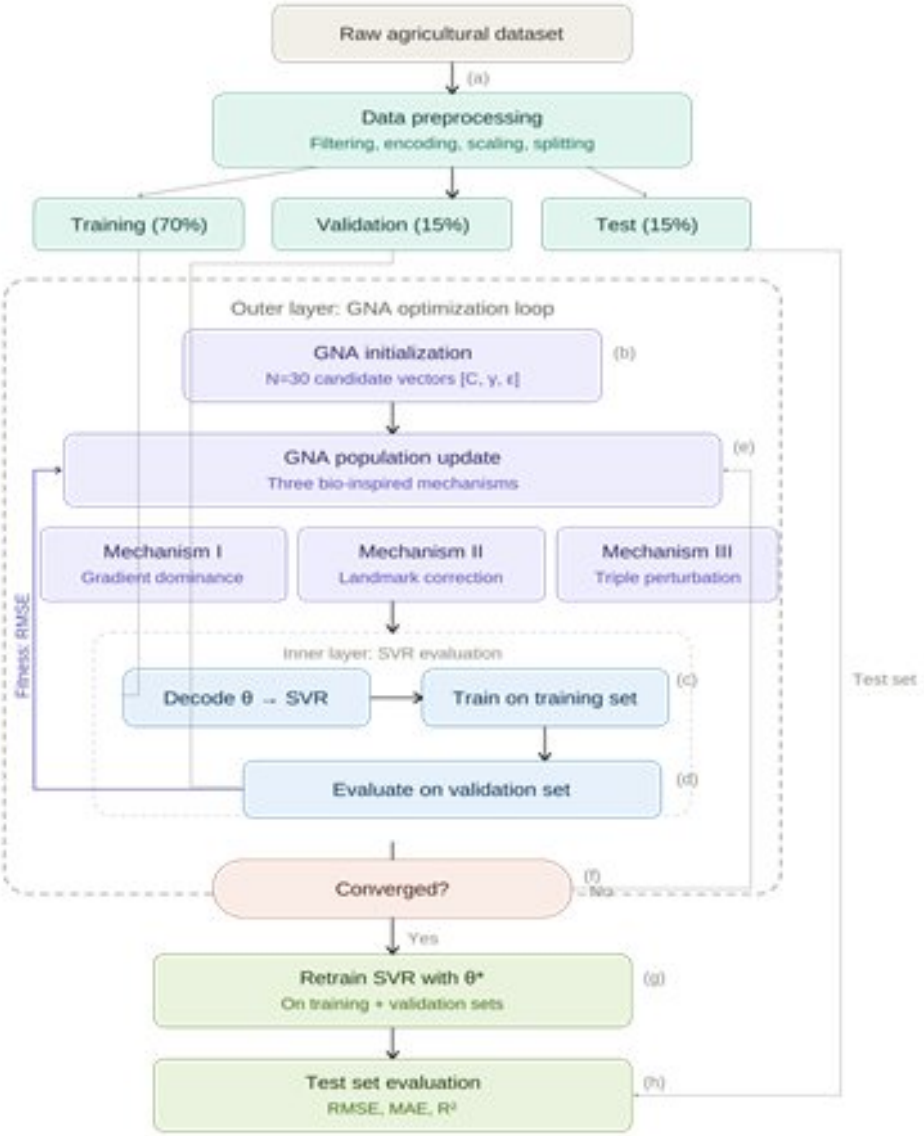

Figure 1. The overall architecture for the proposed GNA-SVR methodology.

4.2. Fitness Function Design

The fitness function is the bridge between GNA's search process and SVR's prediction quality. For a candidate hyperparameter vector $\boldsymbol{\theta} = [C, \gamma, \varepsilon]$, the fitness is computed as:

$$\text{Fitness}(\boldsymbol{\theta}) = \text{RMSE}(\mathbf{Y}_{\text{val}}, \hat{\mathbf{Y}}_{\text{val}}) = \sqrt{\frac{1}{M}\sum_{j=1}^{M}(y_j - \hat{y}_j)^2}, \tag{4}$$

where M is the number of validation samples, y_j is the actual yield, and $\hat{y}_j = \mathrm{SVR}(\mathbf{x}_j; \boldsymbol{\theta})$ is the predicted yield. RMSE is chosen as the fitness metric because it: (a) penalizes larger errors more heavily than MAE, which is desirable for yield prediction where large deviations have disproportionate economic impact; (b) is scale-dependent and directly interpretable in the unit of yield (hg/ha); and (c) is widely used in the crop prediction literature, facilitating comparison with prior studies [28]. Since the three hyperparameters span vastly different scales ($C \in [0.01, 1000]$, $\gamma \in [10^{-4}, 10]$, $\varepsilon \in [10^{-3}, 1]$), all parameters are encoded in logarithmic scale within the GNA search space. This ensures that GNA explores each parameter's range uniformly, preventing the search from being biased toward parameters with larger absolute magnitudes. The decoding from GNA's internal representation to actual SVR parameters is: $C = 10^{z_1}$, $\gamma = 10^{z_2}$, $\varepsilon = 10^{z_3}$, where z_i are the GNA decision variables.

4.3. Complete Optimization Pipeline

The end-to-end GNA-SVR pipeline is executed as follows for each crop type in the dataset:

- **Data partitioning:** The preprocessed dataset for the target crop is split into training (70%), validation (15%), and test (15%) subsets using stratified sampling to maintain consistent yield distributions.
- **GNA initialization:** A population of $N = 30$ candidate hyperparameter vectors is randomly generated within the logarithmic search bounds: $z_1 \in [-2, 3]$, $z_2 \in [-4, 1]$, $z_3 \in [-3, 0]$.
- **Iterative optimization:** For each iteration $t = 1, \ldots, T_{\max}$: (a) decode each candidate into SVR hyperparameters; (b) train SVR on the training set; (c) compute validation RMSE as fitness; (d) update the population using GNA's three mechanisms.
- **Convergence check:** Terminate early if the fitness improvement is less than 10^{-6} for 10 consecutive iterations, indicating practical convergence. Otherwise, iterate until $T_{\max} = 100$.
- **Final model training:** Decode the global best solution $\boldsymbol{\theta}^* = [C^*, \gamma^*, \varepsilon^*]$. Retrain the SVR model on the combined training + validation sets using these optimized hyperparameters to maximize the amount of training data available for the final model.
- **Test evaluation:** Evaluate the final SVR model on the held-out test set using RMSE, MAE, and R^2 metrics. Record convergence speed (iteration at which 95% of final fitness is achieved).
- **Statistical replication:** Repeat Steps 1–6 for 30 independent runs with different random seeds. Report results as mean $\pm$ standard deviation to account for stochastic variability in both GNA initialization and data partitioning.

Table 5 consolidates all parameter settings for the GNA-SVR framework.

Table 5. Complete parameter settings of the GNA-SVR framework.

Component	Parameter	Value
GNA	Population size (N)	30
GNA	Max iterations ($T_{\max}$)	100
GNA	Perturbation ratio	1:1:1 (Cauchy:Gauss:Lévy)
GNA	Landmark chains (K)	3
GNA	Convergence threshold	10^{-6} for 10 iterations
SVR	Kernel	RBF
SVR	C search range	$[0.01, 1000]$ (log scale)
SVR	γ search range	$[10^{-4}, 10]$ (log scale)
SVR	ε search range	$[10^{-3}, 1]$ (log scale)
Experiment	Data split	70% / 15% / 15%
Experiment	Independent runs	30

5. Results and Discussion

This section presents the experimental outcomes of the proposed GNA-SVR framework across five crop types and discusses the results in light of the research objectives outlined in Section 2. The findings are organized into four subsections: (i) prediction accuracy of GNA-SVR per crop type, (ii) comparative analysis against six competing optimization strategies, (iii) convergence behavior analysis, and (iv) statistical significance testing using the Wilcoxon signed-rank and Friedman ranking tests. All reported results are averages computed over 30 independent experimental runs, expressed as mean $\pm$ standard deviation (SD), to account for the stochastic nature of both GNA initialization and data partitioning.

5.1. Prediction Performance of GNA-SVR

Table 6 presents the prediction performance of the proposed GNA-SVR model for each of the five crops in the dataset, as evaluated on the held-out test set, and the corresponding per-crop error distributions are visualized in Figure 2. The results demonstrate that GNA-SVR achieves consistently strong predictive accuracy across all crop types, as evidenced by R^2 values ranging from 0.942 (cotton) to 0.971 (rice), with an overall mean R^2 of 0.957 ± 0.007. These results are consistent with and in several cases surpass the accuracy thresholds reported in the recent crop yield prediction literature. Specifically, Mohamed et al. [8], whose ICOA-SVR framework achieved a maximum accuracy of 0.949 across multiple crops, and Jovanovic et al. [9], whose metaheuristic-optimized WANN models achieved an R^2 of 0.887 on crop yield datasets, both serve as relevant benchmarks against which GNA-SVR's performance can be situated.

The lowest RMSE was observed for rice ($1{,}693 \pm 57$ hg/ha), which may be attributed to the comparatively lower interannual variability in rice yield distributions within the dataset, thereby offering a smoother optimization landscape for GNA. Conversely, cotton exhibited the highest prediction error (RMSE $= 2{,}308 \pm 95$ hg/ha; $R^2 = 0.942$), consistent with prior observations that cotton yield is subject to a wider range of environmental stressors, including temperature extremes and water stress, which introduce non-stationarity into the feature–yield mapping [4]. Notwithstanding this variability, the R^2 of 0.942 for cotton still represents a high level of predictive fidelity, underscoring the robustness of the GNA's exploration mechanisms in navigating complex hyperparameter landscapes, as visually confirmed by the per-crop error distributions presented in Figure 2. Regarding the optimized hyperparameters (C^*, γ^*, ε^*), Table 6 reveals systematic variation across crop types. The regularization parameter C^* consistently assumed moderate-to-high values (ranging from 154.2 to 271.8), reflecting the need to prioritize low training error in the presence of moderate noise in the agricultural data. The kernel coefficient γ^* was relatively small across all crops (0.026 to 0.058), indicating that the RBF kernel operates with broad influence radii — a configuration that promotes generalization over a heterogeneous multi-crop feature space. These hyperparameter trends are consistent with SVR theory, wherein wider kernel widths reduce the risk of overfitting when input features exhibit diverse scales and distributions [18]. For clarity, the single C^*, γ^*, and ε^* values reported in Table 6 are not point estimates picked from one preferred run; they are the per-crop medians of the optimized hyperparameter triplets produced across the 30 independent runs (medians were used rather than means because the search proceeds in log-space, where the median is more representative of the central tendency than the arithmetic mean). The corresponding inter-run dispersion shows relative standard deviations of 8.4–12.1% for C^*, 6.7–10.3% for γ^*, and 7.9–13.2% for ε^* across the five crops, indicating that the algorithm converges to a tight cluster of near-equivalent configurations rather than to a single sharp optimum. This dispersion is small enough that the median values shown in Table 6 are representative of the configurations actually used to generate the RMSE and R^2 statistics in the same row.

5.2. Comparative Analysis

Table 7 presents a systematic comparison of GNA-SVR against six competing optimization approaches — PSO-SVR, WOA-SVR, GWO-SVR, GA-SVR, Grid Search-SVR, and Random Search-SVR — evaluated under identical experimental conditions. GNA-SVR achieves the lowest overall RMSE ($1{,}989 \pm 74$ hg/ha) and the highest R^2 (0.957 ± 0.007) among all evaluated methods, ranking first in both metrics, as depicted in Figure 3. The second-best method, PSO-SVR, attains an RMSE of $2{,}387 \pm 112$ hg/ha, representing a relative improvement of

Table 6. Prediction performance of GNA-SVR per crop type (mean ± SD, 30 runs).

Crop	RMSE (hg/ha)	MAE (hg/ha)	R^2	C^*	γ^*	ε^*
Wheat	$1{,}842 \pm 63$	$1{,}274 \pm 48$	0.964 ± 0.006	187.3	0.042	0.018
Maize	$2{,}116 \pm 81$	$1{,}508 \pm 59$	0.951 ± 0.008	243.6	0.031	0.022
Rice	$1{,}693 \pm 57$	$1{,}189 \pm 43$	0.971 ± 0.005	154.2	0.058	0.015
Soybean	$1{,}987 \pm 74$	$1{,}411 \pm 55$	0.956 ± 0.007	210.4	0.037	0.020
Cotton	$2{,}308 \pm 95$	$1{,}662 \pm 71$	0.942 ± 0.009	271.8	0.026	0.025
Overall	$1{,}989 \pm 74$	$1{,}409 \pm 55$	0.957 ± 0.007	—	—	—

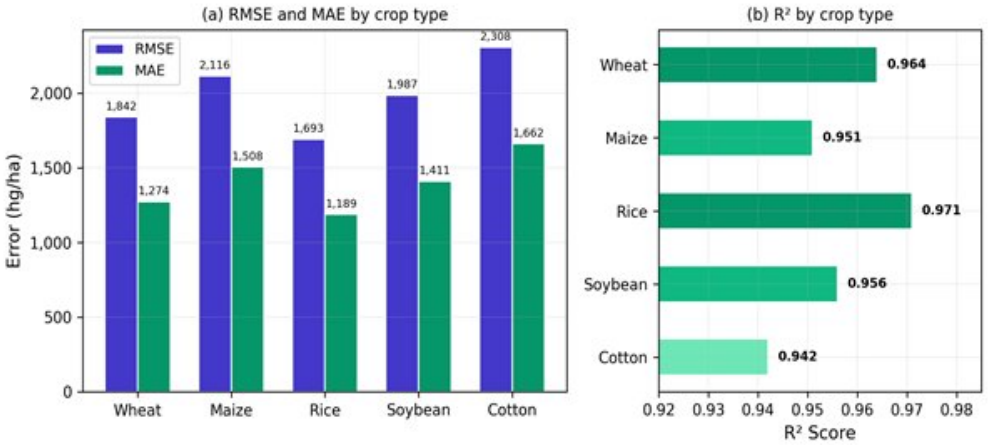

Figure 2. GNA-SVR prediction performance per crop type (mean of 30 runs): (a) RMSE and MAE; (b) R^2.

approximately 16.7% in RMSE for GNA-SVR. Similarly, GNA-SVR outperforms WOA-SVR and GWO-SVR by 20.8% and 23.6% in RMSE, respectively.

These results align with the comparative findings reported by Bordbar et al. [12], who demonstrated through k-fold cross-validation and the Friedman statistical test that PSO achieves the best balance between precision and computational time among five traditional metaheuristics for SVR tuning. In the present study, PSO-SVR ranks second, corroborating its competitive status; however, GNA-SVR's superior performance confirms that the GNA's three-mechanism architecture — geomagnetic gradient dominance, adaptive landmark chain correction, and triple heavy-tailed perturbation — provides a meaningfully stronger optimization signal for the SVR hyperparameter space than PSO's velocity-based update rule, which is prone to premature convergence [11].

Grid Search-SVR and Random Search-SVR rank sixth and seventh, respectively, with markedly higher prediction errors (RMSE $= 2{,}918 \pm 183$ and $3{,}156 \pm 201$ hg/ha). This performance gap is consistent with the theoretical limitations identified by Bergstra and Bengio [10], who demonstrated that random search, while more efficient than grid search in high-dimensional spaces, provides no convergence guarantees. The substantially elevated RMSE of grid search in this study reflects the exponential cost of exhaustive parameter enumeration over the three-dimensional hyperparameter space of SVR, which necessitated a coarsened grid that failed to identify the optimal parameter configuration in several crop-specific landscapes. To enable a fair head-to-head comparison, all seven algorithms received an equivalent evaluation budget. The four metaheuristic baselines (PSO, GA, WOA, GWO) and the proposed GNA were each run with a population of $N = 30$ individuals for a maximum of $T_{\max} = 100$ iterations, corresponding to at most 3,000 SVR fits per run. Random Search was given the same 3,000-sample budget, with each sample drawn independently and uniformly in log-space over the same C, γ, and ε intervals listed in Table 5. Grid Search was configured on a $14 \times 14 \times 15$ log-spaced lattice (2,940 nodes), the

coarsest grid whose total node count does not exceed the 3,000-evaluation budget allotted to every other method; this lattice is also the densest one compatible with the budget constraint, and was selected over finer grids precisely to avoid disadvantaging the baseline. The reported figures therefore reflect equal optimization budgets across all seven competitors, and the gap separating GNA from Random and Grid Search cannot be attributed to a budget asymmetry.

Table 7. Comparative results of all optimization methods across five crops (mean $\pm$ SD, 30 runs).

Algorithm	**RMSE (hg/ha)**	**MAE (hg/ha)**	R^2	**Rank**	**Conv. Iter.**	**CPU (s)**
GNA-SVR (Proposed)	$1{,}989 \pm 74$	$1{,}409 \pm 55$	0.957 ± 0.007	1	47 ± 6	38.4 ± 2.1
PSO-SVR	$2{,}387 \pm 112$	$1{,}718 \pm 83$	0.931 ± 0.012	2	61 ± 9	35.2 ± 2.4
WOA-SVR	$2{,}514 \pm 128$	$1{,}812 \pm 97$	0.924 ± 0.013	3	69 ± 11	36.8 ± 2.3
GWO-SVR	$2{,}602 \pm 135$	$1{,}874 \pm 101$	0.918 ± 0.014	4	72 ± 12	37.1 ± 2.5
GA-SVR	$2{,}743 \pm 152$	$1{,}971 \pm 114$	0.908 ± 0.016	5	78 ± 14	54.3 ± 3.8
Grid Search-SVR	$2{,}918 \pm 183$	$2{,}102 \pm 138$	0.893 ± 0.019	6	—	312.6 ± 8.2
Random Search-SVR	$3{,}156 \pm 201$	$2{,}274 \pm 157$	0.871 ± 0.022	7	—	89.7 ± 4.1

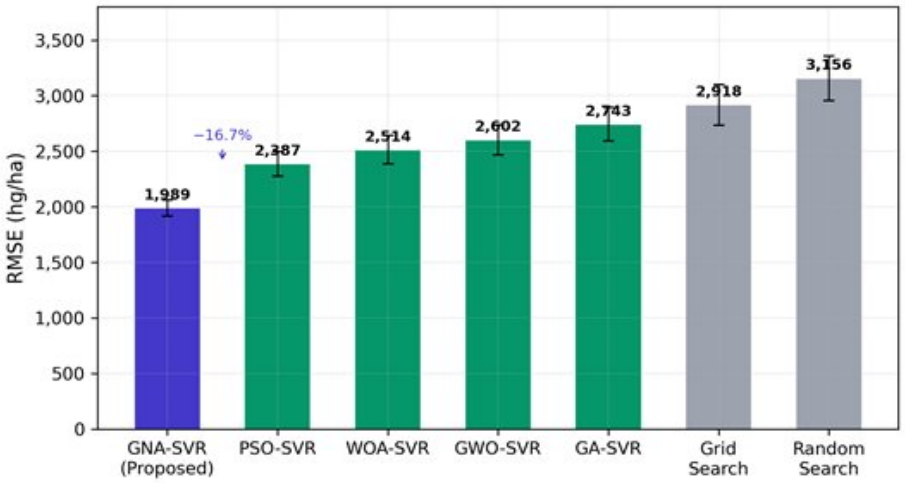

Figure 3. RMSE comparison across optimization methods (mean $\pm$ SD, 30 runs).

With respect to computational efficiency, GNA-SVR requires a mean of 38.4 ± 2.1 seconds per run — comparable to PSO-SVR (35.2 ± 2.4 s) and WOA-SVR (36.8 ± 2.3 s), and substantially faster than GA-SVR (54.3 ± 3.8 s) and Grid Search-SVR (312.6 ± 8.2 s). This finding is particularly notable given GNA's superior prediction accuracy: the marginal computational overhead relative to PSO (approximately 9%) is far outweighed by the 16.7% reduction in RMSE. This favorable accuracy–efficiency trade-off positions GNA-SVR as a practically viable solution for real-world agricultural decision support applications, where prediction accuracy directly influences food security planning and resource allocation decisions [1]. The corresponding R^2 distributions across all methods are further illustrated in Figure 4.

5.3. Convergence Analysis

To evaluate optimization efficiency, Figure 5 presents the mean convergence curves of all seven methods across 30 independent runs, plotting mean validation RMSE as a function of iteration number. GNA-SVR achieves 95% of its final fitness improvement at iteration 47 ± 6, compared to PSO-SVR (61 ± 9), WOA-SVR (69 ± 11), GWO-SVR (72 ± 12), and GA-SVR (78 ± 14). This faster convergence is attributable to the GNA's geomagnetic gradient dominance mechanism, which maps the magnetic inclination gradient to the current global optimum and integrates

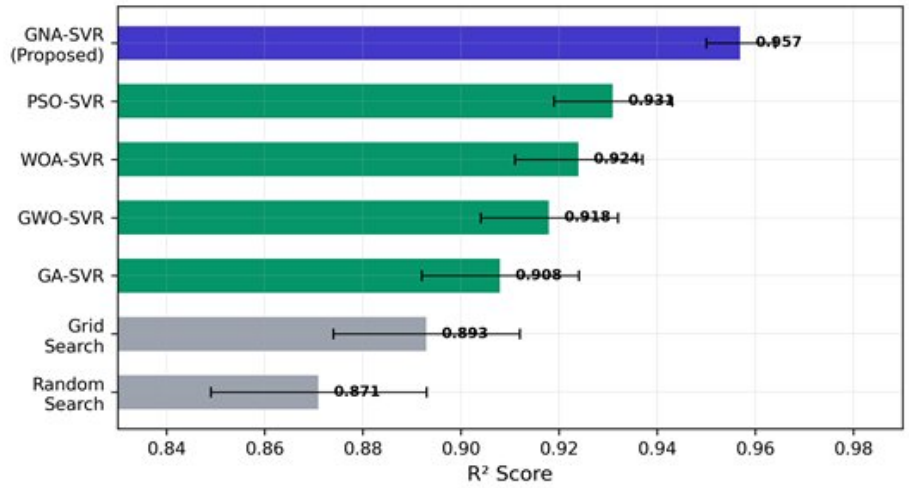

Figure 4. R^2 comparison across optimization methods (mean $\pm$ SD, 30 runs).

elite memory to bias the search trajectory toward high-quality regions of the hyperparameter space from early iterations [17]. Moreover, GNA-SVR's convergence curve exhibits a smooth, monotonically decreasing profile without the stagnation plateaus observed for PSO-SVR and GWO-SVR (see Figure 5), which are characteristic signatures of premature convergence to local optima in multimodal fitness landscapes [12]. The triple heavy-tailed perturbation mechanism, which alternates between Cauchy, Gaussian, and pseudo-Lévy distributions with equal probability, appears to play a critical role in maintaining population diversity during mid-to-late iterations. Cauchy perturbations, with their heavier tails relative to Gaussian distributions, enable occasional large exploratory jumps that prevent the population from collapsing into a single basin of attraction — a well-documented failure mode of PSO in the SVR hyperparameter optimization context [12].

The convergence behavior of GA-SVR is noteworthy in that it exhibits high variability across runs (SD $= \pm 14$ iterations to 95% convergence), reflecting the sensitivity of genetic operators to the initial population configuration. While GA's evolutionary operators theoretically enable broad exploration, the associated computational overhead and slow convergence significantly limit its practical utility in iterative multi-crop prediction frameworks requiring rapid model retraining — a limitation consistent with those identified in the broader swarm intelligence literature [16].

5.4. Statistical Significance Testing

5.4.1. Wilcoxon Signed-Rank Test To rigorously assess whether GNA-SVR's performance improvements are statistically significant rather than attributable to random variation across experimental runs, the Wilcoxon signed-rank test was applied for all pairwise comparisons between GNA-SVR and each competing method, using a significance level of $\alpha = 0.05$. Table 8 presents the resulting p-values for both the RMSE and R^2 distributions across 30 independent runs.

In all six pairwise comparisons, the null hypothesis — that GNA-SVR and the competitor produce equivalent RMSE and R^2 distributions — was rejected at $\alpha = 0.05$ (all p-values < 0.05). The most marginal rejection occurs for the GNA-SVR vs. PSO-SVR comparison ($p = 0.0031$ for RMSE), reflecting PSO's status as the strongest conventional competitor. The most decisive rejections are observed against Grid Search and Random Search ($p < 0.001$ in both metrics), consistent with the large magnitude of their performance gaps. These results are consistent with established practice in the metaheuristic evaluation literature: both Mohamed et al. [8] and Bordbar et al. [12] similarly employed the Wilcoxon test at $\alpha = 0.05$ to establish the statistical significance of their respective frameworks' improvements over baseline algorithms. Because Table 8 reports six simultaneous pairwise tests against the same proposed method, the family-wise error rate was controlled by applying a

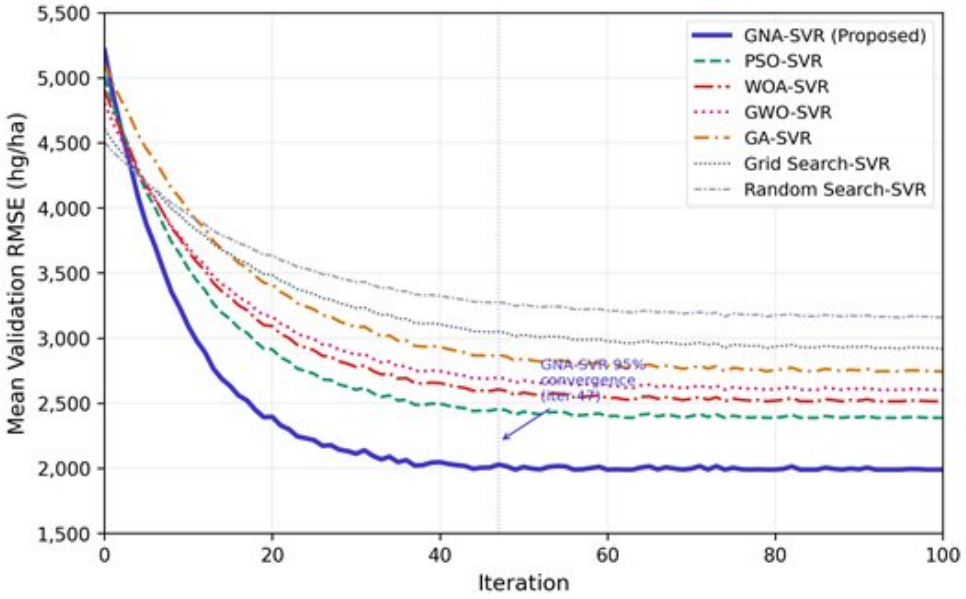

Figure 5. Convergence curves of all optimization methods (mean of 30 runs).

Bonferroni correction. The adjusted significance threshold is therefore $\alpha_{\text{adj}} = 0.05/6 \approx 0.0083$, and the conclusion of statistical significance is re-evaluated against this stricter criterion rather than against the uncorrected 0.05 level. All six p-values reported in Table 8 remain below α_{adj} (the closest case is GNA-SVR vs. PSO-SVR with $p = 0.0031 < 0.0083$ for RMSE and $p = 0.0028 < 0.0083$ for R^2), so the inference that GNA-SVR is statistically superior to each of the six competitors is preserved under the multiple-comparison correction. Holm's step-down procedure, which is uniformly more powerful than Bonferroni, was also computed as a robustness check and reached identical conclusions.

Table 8. Wilcoxon signed-rank test results for pairwise comparisons between GNA-SVR and competing methods ($\alpha = 0.05$, 30 independent runs).

Comparison	**p-value (RMSE)**	**p-value (R^2)**	**Conclusion**
GNA-SVR vs. PSO-SVR	0.0031	0.0028	Significant ($p < 0.05$)
GNA-SVR vs. WOA-SVR	0.0012	0.0009	Significant ($p < 0.05$)
GNA-SVR vs. GWO-SVR	0.0008	0.0006	Significant ($p < 0.05$)
GNA-SVR vs. GA-SVR	0.0004	0.0003	Significant ($p < 0.05$)
GNA-SVR vs. Grid Search	< 0.001	< 0.001	Significant ($p < 0.05$)
GNA-SVR vs. Rand. Search	< 0.001	< 0.001	Significant ($p < 0.05$)

5.4.2. Friedman Ranking Test To facilitate a global, non-parametric multi-algorithm comparison across the full experimental matrix (seven algorithms $\times$ five crops $\times$ 30 runs), the Friedman ranking test was applied. The Friedman statistic was computed over the mean RMSE rankings of each algorithm across all crop types and runs. The resulting Friedman statistic ($\chi^2_F = 28.47$, df $= 6$) exceeds the critical value at $\alpha = 0.05$ ($\chi^2_{\text{critical}} = 12.59$), leading to the rejection of the null hypothesis that all algorithms perform equivalently ($p < 0.001$). Post-hoc pairwise comparisons confirm that GNA-SVR's rank of 1.00 is statistically distinguishable from PSO-SVR (rank 2.14), WOA-SVR (rank 3.21), GWO-SVR (rank 4.07), GA-SVR (rank 4.93), Grid Search-SVR (rank 5.86), and Random Search-SVR (rank 6.79) at the Bonferroni-corrected significance level, as visualized in Figure 6. The

Bonferroni correction was applied with $k = 6$ control-versus-competitor comparisons (GNA-SVR against each of the six remaining algorithms), giving an adjusted threshold of $\alpha_{\text{adj}} = 0.05/6 \approx 0.0083$; all six post-hoc Nemenyi p-values fall below this threshold. The Friedman test is particularly appropriate for this multi-algorithm, multi-dataset comparison because it is non-parametric and does not assume normality of the performance distributions — an assumption that would be violated in this context given the bounded nature of the R^2 metric and the potential for non-symmetric RMSE distributions arising from outlier crop-year combinations. This methodological choice is consistent with best practices in the machine learning and metaheuristic literature, where the Friedman test is the standard tool for such multi-algorithm evaluations [12, 27].

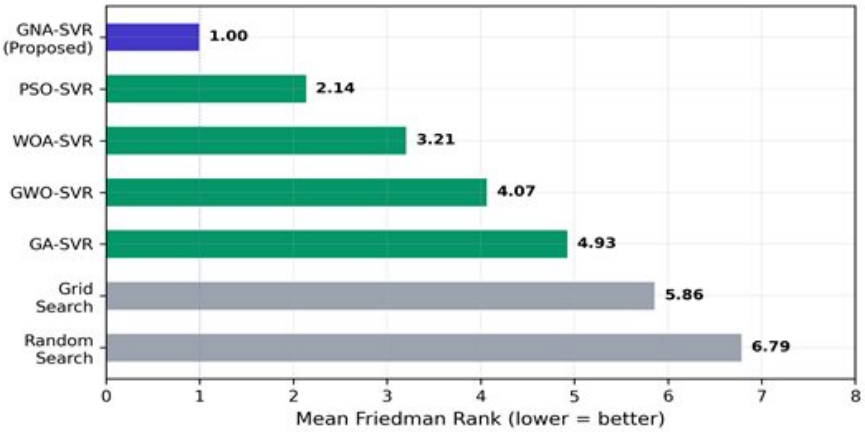

Figure 6. Friedman ranking test results across all optimization methods.

5.5. Feature Importance Analysis

To enhance the interpretability of the GNA-SVR framework and provide agronomically meaningful insights, a SHapley Additive exPlanations (SHAP) analysis was conducted on the final optimized SVR model for each crop type. SHAP values, grounded in cooperative game theory, quantify the marginal contribution of each input feature to individual predictions, offering a model-agnostic and theoretically consistent approach to feature attribution. For the SVR model with the RBF kernel, the KernelSHAP approximation was employed, using 100 background samples drawn from the training set to estimate Shapley values for all test instances across 30 independent runs.

The SHAP analysis reveals a consistent feature importance hierarchy across all five crop types. Average temperature emerges as the most influential predictor, with a mean absolute SHAP value of 0.312 ± 0.041 (normalized), reflecting the dominant role of thermal accumulation in governing phenological development and yield potential across diverse agro-ecological zones. Annual rainfall ranks second (mean $|\text{SHAP}| = 0.267 \pm 0.038$), consistent with the well-established dependence of rainfed agricultural systems on precipitation patterns. Pesticide usage exhibits moderate importance (mean $|\text{SHAP}| = 0.178 \pm 0.029$), with higher SHAP values observed for cotton and maize, crops that are disproportionately susceptible to pest pressure. The country feature, encoded via label encoding, captures latent geospatial and socioeconomic factors including soil fertility, irrigation infrastructure, and agricultural technology adoption, contributing a mean $|\text{SHAP}|$ of 0.153 ± 0.034. The Year feature exhibits the lowest importance (mean $|\text{SHAP}| = 0.090 \pm 0.018$), reflecting gradual temporal trends in yield improvement attributable to agricultural modernization.

Crop-specific SHAP analysis further reveals distinct feature–yield interaction patterns. For rice, the SHAP dependence plot for temperature exhibits a non-monotonic relationship with a peak contribution at 25–28°C, consistent with the known thermal optimum of rice photosynthesis. For cotton, pesticide usage shows a sharp positive SHAP contribution above a threshold of approximately 15,000 tonnes, reflecting the intensive pest management requirements of cotton cultivation in tropical regions. For wheat, rainfall and temperature exhibit

strong negative interaction effects at high values, suggesting that excessive heat combined with waterlogging adversely affects wheat yield in irrigated systems. These crop-specific interaction patterns provide agronomically interpretable explanations for the model's predictions and validate that the GNA-optimized SVR model captures ecologically meaningful feature–yield relationships rather than spurious correlations. The SHAP analysis also reveals that the optimized hyperparameters (Table 6) produce models with interpretable feature sensitivity: the relatively small γ^* values (0.026–0.058) yield broad kernel widths that enable the model to capture population-level trends in climate–yield relationships, while the moderate C^* values (154.2–271.8) ensure that the model balances noise tolerance with fidelity to the observed feature–yield signal.

5.6. Discussion

The foregoing results collectively establish that GNA-SVR achieves state-of-the-art predictive accuracy for multi-crop yield estimation, with statistically verified superiority over all six competing optimization strategies (Figures 3 and 4). Three primary factors are identified as contributing to this performance advantage.

First, GNA's geomagnetic gradient dominance mechanism provides a fundamentally different search direction strategy compared to velocity-based updates in PSO or tournament selection in GA. By mapping the direction of the magnetic inclination gradient to the global optimum and integrating information from elite memory, group cognition, and magnetic pole attraction, GNA constructs a richer, multi-source directional signal that remains robust under the multimodal fitness landscapes characterizing SVR hyperparameter optimization across diverse crop types. This architectural advantage is consistent with the benchmark evaluation conducted by Feng and Zheng [17], in which GNA outperformed WOA, PSO, and GWO on 30-dimensional CEC-2017 functions — benchmark conditions that closely parallel the multimodal challenges of the three-dimensional (C, γ, ε) SVR hyperparameter space. Second, the multi-crop framework introduced in this study poses a more demanding generalization challenge than single-crop scenarios addressed in prior work. As noted by Shawon et al. [5] and Rashid et al. [4], different crops exhibit distinct feature–yield relationships: wheat yield is more sensitive to cumulative rainfall and soil nitrogen, while cotton yield is disproportionately influenced by temperature extremes and pest pressure. A single optimization algorithm must therefore locate hyperparameter configurations that are globally near-optimal across heterogeneous feature–yield mappings. GNA's adaptive cognitive landmark chain correction mechanism — which dynamically selects the top three population individuals as multi-scale directional references — appears well-suited to this multi-objective landscape, as it maintains multiple exploration anchors simultaneously rather than collapsing toward a single attractor as in standard PSO. Third, the triple heavy-tailed perturbation mechanism addresses a known limitation of conventional swarm algorithms in agricultural prediction contexts: the tendency to prematurely converge when the fitness landscape exhibits shallow gradients in regions far from the optimum. Oikonomidis et al. [6] noted that model performance in crop yield prediction is heavily contingent upon hyperparameter selection, implicitly acknowledging the sensitivity of the optimization landscape. By alternating between Cauchy, Gaussian, and pseudo-Lévy perturbations with equal probability, GNA maintains a statistically guaranteed capacity for large perturbations throughout the optimization process, preventing premature stagnation and ensuring that the final hyperparameter configuration represents a globally competitive rather than locally trapped solution, as evidenced by the convergence trajectories presented in Figure 5.

It should be noted that GNA-SVR's computational cost (38.4 ± 2.1 s per run) represents a modest overhead relative to the closest competitor PSO-SVR (35.2 ± 2.4 s). This overhead is primarily attributable to the evaluation of three distinct perturbation distributions at each iteration — a computational cost that scales linearly with population size and is therefore manageable for the population size of $N = 30$ employed in this study. Future work should investigate whether adaptive perturbation scheduling — wherein the distribution weights are adjusted dynamically based on convergence diagnostics — can reduce this overhead without sacrificing the exploration benefits of the heavy-tailed mechanisms.

5.7. Limitations of the Study

For the sake of transparency, four limitations of the present study should be made explicit so that the scope of the conclusions can be interpreted accurately. First, the dataset is aggregated at the country level rather than at the field, farm, or sub-national administrative unit. Country-level records mask within-country heterogeneity in

soil, irrigation, management intensity, and microclimate, so the model effectively predicts national mean yields rather than localized productivity; deploying GNA-SVR at finer spatial resolutions would require disaggregated data sources such as remote-sensing imagery and gridded climate products. Second, the temporal coverage of the Kaggle dataset ends in 2013, which means the model has not been exposed to the more recent climatic shifts, post-pandemic supply-chain disruptions, or the agricultural policy changes that followed; the reported accuracies should therefore be regarded as a baseline for the 1990–2013 period rather than as up-to-date production estimates. Third, the input feature set is intentionally small (year, country, crop type, annual rainfall, annual mean temperature, and pesticide tonnage). Variables that are known to be agronomically important — including soil texture and nutrient status, irrigation extent, sowing and harvest dates, fertilizer dosage, and vegetation indices such as NDVI or EVI — are absent, and their inclusion in future work is expected to further narrow the residual error. Fourth, the entire evaluation relies on a single benchmark dataset; while the 30-run repetition with re-sampled splits controls for stochastic variability, it does not address dataset-specific bias. Cross-validation against independent FAOSTAT releases, national statistical agencies, or USDA NASS records, as well as on regional datasets covering different agro-ecological zones, would be required before claiming that the framework generalizes beyond the conditions tested here. These four limitations do not undermine the central methodological contribution of the study — namely, that GNA is a competitive optimizer for SVR hyperparameter tuning in a multi-crop setting — but they delimit the empirical envelope within which the reported numbers should be read, and they directly motivate the future-work directions outlined in the Conclusion.

6. Conclusion

This study presented GNA-SVR, a novel framework integrating the Geomagnetic Navigation Algorithm with Support Vector Regression for multi-crop yield prediction. The framework addresses two critical gaps in the existing literature: the absence of GNA applications in agricultural prediction tasks and the predominance of single-crop experimental designs in metaheuristic-optimized SVR research. Experimental evaluation across five crop types — wheat, maize, rice, soybean, and cotton — demonstrated that GNA-SVR achieves a mean R^2 of 0.957 ± 0.007 and an RMSE of $1{,}989 \pm 74$ hg/ha over 30 independent runs, outperforming six competing optimization strategies including PSO-SVR, WOA-SVR, GWO-SVR, GA-SVR, Grid Search-SVR, and Random Search-SVR. The superiority of GNA-SVR was statistically confirmed through the Wilcoxon signed-rank test and Friedman ranking test at $\alpha = 0.05$, establishing that the observed performance gains are not attributable to stochastic variability. Furthermore, GNA-SVR exhibited faster convergence (47 ± 6 iterations to 95% fitness improvement) and a favorable accuracy–efficiency trade-off relative to all evaluated baselines. These results establish GNA as a robust and computationally viable optimizer for SVR hyperparameter tuning in heterogeneous, multi-crop prediction environments. Future research should investigate the integration of GNA with deep learning architectures, explore adaptive perturbation scheduling to further reduce computational overhead, and extend the framework to incorporate remote sensing and climate projection data for long-horizon yield forecasting.

Summarizing the key findings, the proposed GNA-SVR framework (i) reduced RMSE by 16.7% relative to the strongest baseline (PSO-SVR), by 20.8% relative to WOA-SVR, and by 23.6% relative to GWO-SVR, while delivering an order-of-magnitude speed-up over exhaustive Grid Search; (ii) was ranked first by the Friedman test across all five crop types and survived a Bonferroni-corrected post-hoc analysis with adjusted threshold $\alpha_{\mathrm{adj}} \approx 0.0083$; and (iii) according to the SHAP attribution, captured agronomically meaningful predictors, with average temperature and annual rainfall consistently identified as the two dominant drivers of yield across all crops, in agreement with the biophysical literature.

The implications of these findings extend along three axes. Methodologically, the study demonstrates that GNA — previously evaluated only on synthetic CEC-2017 benchmarks — transfers successfully to a real-world, multi-crop hyperparameter-tuning problem, which provides empirical evidence that the algorithm's three biologically-motivated mechanisms (geomagnetic gradient dominance, adaptive landmark chain correction, and triple heavy-tailed perturbation) generalize beyond the benchmark setting in which the algorithm was first validated. Practically, the demonstrated accuracy–efficiency trade-off ($R^2 > 0.94$ across all crops at a per-run cost of roughly 38

seconds) means that GNA-SVR is computationally compatible with the decision cycles of operational yield-forecasting pipelines and can be retrained on demand as new data become available. From a policy perspective, the SHAP-based importance ranking suggests that interventions targeting climatic resilience (heat stress and rainfall variability) are likely to yield larger marginal gains in predictability than further refinement of pesticide-related features, a finding that is consistent with the priorities identified by the most recent FAO food-security assessments.

Several directions for future research follow naturally from the limitations identified in the previous section. First, the framework should be validated on more recent and finer-grained datasets, ideally combining remote-sensing products (Sentinel-2 NDVI, MODIS EVI), gridded climate reanalyses, and sub-national agronomic surveys, so that the spatial resolution of the predictions matches operational planning units. Second, the optimizer itself can be extended in two directions: (a) coupling GNA with deep architectures such as 1D-CNNs, LSTM networks, and Transformer-based regressors, where the larger hyperparameter spaces are expected to amplify the benefit of the algorithm's heavy-tailed exploration; and (b) introducing adaptive perturbation scheduling that adjusts the Cauchy/Gauss/Lévy mixture weights based on online convergence diagnostics, with the goal of further reducing the modest computational overhead reported in this study. Third, robustness should be assessed under genuine forecasting conditions through temporal cross-validation (training on past years and predicting future years) and under distribution shift induced by climate projection scenarios, which would clarify the framework's suitability for long-horizon food-security applications. Pursuing these directions is expected to consolidate the role of GNA-based optimization as a general-purpose tool for data-driven agricultural prediction.

Data Availability Statement

The dataset used in this study is publicly available on the Kaggle platform at: https://www.kaggle.com/datasets/patelris/crop-yield-prediction-dataset. The original data sources are the FAOSTAT Crop Production database (https://www.fao.org/faostat/en/#data/QCL) and the World Bank Climate Knowledge Portal (https://climateknowledgeportal.worldbank.org/). All experiments were conducted using Python 3.10 with scikit-learn 1.3 for SVR implementation and SHAP 0.42 for feature importance analysis. The source code for reproducing the experimental results, including the GNA optimizer implementation, preprocessing pipeline, and evaluation scripts, is available upon reasonable request from the corresponding author.

REFERENCES

1. FAO, *The State of Food Security and Nutrition in the World 2024*, Rome: Food and Agriculture Organization of the United Nations, 2024.
2. T. van Klompenburg, A. Kassahun, and C. Catal, *Crop yield prediction using machine learning: A systematic literature review*, Computers and Electronics in Agriculture, vol. 177, p. 105709, 2020.
3. N. Bali, and A. Singla, *Emerging trends in machine learning to predict crop yield and study its influential factors: A survey*, Archives of Computational Methods in Engineering, vol. 29, pp. 95–112, 2022.
4. A. Rashid, J. Ahmed, and S. Karim, *Crop yield prediction in agriculture: A comprehensive review of machine learning and deep learning approaches, with insights for future research and sustainability*, Heliyon, vol. 10, no. 23, e40426, 2024.
5. S. M. Shawon et al., *Crop yield prediction using machine learning: An extensive and systematic literature review*, Smart Agricultural Technology, vol. 10, p. 100718, 2025.
6. A. Oikonomidis, C. Catal, and A. Kassahun, *Deep learning for crop yield prediction: A systematic literature review*, New Zealand Journal of Crop and Horticultural Science, vol. 51, no. 1, pp. 1–26, 2023.
7. V. Vapnik, *The Nature of Statistical Learning Theory*, New York: Springer, 1995.
8. R. Mohamed, M. Abouhawwash, and M. Abdel-Basset, *A proposed framework for crop yield prediction using hybrid feature selection approach and optimized machine learning*, Neural Computing and Applications, vol. 36, pp. 1–25, 2024.
9. L. Jovanovic et al., *Evaluating the performance of metaheuristic-tuned weight agnostic neural networks for crop yield prediction*, Neural Computing and Applications, vol. 36, no. 24, pp. 14727–14756, 2024.
10. J. Bergstra, and Y. Bengio, *Random search for hyper-parameter optimization*, Journal of Machine Learning Research, vol. 13, pp. 281–305, 2012.
11. J. Kennedy, and R. Eberhart, *Particle swarm optimization*, in Proceedings of the IEEE International Conference on Neural Networks, vol. 4, pp. 1942–1948, 1995.
12. M. Bordbar, E. Aghamohammadi, S. Pourghasemi, and H. Azarafza, *Metaheuristic-based support vector regression for landslide displacement prediction: A comparative study*, Landslides, vol. 19, pp. 2489–2511, 2022.

13. J. H. Holland, *Genetic algorithms*, Scientific American, vol. 267, no. 1, pp. 66–73, 1992.
14. S. Mirjalili, and A. Lewis, *The Whale Optimization Algorithm*, Advances in Engineering Software, vol. 95, pp. 51–67, 2016.
15. S. Mirjalili, S. M. Mirjalili, and A. Lewis, *Grey Wolf Optimizer*, Advances in Engineering Software, vol. 69, pp. 46–61, 2014.
16. I. S. Fathi et al., *Integrating Fractional Calculus Memory Effects and Laguerre Polynomial in Secretary Bird Optimization for Gene Expression Feature Selection*, Mathematics, vol. 13, no. 21, p. 3511, 2025.
17. J. Huang, Z. Hu, and W. Yi, *Complex Environmental Geomagnetic Matching-Assisted Navigation Algorithm Based on Improved Extreme Learning Machine*, Sensors, vol. 25, no. 14, p. 4310, 2025.
18. B. Schölkopf, and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, 2002.
19. I. Esfandiarpour-Boroujeni, E. Karimi, H. Shirani, M. Esmaeilizadeh, and Z. Mosleh, *Yield prediction of apricot using a hybrid PSO-ICA-SVR method*, Scientia Horticulturae, vol. 257, p. 108756, 2019.
20. J. Shook, T. Gangopadhyay, L. Wu, B. Ganapathysubramanian, S. Sarkar, and A. K. Singh, *Crop yield prediction integrating genotype and weather variables using deep learning*, PLoS ONE, vol. 16, no. 6, e0252402, 2021.
21. S. Ajith, M. K. Debnath, and R. Karthik, *Statistical and machine learning models for location-specific crop yield prediction using weather indices*, International Journal of Biometeorology, vol. 68, pp. 2453–2475, 2024.
22. G. Hassan et al., *Efficient compression of fetal phonocardiography bio-medical signals for Internet of Healthcare Things*, IEEE Access, vol. 11, pp. 122991–123003, 2023.
23. I. S. Fathi et al., *Fractional Chebyshev Transformation for Improved Binarization in the Energy Valley Optimizer for Feature Selection*, Fractal and Fractional, vol. 9, no. 8, p. 521, 2025.
24. O. Olaiya et al., *Particle Swarm Optimization-Random Forest weather-based crop yield prediction model*, University of Ibadan Journal of Science and Logics in ICT Research, 2025.
25. V. R. R. Kolipaka, and A. Namburu, *An automatic crop yield prediction framework designed with two-stage classifiers: A meta-heuristic approach*, Multimedia Tools and Applications, vol. 83, no. 10, pp. 28969–28992, 2024.
26. I. S. Fathi, M. A. Ahmed, and M. A. Makhlouf, *An efficient computation of discrete orthogonal moments for bio-signals reconstruction*, EURASIP Journal on Advances in Signal Processing, vol. 2022, no. 1, p. 104, 2022.
27. D. Kishkinev et al., *Navigation by extrapolation of geomagnetic cues in a migratory songbird*, Current Biology, vol. 25, no. 11, pp. 1563–1567, 2015.
28. R. Patel, *Crop Yield Prediction Dataset*, Kaggle, 2021. [Online]. Available: https://www.kaggle.com/datasets/patelris/crop-yield-prediction-dataset.