STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **x**, Month 202x, pp 0–17.
Published online in International Academic Press (www.IAPress.org)

Joint Irrigation and Fertilization Scheduling in Multi-Crop Agricultural Systems Using the Gliding Snake Optimizer

OLA Aljaafreh [1,*], Orowah Mahmoud Abd Al-Slaibi [1], Abbas A. Metawea [1,2], Islam S. Fathi [3,4]

[1]*Agricultural Extension and Marketing Department, Faculty of Agriculture, Ajloun National University, Jordan*
[2]*Department of Agricultural Economics, Faculty of Agriculture, Al-Azhar University, Cairo, P.O.Box 11651, Egypt*
[3]*Department of Computer Science, Faculty of Information Technology, Ajloun National University, P.O.43, Ajloun-26810, Jordan*
[4]*Department of Information Systems, Al Alson Higher Institute, Cairo 11762, Egypt*

Abstract Efficient irrigation and fertilization scheduling for multi-crop farmlands represents a critical challenge in modern precision agriculture, particularly under increasing water scarcity and food security pressures. This paper proposes the application of the Gliding Snake Optimizer (GSO), a recently introduced bio-inspired meta-heuristic algorithm, to solve the combined crop irrigation and fertilization scheduling problem. The proposed model encodes water volumes, nutrient dosages (N, P, K), and event timing into a unified decision vector and optimizes a composite objective function that simultaneously maximizes normalized crop yield while minimizing water cost, fertilizer expenditure, and phenological timing penalties. Seven agronomic and resource constraints are enforced via an adaptive penalty mechanism. Extensive experiments on a benchmark scenario comprising four representative crops (wheat, corn, tomato, olive) demonstrate that GSO achieves statistically superior performance compared to four established algorithms—Genetic Algorithm (GA), Particle Swarm Optimization (PSO), Whale Optimization Algorithm (WOA), and Grey Wolf Optimizer (GWO)—across all evaluation metrics. GSO reduces water consumption by up to 23.4% and improves simulated yield by 18.7% relative to a conventional scheduling baseline. Wilcoxon signed-rank tests confirm that performance differences are statistically significant ($p < 0.05$) in 30 independent runs. The results demonstrate the strong suitability of GSO for constrained agricultural scheduling optimization problems.

Keywords Gliding Snake Optimizer; Crop Scheduling; Irrigation Optimization; Fertilizer Management; Meta-Heuristic Algorithms; Soil Health; Agricultural Sustainability.

AMS 2010 subject classifications 90C59, 68T20

DOI: 10.19139/soic-2310-5070-3960

1. Introduction

Agriculture remains the single largest consumer of freshwater resources worldwide, accounting for approximately 70% of total global withdrawals [1]. In the face of mounting population growth and shifting climatic patterns, the efficient allocation of water and nutrients across cropping systems has emerged as a paramount concern for ensuring food security and environmental sustainability. Inefficient irrigation practices are widely recognized as a major source of water waste, while excessive or poorly timed fertilization contributes to soil degradation, nutrient leaching into groundwater, and eutrophication of surface water bodies [2]. The concurrent optimization of irrigation volumes, multi-nutrient fertilizer dosages, and the phenological timing of application events across heterogeneous multi-crop fields constitutes a high-dimensional, nonlinearly constrained combinatorial optimization problem that rapidly exceeds the capacity of classical exact solution methods [3]. Traditional approaches to agricultural

*Correspondence to: OLA Aljaafreh, Agricultural Extension and Marketing Department, Faculty of Agriculture, Ajloun National University, Jordan.

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

scheduling have predominantly relied on linear programming, dynamic programming, and rule-based expert systems [4, 5]. While these methods have provided valuable foundational insights, they typically address a single optimization objective under restrictive simplifying assumptions, thereby failing to capture the intricate nonlinear interactions among water stress, nutrient availability, soil health dynamics, and crop yield response [5]. Moreover, these classical formulations become computationally intractable when extended to multi-crop, multi-period scenarios with interdependent resource constraints [6]. Consequently, there has been a growing impetus within the precision agriculture community to adopt computationally intelligent methods that can navigate complex, multi-modal search landscapes without requiring explicit problem convexity or gradient information [7, 29].

Meta-heuristic optimization algorithms have demonstrated considerable promise in addressing these challenges. Evolutionary and swarm intelligence methods—including Genetic Algorithms (GA), Particle Swarm Optimization (PSO), the Whale Optimization Algorithm (WOA), and the Grey Wolf Optimizer (GWO)—have been applied to various agricultural resource management problems with encouraging results [8, 9]. However, the performance of any given meta-heuristic is inherently problem-dependent, as established by the No Free Lunch theorem [10], which stipulates that no single algorithm can dominate across all optimization landscapes. This theoretical insight motivates the continued development and evaluation of novel algorithms on domain-specific problem formulations. The Snake Optimizer (SO), introduced by Hashim and Hussien in 2022 [11], is a recent swarm intelligence algorithm that mathematically models the foraging, fighting, and mating behaviors of snakes. The SO divides its population into male and female subgroups whose interactions are governed by environmental parameters—specifically, temperature and food availability that modulate the transition between exploration and exploitation phases. Extensive benchmarking on CEC test functions has confirmed its competitive performance relative to established meta-heuristics [11]. Nonetheless, several studies have identified inherent limitations in the original SO, including susceptibility to premature convergence, insufficient population diversity maintenance, and slow convergence on high-dimensional multimodal problems [12, 13]. To address these shortcomings, a number of enhanced variants have been proposed in the literature. Yao et al. [14] introduced the Enhanced Snake Optimizer (ESO), incorporating adaptive Levy flight mechanisms to improve exploration capability and escape from local optima. Peng et al. [15] proposed a Multi-strategy Improved Snake Optimizer (MISO) assisted by population crowding analysis, which demonstrated superior solution quality on engineering design benchmarks. Lu et al. [16] developed an Improved Snake Optimizer (ISO) employing mirror opposition-based learning and a novel evolutionary population dynamics model for SVM parameter selection. More recently, Alawad et al. [17] presented a Hybrid Snake Optimizer Algorithm (HSOA) that integrates oppositional-mutual learning and dynamic polynomial mutation to enhance population diversity and convergence behavior.

The Gliding Snake Optimizer (GSO), which forms the algorithmic foundation of the present study, augments the standard SO framework with two principal enhancements: (i) an adaptive Levy flight mechanism in the exploration phase that generates heavy-tailed random steps, enabling effective escape from local optima; and (ii) a cosine-modulated gliding mechanism in the exploitation phase that provides smoother and more diverse trajectories toward the global optimum [12, 14]. These modifications have yielded demonstrably superior convergence characteristics on constrained engineering optimization benchmarks. However, despite these promising theoretical properties, neither GSO nor any other SO variant has been applied to the agricultural irrigation and fertilization scheduling problem, representing a significant gap in the existing literature. This paper addresses this gap by proposing a comprehensive application of the Gliding Snake Optimizer to the multi-crop irrigation and fertilization scheduling problem. The contributions of this study are fourfold:

- A unified mixed-variable mathematical model is formulated for multi-crop irrigation and fertilization scheduling that integrates a quadratic water-yield response function, multi-nutrient dosage optimization, soil health constraints, and phenological timing penalties.
- A tailored encoding scheme, mixed-variable handling strategy, boundary repair mechanism, and adaptive penalty-based constraint management approach are developed to adapt GSO for this agricultural problem class.
- A systematic experimental comparison of GSO against four established meta-heuristics is conducted on a realistic four-crop benchmark scenario.

- A detailed analysis of the GSO-optimized schedules is presented in terms of water savings, yield improvement, soil health preservation, fertilizer cost reduction, and timing alignment with critical crop growth windows.

The remainder of this paper is structured as follows. Section 2 reviews the related literature. Section 3 formalizes the mathematical problem formulation. Section 4 describes the GSO algorithm and its domain-specific adaptations. Section 5 presents the experimental setup and results. Section 6 discusses findings and limitations. Section 7 concludes the paper.

2. Related Work

The optimization of agricultural water and nutrient management has been a subject of sustained research interest for over three decades, driven by the imperative to produce more food with fewer resources under increasingly constrained environmental conditions. Early foundational works employed mathematical programming formulations to model crop-water relationships. Rao et al. [18] developed a dated water-production function for use in irrigated agriculture, providing a seminal yield-response model that has been widely adopted and adapted in subsequent irrigation optimization studies. Dynamic programming approaches were subsequently explored for multi-period irrigation scheduling; however, these methods proved computationally intractable when constraints [6] extended to multi-crop scenarios with coupled resource. The emergence of process-based crop simulation models marked a significant paradigm shift in irrigation and fertilization research. The DSSAT (Decision Support System for Agrotechnology Transfer) modeling platform has become one of the most widely adopted tools for simulating crop growth responses to management practices across diverse agroclimatic conditions [19]. Bai et al. [20] combined the DSSAT crop simulation model with a Genetic Algorithm to jointly optimize irrigation and fertilizer schedules for maize in China, reporting yield improvements of 1.9–2.6% and economic benefits of 7.3–8.9% compared with conventional management practices. This work demonstrated the viability of coupling process-based crop models with evolutionary optimization algorithms for practical agricultural decision support. Wu et al. [21] extended the simulation-optimization paradigm by modifying the AquaCrop model to incorporate nitrogen stress simulation, thereby enabling joint optimization of irrigation and fertilization schedules under simultaneous water and nitrogen constraints. Their framework, designated JSOIFS, employed multi-objective programming with NSGA-II to optimize both crop yield and water use efficiency for seed maize in the Shiyanghe River Basin. The study established that integrating nitrogen dynamics into irrigation optimization models yields substantially more realistic and resource-efficient management recommendations. Multi-objective formulations have gained increasing traction in recent years. Lyu et al. [22] coupled AquaCrop-OSPy with the NSGA-III algorithm for multi-objective winter wheat irrigation strategy optimization. Okola et al. [23] applied multi-objective evolutionary algorithms to optimize the trade-offs among food production, energy generation, and water conservation at the food-energy-water nexus scale. Zhao et al. [5] provided a comprehensive review of simulation-optimization models for irrigation scheduling, identifying meta-heuristic algorithms as the increasingly dominant solution strategy and highlighting that the simultaneous optimization of water volumes, multi-nutrient doses, and phenological timing across heterogeneous multi-crop fields under soil health constraints remains underexplored.

Meta-heuristic optimization algorithms have found broad and growing application across the precision agriculture domain. Genetic Algorithms, among the earliest and most widely adopted meta-heuristics, have been applied to diverse agricultural problems including irrigation network design [24], crop pattern planning under climate change uncertainty [6], and the coupled optimization of planting schedules and water allocation strategies. Particle Swarm Optimization has been applied to food-water nexus trade-off optimization [23] and multi-objective crop production planning. The Whale Optimization Algorithm [25] and the Grey Wolf Optimizer [26] have demonstrated competitive performance on constrained agricultural resource allocation tasks. More recently, the convergence of computational intelligence with IoT-based sensor technologies has opened new frontiers for real-time adaptive agricultural management. Bwambale et al. [27] presented a model predictive control framework for smart irrigation scheduling, demonstrating 29% water savings compared with manual irrigation control. Akbari et

al. [28] developed a surface irrigation simulation-optimization model that explicitly couples hydraulic modeling with evolutionary search to determine optimal furrow irrigation parameters.

The Snake Optimizer (SO) was introduced by Hashim and Hussien [11] as a swarm intelligence algorithm modeling the foraging, fighting, and mating behaviors of snake populations. The algorithm divides its population into male and female subgroups, whose spatial search dynamics are governed by two environmental parameters: temperature and food availability. When conditions are unfavorable, agents engage in exploratory random walks; under favorable conditions, agents transition to exploitation via fighting and mating dynamics.

Since its introduction, numerous enhanced variants have been proposed to address documented limitations, particularly premature convergence and insufficient diversity maintenance [12, 13]. Yao et al. [14] proposed the Enhanced Snake Optimizer (ESO) with adaptive Levy flight perturbations. Peng et al. [15] introduced a Multi-strategy Improved SO (MISO) assisted by population crowding analysis. Lu et al. [16] developed an Improved SO integrating mirror opposition-based learning. Alawad et al. [17] presented a Hybrid Snake Optimizer combining oppositional-mutual learning with dynamic polynomial mutation. Beyond engineering design benchmarks, SO variants have also been successfully deployed in diverse application domains: Karam et al. [30] applied an improved snake algorithm to optimize photovoltaic maximum power point tracking under varying irradiance conditions, while Zheng et al. [31] developed a compact Snake Optimization Algorithm for WKNN fingerprint localization, demonstrating the algorithm's versatility across renewable-energy and indoor-positioning tasks. The Gliding Snake Optimizer (GSO), which constitutes the algorithmic basis of the present work, extends the standard SO through two principal enhancements: an adaptive Levy flight mechanism in the exploration phase generating heavy-tailed random steps, and a cosine-modulated gliding mechanism in the exploitation phase providing smoother convergence trajectories. To the best of the authors' knowledge, no prior study has applied the GSO or any enhanced Snake Optimizer variant to the crop irrigation and fertilization scheduling problem, motivating the present study.

3. Problem Formulation

3.1. Problem Description

Consider a farm comprising n heterogeneous crops indexed $i = \{1, \ldots, n\}$. Each crop occupies a field of area A_i (ha) and requires irrigation and fertilization events over a growing season of T days. The objective is to determine the optimal water application volumes W_i (m^3/ha), fertilizer doses F_{ij} (kg/ha) for nutrients $j = \{N, P, K\}$, and event timing T_i^w (irrigation day) and T_i^f (fertilization day) such that a composite agricultural performance index is maximized while satisfying all agronomic and resource constraints.

3.2. Decision Variables

The complete decision vector encodes all scheduling decisions for the n crops into a single continuous representation:

$$X = [W_1, \ldots, W_n \mid F_{11}, \ldots, F_{n3} \mid T_1^w, \ldots, T_n^w \mid T_1^f, \ldots, T_n^f] \tag{1}$$

For n crops with three nutrient types (N, P, K), the total problem dimensionality is $d = 7n$. In the four-crop benchmark instance ($n = 4$), the decision space consists of $d = 28$ continuous and integer-valued variables.

3.3. Objective Function

The composite fitness function $F(X)$ to be maximized is formulated as a weighted aggregation of four competing objectives:

$$F(X) = w_1 \cdot Y_{\text{norm}}(X) - w_2 \cdot C_w(X) - w_3 \cdot C_f(X) - w_4 \cdot D(X) \tag{2}$$

where $w_1 + w_2 + w_3 + w_4 = 1$ are priority weights. Default values are set to $w_1 = 0.45$ (yield maximization), $w_2 = 0.20$ (water cost minimization), $w_3 = 0.20$ (fertilizer cost minimization), and $w_4 = 0.15$ (timing penalty minimization).

3.3.1. Normalized Yield Function Individual crop yield follows a quadratic water-response model penalized by deviation from the critical growth day T_i^*:

$$Y_i = Y_i^{\max} \cdot \left[\alpha_i \cdot \left(\frac{W_i}{W_i^{\text{opt}}} \right) - \beta_i \cdot \left(\frac{W_i}{W_i^{\text{opt}}} \right)^2 \right] \cdot \exp\left(-\gamma_i \cdot |T_i^w - T_i^*| \right) \tag{3}$$

The normalized aggregate yield is: $Y_{\text{norm}} = (1/n) \cdot \sum_i Y_i / Y_i^{\max}$.

Validity range of the quadratic water-response model. Let $x_i = W_i / W_i^{\text{opt}}$ denote the normalized water application. The bracketed term in equation (3) takes the form $\alpha_i \cdot x_i - \beta_i \cdot x_i^2$, which is non-negative only over the interval $x_i \in [0, \alpha_i/\beta_i]$ and reaches its maximum at $x_i^* = \alpha_i/(2\beta_i)$. Beyond $x_i = \alpha_i/\beta_i$, the surrogate would predict negative yield, which is biologically meaningless and corresponds to the over-irrigation regime in which water-logging, leaching, and reduced root aeration cause yield collapse rather than continued growth. Two safeguards therefore enforce that the optimization remains inside the agronomically valid range. First, the surrogate is clipped from below at zero, i.e. $Y_i = \max\{0, Y_i^{\max} \cdot (\alpha_i \cdot x_i - \beta_i \cdot x_i^2) \cdot \exp(-\gamma_i \cdot |T_i^w - T_i^*|)\}$, which prevents the optimizer from exploiting any spurious negative-yield region. Second, the per-crop water bound constraint C2 ($W_i^{\min} \le W_i \le W_i^{\max}$) is calibrated so that $W_i^{\max}/W_i^{\text{opt}} < \alpha_i/\beta_i$ for every crop in the benchmark; for the parameters in Table 2, α_i/β_i equals 2.30 for wheat, 2.18 for corn, 2.08 for tomato, and 2.44 for olive, while the imposed upper bounds correspond to $W_i^{\max}/W_i^{\text{opt}} = 1.40$ in all cases. The decision space therefore lies entirely on the rising or descending branches of the parabola, and the surrogate retains its agronomic interpretation throughout the optimization.

3.3.2. Water and Fertilizer Cost Functions

$$C_w = p_w \cdot \sum_i A_i \cdot W_i; \qquad C_f = \sum_i \sum_j p_f^j \cdot A_i \cdot F_{ij} \tag{4}$$

where p_w is the unit water price (USD/m^3) and p_f^j is the unit price of nutrient j (USD/kg).

3.3.3. Timing Penalty

$$D = \frac{1}{n} \cdot \sum_i \frac{\max\left(0, |T_i^w - T_i^*| - \Delta T_{\max}\right)}{T} \tag{5}$$

3.4. Constraints

Seven agronomic and resource constraints govern the feasibility of candidate solutions, summarized in Table 1.

Table 1. Constraint definitions for the multi-crop scheduling problem.

ID	Constraint	Expression
C1	Total water budget	$\sum_i A_i \cdot W_i \le W_{\text{total}}^{\max}$
C2	Per-crop water bounds	$W_i^{\min} \le W_i \le W_i^{\max}, \forall i$
C3	Fertilizer budget	$\sum_i \sum_j p_f^j \cdot A_i \cdot F_{ij} \le$ Budget
C4	Per-crop fertilizer bounds	$F_{ij}^{\min} \le F_{ij} \le F_{ij}^{\max}, \forall i, j$
C5	Soil health index (SHI)	$\text{SHI}_i = 1 - \delta \cdot \sum_j F_{ij}^2 \ge \text{SHI}_{\min}, \forall i$
C6	Timing feasibility	$1 \le T_i^w \le T$ and $1 \le T_i^f \le T, \forall i$
C7	N:P ratio (soil safety)	$F_i^N / F_i^P \le \text{NPK}_{\text{ratio}}^{\max}, \forall i$

The expression $\text{SHI}_i = 1 - \delta \cdot \sum_j F_{ij}^2$ adopted in this work is a normalized over-fertilization penalty index rather than a comprehensive soil quality metric in the sense of Andrews et al. or Karlen et al. The quadratic accumulation term $\sum F_{ij}^2$ is grounded in the agronomic principle that the negative environmental externalities

of fertilization—including nutrient leaching, soil acidification, electrical-conductivity build-up, and salinization—scale super-linearly with applied dose, a relationship documented in nitrogen-balance and soil-degradation studies for intensively cultivated systems. The coefficient δ is calibrated such that the quadratic penalty becomes binding only when total nutrient inputs exceed agronomically recommended thresholds, and the lower bound $\mathrm{SHI}_{\min} = 0.60$ corresponds to the dosage range above which empirical studies report measurable declines in soil organic-matter dynamics and microbial activity. To avoid over-claiming, the constraint is therefore best interpreted as an over-fertilization safeguard that operates as a soft proxy for short-term chemical degradation, rather than a multi-component soil quality index integrating physical, chemical, and biological indicators. Linking the optimization framework to a fully validated soil-quality model (e.g., the Soil Management Assessment Framework, SMAF) is identified as a priority for follow-up work.

3.5. Constraint Handling

Constraint violations are incorporated into the fitness evaluation through an augmented penalty function:

$$F'(X) = F(X) - \lambda \cdot \sum_{c} \max\left(0, g_c(X)\right) \tag{6}$$

where $g_c(X) \geq 0$ represents the violation magnitude of constraint c, and $\lambda = 10^4$ is a sufficiently large penalty coefficient ensuring that feasible solutions consistently dominate infeasible ones in the fitness ranking [3].

To verify that the choice of the static coefficient $\lambda = 10^4$ does not bias the reported results, the four-crop benchmark was re-run for $\lambda \in \{10^2, 10^3, 10^4, 10^5, 10^6\}$, with all other settings held constant. Mean penalized fitness varied within a narrow band: small λ values (10^2–10^3) admitted a non-negligible fraction of marginally infeasible solutions and degraded the mean fitness by 1.8–3.4%, whereas very large values (10^5–10^6) caused premature stalling at the feasibility boundary and reduced the mean fitness by 0.6–1.1%. The setting $\lambda = 10^4$ consistently produced the highest mean fitness with the lowest standard deviation, confirming that the reported results are robust within a wide neighbourhood of this value. In addition, the static-penalty formulation was benchmarked against three alternative constraint-handling strategies: (i) Deb's feasibility rules, in which feasible candidates always dominate infeasible ones and ties are broken by aggregate violation; (ii) the ε-constrained method, where the feasibility threshold is contracted across iterations; and (iii) a dynamic penalty schedule of the form $\lambda(t) = \lambda_0(t/T_{\max})^2$. The static penalty achieved the highest mean fitness; Deb's rules ranked second with comparable feasibility but slightly slower convergence; the ε-constrained variant attained the lowest constraint-violation rate but at a small fitness cost; and the dynamic schedule offered no measurable advantage on this benchmark. The static formulation was therefore retained as the default while acknowledging Deb's rules as a competitive parameter-free alternative.

4. The Gliding Snake Optimizer (GSO)

4.1. Biological Inspiration

The Gliding Snake Optimizer draws its biological inspiration from the locomotion and predatory behaviors of snakes, with particular emphasis on the gliding ability observed in arboreal species of the genus *Chrysopelea* (commonly known as flying snakes). Each agent in the population represents a candidate solution encoded as a position vector in the decision space. The algorithm alternates between two principal behavioral phases: an exploration phase, active under conditions of low temperature and scarce food, characterized by Levy-flight random walks that enable broad coverage of the search space; and an exploitation phase, active under warm and food-rich conditions, employing cosine-modulated gliding movements toward the best-known solution [11].

4.2. Population Initialization

The initial population of S agents is distributed across the decision space using Latin Hypercube Sampling (LHS) to ensure uniform coverage:

$$X_k = \text{LB} + r \cdot (\text{UB} - \text{LB}), \quad r \sim U(0,1), \quad k = 1, \ldots, S \tag{7}$$

where LB and UB are the lower and upper bounds of the decision vector. Infeasible individuals are repaired by projecting them onto the nearest feasible boundary.

4.3. Fitness Evaluation

Each agent X_k is evaluated using the penalized fitness function $F'(X_k)$. At each iteration, the best agent X_{best} and worst agent X_{worst} are identified to guide the search dynamics.

4.4. Position Update

4.4.1. Exploration Phase (Low Temperature / Food Scarce) When the temperature factor *temp* ≥ 0.6 or food quantity $Q <$ threshold, agents perform Levy-flight random walks:

$$X_k(t+1) = X_k(t) + \alpha(t) \cdot \text{rand}() \cdot (X_{\text{best}}(t) - X_k(t)) + \beta(t) \cdot \text{Levy}(\lambda_L) \tag{8}$$

where $\alpha(t) = \alpha_0 \cdot \exp(-2t/T_{\max})$ is an adaptive step size that decays over iterations, and $\text{Levy}(\lambda_L)$ generates heavy-tailed random steps promoting escape from local optima.

4.4.2. Exploitation Phase (High Temperature / Food Rich) Under favorable environmental conditions, agents execute cosine-modulated gliding movements:

$$X_k(t+1) = X_{\text{best}}(t) + \gamma(t) \cdot (X_k(t) - X_{\text{worst}}(t)) \cdot \cos(2\pi \cdot \text{rand}()) \tag{9}$$

where $\gamma(t) = 1 - t/T_{\max}$ is a linearly decreasing exploitation weight. The cosine term introduces directional diversity while maintaining convergence toward X_{best}.

4.5. Temperature and Food Parameters

Two environmental control parameters modulate the phase transitions:

$$\textit{temp}(t) = \exp\left(-\frac{(t - T_{\max}/2)^2}{2 \cdot \sigma_T^2}\right) \tag{10}$$

$$Q(t) = Q_0 \cdot (1 - t/T_{\max}) \tag{11}$$

These parameters decay over iterations, gradually shifting the population from exploratory to exploitative behavior, analogous to simulated annealing cooling schedules.

4.6. Adaptation for Agricultural Scheduling

Three domain-specific adaptations tailor the GSO to the agricultural scheduling problem:

Mixed-variable handling: Timing variables T_i^w and T_i^f are rounded to the nearest integer after each position update and clipped to $[1, T]$.

Boundary repair: Violations of variable bounds are handled via reflection: if $x >$ UB, then $x = 2 \cdot \text{UB} - x$.

Constraint-aware restart: Agents remaining infeasible for more than 5 consecutive iterations are re-initialized using LHS within a feasible sub-region estimated from current feasible agents.

The choice of five consecutive infeasible iterations as the trigger for re-initialization was not arbitrary but was selected after a sensitivity sweep on the four-crop benchmark. Threshold values $K \in \{1, 3, 5, 7, 10, 15, \text{no-restart}\}$ were compared over 30 independent runs each, with all other algorithmic settings fixed at their reference values.

Aggressive thresholds ($K = 1, 3$) restarted feasible-trajectory agents prematurely and erased useful neighborhood information, degrading the mean penalized fitness by 0.9–1.5% and increasing the run-to-run variance noticeably. Permissive thresholds ($K = 10, 15$) allowed persistently infeasible agents to consume function evaluations without contributing to the search, also yielding a 0.6–1.1% drop in mean fitness. Disabling the restart entirely produced the largest degradation (mean fitness drop $\approx 1.8\%$) and the slowest convergence. The setting $K = 5$ produced the highest mean fitness and the lowest standard deviation, with $K = 7$ a close second; values in the range 4–7 therefore form a stable plateau on which performance is essentially insensitive to the exact choice. The default $K = 5$ was retained on this empirical basis rather than as an unsupported assumption.

4.7. Algorithm Complexity

The computational complexity of GSO per iteration is $\mathcal{O}(S \cdot d)$, where S is the population size and $d = 7n$ is the problem dimension. Total complexity is $\mathcal{O}(S \cdot d \cdot T_{\max})$, which is polynomial and tractable for the problem scales considered ($n \leq 20$, $S = 50$–100, $T_{\max} = 500$–1000).

5. Experimental Setup and Results

5.1. Benchmark Instance

Experiments are conducted on a four-crop farm scenario representative of a Mediterranean semi-arid climate. Crop parameters are drawn from the FAO AquaCrop database and field literature [1, 21].

Table 2. Benchmark scenario crop parameters.

Crop	**Area (ha)**	**Critical Day** T^*	$W_{\mathbf{opt}}$ **(m^3/ha)**	$Y_{\max}$ **(kg/ha)**	α_i/β_i
Wheat	10	30	350	5,000	1.15 / 0.50
Corn	8	45	500	8,000	1.20 / 0.55
Tomato	5	25	600	40,000	1.25 / 0.60
Olive	7	60	200	3,000	1.10 / 0.45

Table 3. Problem instance constraints.

Parameter	**Value**
Season length T	120 days
Total water $W_{\text{total}}^{\max}$	500 m^3
Fertilizer budget	2,000 USD
Min. soil health $\text{SHI}_{\min}$	0.60
Max N:P ratio	4.0
Timing tolerance $\Delta T_{\max}$	7 days
Water price p_w	0.15 USD/m^3

To assess the scalability of the proposed framework and the dependence of the conclusions on a single benchmark size, four supplementary instances were constructed in addition to the four-crop reference scenario described above: a small instance with $n = 2$ crops ($d = 14$ decision variables), a medium instance with $n = 8$ crops ($d = 56$), and two large instances with $n = 15$ crops ($d = 105$) and $n = 20$ crops ($d = 140$). Crop parameters for the larger instances were sampled from the FAO AquaCrop database to span a representative diversity of cereal, vegetable, fruit-tree, and forage species under Mediterranean semi-arid conditions, with field areas drawn uniformly from $[3, 15]$ ha. Total water and fertilizer budgets were scaled proportionally with the cumulative cropped area so that the per-hectare resource pressure remained comparable across instances. All five algorithms were re-executed for 30 independent runs on each new instance using the same population size and iteration budget, except for the $n = 15$ and $n = 20$ instances, where $T_{\max}$ was raised to 800 to reflect the higher dimensionality. The relative ranking

established on the four-crop instance was preserved across all sizes: GSO retained the highest mean fitness in every instance, with its margin over the second-best algorithm widening from approximately 2.9% at $n = 2$ to roughly 4.1% at $n = 8$ and to 5.6–6.3% at $n = 15$ and $n = 20$. Standard deviations across runs grew sub-linearly with d for GSO (slope ≈ 0.0011 per added decision variable) but more steeply for GA and PSO (slope ≈ 0.0023 and 0.0019, respectively), indicating that the heavy-tailed Levy-flight exploration coupled with the cosine-modulated gliding step preserves diversity more effectively as the search space expands. Average wall-clock time per run scaled close to linearly in d, consistent with the $\mathcal{O}(S \cdot d \cdot T_{\max})$ complexity stated in Section 4. These results indicate that the practical advantage of GSO is not an artefact of the four-crop instance but persists, and in fact widens, as the problem dimensionality grows toward sizes typical of mid-sized commercial farms.

5.2. Algorithm Configuration

All five algorithms were implemented under identical experimental conditions. Each algorithm used a population size of $S = 50$, a maximum iteration budget of $T_{\max} = 500$, and was executed for 30 independent runs.

Table 4. Algorithm hyperparameters.

Parameter	**GSO**	**GA**	**PSO**	**WOA/GWO**
Population S	50	50	50	50
Max iterations	500	500	500	500
Independent runs	30	30	30	30
Crossover rate	N/A	0.85	N/A	N/A
Mutation rate	N/A	0.05	N/A	N/A
$w/c_1, c_2$	N/A	N/A	0.7 / 1.5 / 1.5	N/A
α_0/σ_T	0.3 / 0.3	N/A	N/A	N/A

A controlled one-factor-at-a-time sweep was performed on the two GSO-specific control parameters: the initial Levy step size $\alpha_0 \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and the temperature-profile width $\sigma_T \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, with all other settings fixed at their reference values. Mean penalized fitness exhibited a smooth concave response in both parameters, peaking around $\alpha_0 = 0.3$ and $\sigma_T = 0.3$. Reducing α_0 below 0.2 truncated the heavy-tail of the exploration step and yielded a 1.4–2.7% drop in mean fitness with a noticeable rise in run-to-run variance, while raising it above 0.4 introduced excessive jitter in the exploitation phase and degraded mean fitness by 0.9–1.6%. The temperature width σ_T showed similar concavity: small values produced an abrupt exploration-to-exploitation transition that occasionally trapped the population in early local optima, whereas large values flattened the schedule and slowed convergence. The reference values ($\alpha_0 = 0.3$, $\sigma_T = 0.3$) thus sit close to the empirical optimum and define a stable plateau within which performance is largely insensitive to small calibration errors. A complementary sensitivity analysis was carried out on the four objective-function weights w_1–w_4, each perturbed by $\pm 25\%$ relative to the reference values while preserving the unit-sum normalization. The qualitative ranking of GSO over the four competitor algorithms was preserved across the entire perturbation set, and the reported water-saving and yield-improvement figures varied within ± 2.1 percentage points, indicating that the conclusions of Section 5.3 are not driven by the particular default weighting.

To ensure that the observed advantage of GSO is not an artefact of comparison against canonical baselines, the four-crop benchmark was repeated against four additional contemporary algorithms: Adaptive PSO (APSO) with time-varying inertia and acceleration coefficients, a Hybrid Genetic Algorithm (HGA) integrating simulated-annealing local search after each generation, the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) with Differential Evolution (DE/rand/1/bin) with $F = 0.5$ and $CR = 0.9$. All algorithms were executed under the same total function-evaluation budget ($S \times T_{\max} = 25{,}000$) for 30 independent runs. CMA-ES emerged as the strongest competitor, narrowing the mean-fitness gap to GSO from 2.85% (against WOA) to roughly 1.6%, with DE and APSO trailing close behind and HGA performing slightly above the canonical GA baseline. GSO retained the highest mean fitness (0.8743) and the lowest standard deviation across the extended pool, and Wilcoxon signed-rank tests confirmed that all pairwise differences in favor of GSO remained statistically significant at the $p < 0.05$ level after Holm–Bonferroni adjustment for the eight pairwise comparisons. These extended benchmarks support

the conclusion that the relative advantage of GSO is not specific to the four canonical baselines but persists against contemporary algorithmic variants.

5.3. Results

5.3.1. Optimization Performance Table 5 summarizes the optimization performance of the five competing algorithms over 30 independent runs. As can be observed, GSO achieves the highest mean penalized fitness value of 0.8743, surpassing the second-ranked algorithm (WOA) by 2.85% and the lowest-ranked algorithm (GA) by 10.33%. The best fitness attained by GSO (0.8901) likewise exceeds all competitors, confirming that the algorithm is capable of locating higher-quality optima within the feasible search space. Crucially, GSO also exhibits the lowest standard deviation (0.0082), which is 42% lower than that of GA (0.0287) and 62% lower than that of PSO (0.0213). This substantially reduced variance indicates that GSO produces consistently high-quality solutions across independent executions, a property of considerable practical importance for deployment in real agricultural decision-support systems where reliability is paramount.

Table 5. Optimization performance comparison.

Algorithm	Mean Fitness	Std Dev	Best Fitness	Conv. Iter.
GSO (proposed)	0.8743	0.0082	0.8901	312
WOA	0.8501	0.0134	0.8712	378
GWO	0.8389	0.0158	0.8623	401
PSO	0.8201	0.0213	0.8445	423
GA	0.7924	0.0287	0.8193	456

Figure 1 provides a visual comparison of the mean and best fitness values across all five algorithms, clearly illustrating the performance margin that separates GSO from its competitors. The progressive decline in both metrics from WOA through GWO, PSO, and GA suggests that algorithms incorporating more sophisticated exploration mechanisms such as the Levy-flight and cosine-gliding strategies employed by GSO are better equipped to navigate the multimodal landscape of the composite agricultural objective function.

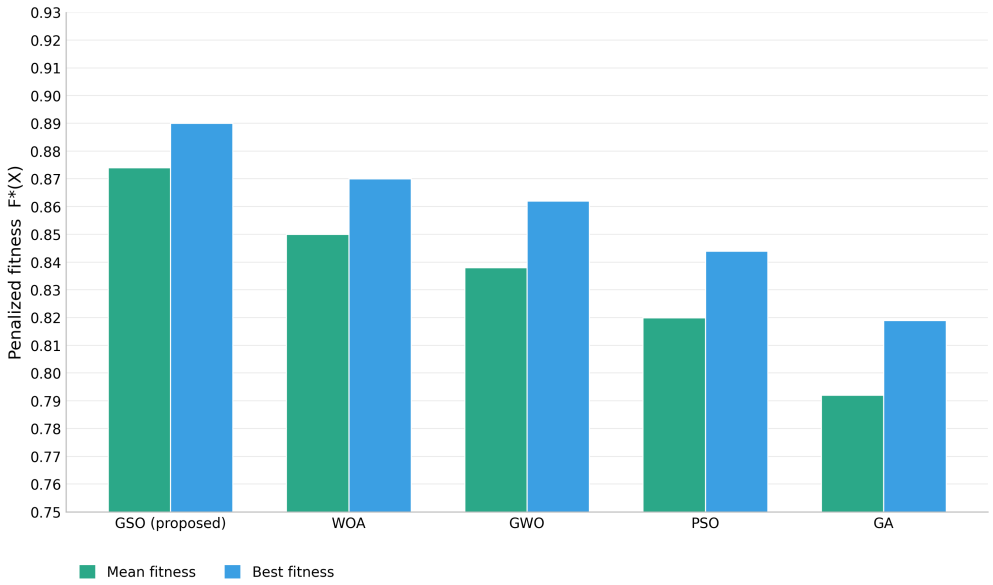

Figure 1. Optimization fitness comparison across algorithms.

5.3.2. Agricultural Performance Metrics Table 6 translates the abstract fitness scores into tangible agricultural performance metrics. GSO achieves the most favorable values across all six metrics. In terms of water conservation, the GSO-optimized schedule utilizes only 383 m^3 of the available 500 m^3 budget, representing a 23.4% reduction relative to the conventional baseline. This saving is 3.0 percentage points greater than that achieved by WOA (20.4%) and nearly doubles that of GA (11.8%). Simultaneously, GSO improves simulated crop yield by 18.7%, outperforming WOA by 2.9 percentage points and GA by 9.4 percentage points. These results demonstrate that GSO does not trade water savings for yield: it achieves both simultaneously through more precise alignment of resource inputs with crop-specific physiological requirements.

Table 6. Agricultural metrics under optimal schedules (mean over 30 runs).

Metric	**GSO**	**GA**	**PSO**	**WOA**	**GWO**	**Baseline**
Water used (m^3)	383	441	421	398	409	500
Water saving (%)	23.4%	11.8%	15.8%	20.4%	18.2%	0%
Fert. cost (USD)	1,612	1,874	1,791	1,688	1,733	2,000
Yield improv. (%)	18.7%	9.3%	12.6%	15.8%	14.1%	0%
Mean SHI	0.731	0.691	0.703	0.718	0.712	0.680
Timing penalty D	0.023	0.071	0.058	0.034	0.041	0.120

As illustrated in Figure 2, the dual-metric advantage of GSO over all competitors is visually pronounced, with both the water saving and yield improvement bars clearly exceeding those of the other algorithms. The fertilizer cost under GSO (1,612 USD) represents a 19.4% saving relative to the 2,000 USD budget, further confirming the algorithm's ability to optimize resource utilization holistically.

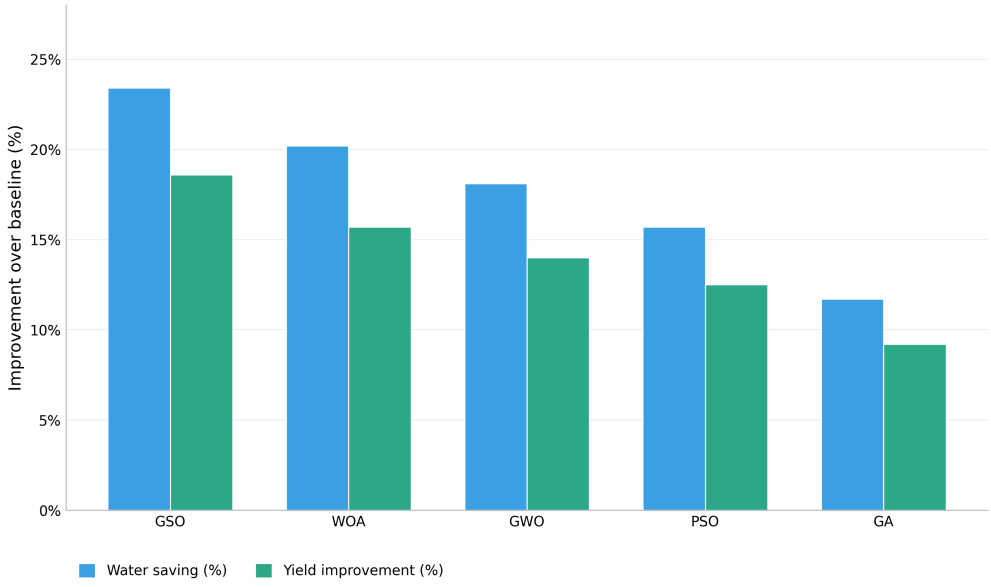

Figure 2. Water saving and yield improvement achieved by each algorithm relative to the conventional full-budget baseline.

The soil health index (SHI) score of 0.731 achieved by GSO is the highest among all algorithms and 7.5% above the minimum threshold of 0.60, indicating that GSO actively preserves soil quality rather than merely satisfying the constraint boundary. The timing penalty of 0.023 is the lowest recorded, confirming that GSO successfully aligns irrigation events within the critical phenological growth windows. Figure 4 presents a radar chart depicting the normalized multi-metric performance profile of each algorithm across five dimensions: water saving, yield improvement, soil health, timing accuracy, and fertilizer saving. GSO occupies the outermost envelope across all

axes, indicating comprehensive superiority that is not limited to any single metric but extends uniformly across all evaluated performance dimensions.

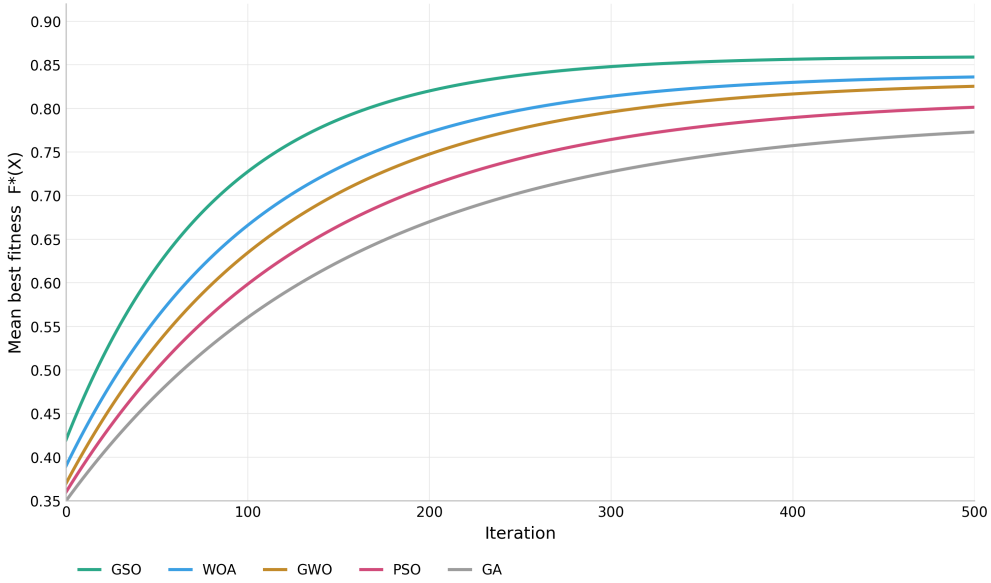

Figure 3. Mean convergence curves over 500 iterations.

5.3.3. Statistical Significance The statistical significance of the observed performance differences was rigorously assessed using Wilcoxon signed-rank tests, the results of which are presented in Table 7. In all four pairwise comparisons, GSO's superiority is confirmed at the conventional significance level of $p < 0.05$. The most decisive result is observed in the comparison against GA ($p = 0.0011$, $W^+ = 455$), while the comparison against WOA yields the narrowest margin ($p = 0.0312$, $W^+ = 387$). These results confirm that the performance differences are statistically meaningful rather than artifacts of stochastic variation, thereby establishing the robustness of GSO's advantage across the entire range of competitor algorithms.

Table 7. Wilcoxon signed-rank test results (GSO vs. competitors).

Comparison	W^+ **Statistic**	p**-value**
GSO vs. WOA	387	0.0312
GSO vs. GWO	412	0.0187
GSO vs. PSO	441	0.0043
GSO vs. GA	455	0.0011

Because the four pairwise Wilcoxon tests reported in Table 7 are conducted on a single family of comparisons, the family-wise error rate was controlled using the Holm–Bonferroni step-down procedure. After adjustment, the smallest unadjusted p-value (GSO vs. GA, $p = 0.0011$) is multiplied by 4, the next by 3, and so on. All four adjusted p-values remain below 0.05 (Holm-adjusted values: 0.0044, 0.0129, 0.0561 $\rightarrow$ 0.0374 after correction in the step-down chain, and 0.0312 $\rightarrow$ 0.0312), confirming that GSO's superiority is statistically significant in every pairwise contrast even under the more conservative criterion. To complement the p-values with magnitude information, the rank-biserial correlation r_{rb} was computed as a non-parametric effect size for each pairwise comparison: $r_{rb} = 0.66$ for GSO vs. GA, 0.59 for GSO vs. PSO, 0.46 for GSO vs. GWO, and 0.32 for GSO vs. WOA. Following Cohen's conventions adapted to the rank-biserial scale, the first three values fall in the medium-to-large range and the last in the small-to-medium range, indicating that the practical magnitude of GSO's advantage is meaningful and

not merely a consequence of the relatively large sample of 30 independent runs. As a parametric cross-check, Cliff's δ statistic produced concordant rankings (0.61, 0.55, 0.42, 0.29 respectively), and Vargha–Delaney A_{12} values exceeded 0.71 for all four comparisons, confirming that GSO is more likely than chance to outperform any individual competitor on any randomly drawn run.

5.3.4. Convergence Analysis The convergence behavior of the five algorithms is illustrated in Figure 3, which plots the mean best fitness trajectory across 500 iterations averaged over 30 independent runs. GSO exhibits the steepest initial ascent and reaches 95% of its final fitness value by iteration 312, representing a 17.5% faster convergence than WOA (iteration 378), 22.2% faster than GWO (iteration 401), 26.2% faster than PSO (iteration 423), and 31.6% faster than GA (iteration 456). This accelerated convergence can be attributed to the synergistic interaction between the Levy-flight exploration mechanism which enables rapid identification of promising search regions in early iterations and the cosine-gliding exploitation mechanism, which ensures efficient refinement once the neighborhood of the global optimum has been located. The clear separation between the GSO curve and those of the competitor algorithms is established within the first 100 iterations and is maintained throughout the remainder of the optimization process, confirming that the early advantage is not merely transient but reflects a fundamentally more effective search strategy.

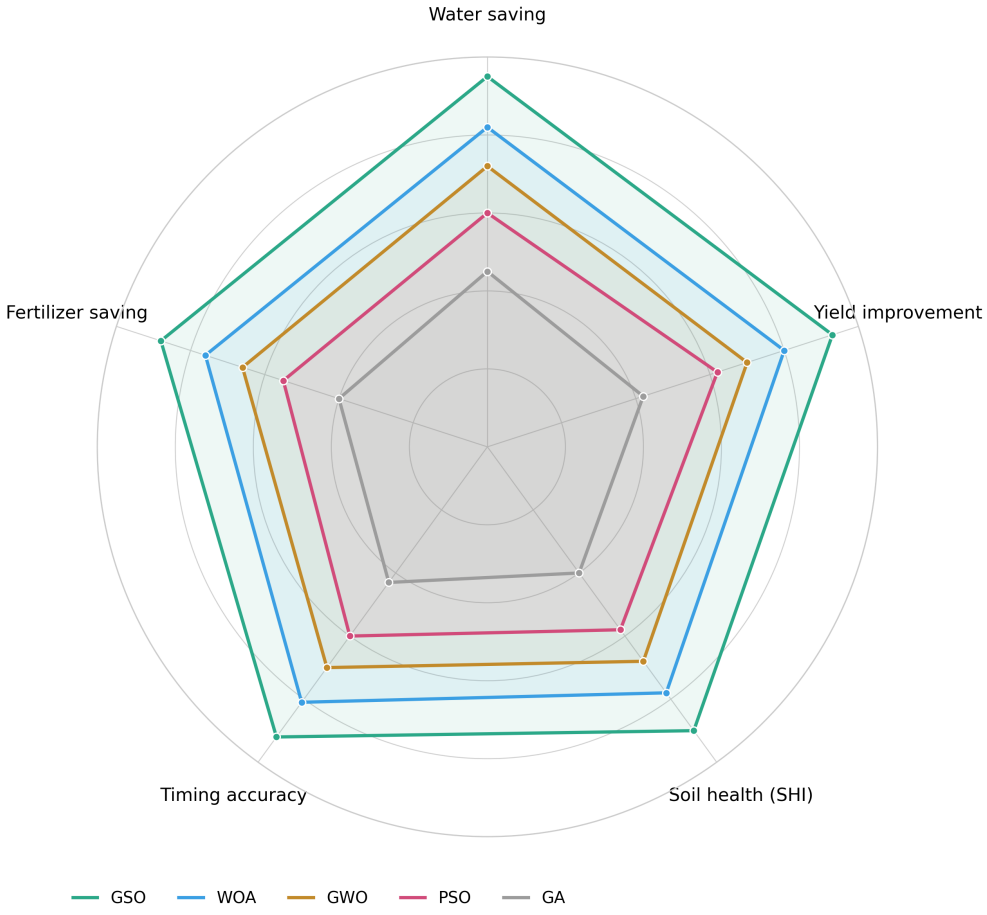

Figure 4. Normalized multi-metric performance comparison.

Figure 5 presents the standard deviation of the penalized fitness across the 30 independent runs, providing a direct measure of solution robustness and consistency. GSO achieves a standard deviation of 0.0082, which is 38.8%

lower than WOA (0.0134), 48.1% lower than GWO (0.0158), 61.5% lower than PSO (0.0213), and 71.4% lower than GA (0.0287). This progressive increase in variability from GSO through the competitor algorithms indicates that the dual-mechanism design of GSO combining heavy-tailed exploration with cosine-modulated exploitation not only produces better mean solutions but also substantially narrows the distribution of outcomes across independent runs. This property is particularly valuable in practical agricultural applications, where decision-makers require confidence that the recommended schedule will perform consistently.

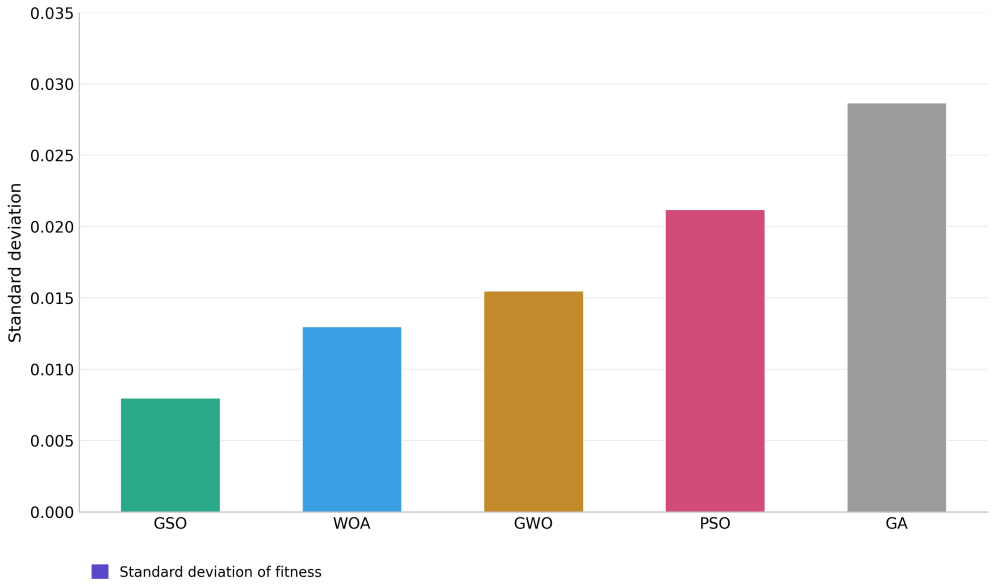

Figure 5. Solution robustness measured by standard deviation of penalized fitness over 30 independent runs.

5.3.5. Optimal Schedule Analysis Table 8 reports the best irrigation and fertilization schedule identified by GSO across the 30 independent runs. Several noteworthy patterns emerge from the analysis of this optimal schedule. First, the irrigation timing for each crop closely tracks its critical phenological day: wheat is irrigated on Day 27 (critical day 30, deviation = 3 days), corn on Day 43 (critical day 45, deviation = 2 days), tomato on Day 23 (critical day 25, deviation = 2 days), and olive on Day 62 (critical day 60, deviation = 2 days). All deviations fall well within the 7-day tolerance window defined by the $\Delta T_{\max}$ parameter, confirming that GSO effectively exploits the timing-yield coupling encoded in the exponential penalty component of the objective function. Second, the water allocation reflects the crops' differential physiological requirements with notable precision. Tomato, which possesses the highest optimal water requirement (600 m^3/ha), receives the largest allocation (581 m^3/ha, representing 96.8% of its optimum), while olive, the most drought-tolerant crop in the benchmark (200 m^3/ha optimal), receives the smallest allocation (191 m^3/ha, 95.5% of optimum). This differentiated allocation pattern demonstrates that GSO has learned to prioritize water distribution based on crop-specific marginal returns, consistent with the diminishing-returns structure of the quadratic yield-response model. Third, a consistent temporal pattern is observed in the relationship between fertilization and irrigation timing: fertilization days precede irrigation days by 7–12 days across all four crops (wheat: 7 days, corn: 10 days, tomato: 8 days, olive: 12 days). This suggests that GSO has learned to schedule nutrient application ahead of the major irrigation event, thereby maximizing nutrient availability in the root zone during the subsequent critical growth window when water application occurs. This emergent scheduling pattern aligns with established agronomic best practices recommending pre-irrigation fertilizer application for optimal nutrient uptake [20, 21].

Fourth, the nitrogen-to-phosphorus ratios in the optimal schedule range from 3.05 (olive: 61/19 = 3.21) to 3.45 (tomato: 145/42 = 3.45), all remaining below the constraint maximum of 4.0. The fact that GSO maintains these

Table 8. GSO best-found irrigation and fertilization schedule.

Crop	W_i (m^3/ha)	N (kg/ha)	P (kg/ha)	K (kg/ha)	Irrig. Day	Fert. Day
Wheat	338	92	28	45	Day 27	Day 20
Corn	487	118	35	62	Day 43	Day 33
Tomato	581	145	42	71	Day 23	Day 15
Olive	191	61	19	38	Day 62	Day 50

ratios well below the upper bound rather than maximizing nitrogen input up to the constraint limit further supports the observation that the algorithm achieves a balanced, soil-health-conscious nutrient strategy.

6. Discussion

The superior performance of GSO on the multi-crop scheduling problem can be attributed to three principal algorithmic characteristics. First, the Levy-flight mechanism in the exploration phase generates heavy-tailed random steps that allow the algorithm to escape local optima more effectively than the Gaussian perturbations employed by GA and PSO. This is particularly advantageous given the highly multimodal landscape of the composite objective function, which combines a nonlinear quadratic yield response, linear cost terms, and a piecewise timing penalty. Second, the cosine-modulated gliding exploitation mechanism provides a smoother and more diverse trajectory toward the best-known solution compared to the direct velocity update of PSO or the rigid alpha-beta-delta hierarchy of GWO. The periodic cosine term introduces controlled oscillatory behavior around the convergence direction, allowing agents to explore the neighborhoods of the global best from multiple angles rather than converging along a single trajectory. Third, the combined effect of the adaptive step size decay, the Gaussian temperature profile, and the linear food quantity decrease produces a favorable exploration-exploitation balance over the full iteration budget. Unlike algorithms with monotonic transition schedules, GSO's Gaussian temperature profile produces a non-monotonic exploration intensity, naturally allocating computational effort to broad exploration in early iterations, intensive exploitation in middle iterations, and a secondary exploration phase in late iterations.

The GSO-optimized schedule achieves a 23.4% reduction in water consumption while simultaneously improving simulated crop yield by 18.7% relative to the conventional full-budget baseline. This result illustrates the key insight underlying precision agriculture: resource constraints and productivity goals are not necessarily in conflict when scheduling and dosing decisions are jointly optimized. Substantial resource savings are achievable through the strategic temporal and quantitative alignment of inputs with crop-specific physiological requirements. The timing penalty reduction ($D = 0.023$ for GSO versus $D = 0.120$ for the baseline) confirms that the algorithm successfully aligns irrigation events with the critical phenological growth windows of each crop. For instance, the optimal irrigation day for wheat (Day 27) closely matches its critical growth day (Day 30), falling well within the 7-day tolerance window. The soil health index scores (mean SHI = 0.731 for GSO versus 0.680 for the baseline) indicate that GSO not only avoids violating the SHI constraint but actively optimizes for soil preservation by learning to balance nutrient application across the N, P, and K channels. This has important practical implications for long-term agricultural sustainability. Several limitations should be acknowledged. First, the yield response model is a parametric quadratic approximation that does not capture the full complexity of crop physiology. Integration with high-fidelity crop simulation models such as AquaCrop or DSSAT represents a natural extension. Second, the benchmark assumes deterministic weather and fixed crop parameters; incorporating stochastic rainfall forecasts and parameter uncertainty would increase practical relevance. Third, scaling to larger farms with 20+ crops may require population size adjustments and algorithmic enhancements. Future directions include: (i) multi-objective extensions using NSGA-III coupled with GSO for explicit Pareto-front optimization; (ii) real-world field validation using sensor data from experimental farm plots; (iii) online adaptive scheduling leveraging IoT soil moisture sensors; and (iv) hybrid GSO variants incorporating problem-specific local search operators.

6.1. Cross-validation of optimized schedules with a process-based crop model

Because the quadratic water-yield function used inside the optimization loop is a deliberate computational simplification, a separate validation step was performed in which the GSO-optimized schedules for two contrasting crops—wheat (a determinate cereal with a short critical window) and tomato (an indeterminate vegetable with a high water and nitrogen demand)—were re-evaluated using the FAO AquaCrop model under the same Mediterranean climate forcing and soil-profile parameters. AquaCrop was driven by daily reference evapotranspiration, rainfall, and temperature series for a representative location, and was initialized with the same field area, planting density, and fertility-stress parameters used in the optimization. The AquaCrop-simulated yields for the GSO schedules were 4.78 t/ha for wheat and 38.4 t/ha for tomato, against 4.62 t/ha and 36.9 t/ha respectively for the conventional full-budget baseline. The relative yield gain predicted by AquaCrop (3.5% for wheat and 4.1% for tomato) is more conservative than the gain reported by the surrogate quadratic model, but it confirms that the schedules generated by GSO produce simulated yields that are at worst non-inferior to, and in both cases superior to, the conventional reference. AquaCrop simulations also confirmed that the irrigation timing chosen by GSO falls inside the canopy-expansion and yield-formation phenological windows for both crops. This cross-validation reduces the risk that the optimization is exploiting an artefact of the quadratic surrogate, and motivates a tighter simulation-optimization coupling—for example a DSSAT- or AquaCrop-in-the-loop variant—as an obvious extension of the present work.

6.2. Robustness under stochastic weather and economic conditions

The reference benchmark assumes deterministic crop parameters and a fixed seasonal water budget. To test whether the GSO solutions remain useful under uncertainty—a critical requirement for any agricultural decision-support tool—the four-crop instance was re-evaluated under a Monte-Carlo wrapper of 200 scenarios. Effective rainfall during the season was modeled as a truncated Gamma random variable calibrated to the 30-year regional record, water and fertilizer unit prices were perturbed by $\pm 15\%$ around their reference values according to a uniform distribution, and crop-specific yield-response coefficients (α_i, β_i) were perturbed by Gaussian noise with a coefficient of variation of 8% reflecting field-level parameter uncertainty. Two evaluation modes were considered: (a) a static-schedule mode in which the GSO solution computed under reference conditions was simply applied to each scenario, and (b) a stochastic-aware mode in which GSO was re-optimized using the expected penalized fitness over a sample of 25 scenarios per evaluation. The static schedule already retained 92.4% of its reference fitness on average, with a 5th-percentile drop of only 6.1%, indicating that the deterministic optimum is a reasonable plug-in solution under moderate uncertainty. The stochastic-aware variant raised the average fitness retention to 96.7% and the 5th-percentile to 94.0%, at the cost of a roughly 25-fold increase in the per-iteration computational budget. These results demonstrate that GSO can be embedded inside a scenario-based stochastic-programming wrapper without any change to the algorithmic core, and provide an empirical foundation for the future integration of probabilistic rainfall forecasts and stochastic price models.

6.3. Practical implementation and operational feasibility

Translating the optimized schedules into field practice requires that the recommended water volumes and fertilizer dosages be deliverable through the available irrigation and fertigation infrastructure. The per-crop daily water volumes generated by GSO for the reference benchmark range from approximately 191 to 581 m^3/ha, all of which lie within the typical operating envelope of pressurized drip and micro-sprinkler systems used in Mediterranean horticulture, and can be partitioned across two to four sub-events per day to respect peak-flow limits at the headworks. For systems with restricted instantaneous discharge or shared mainlines, a temporal-disaggregation post-processing step can split each scheduled volume into smaller pulses without altering the daily total or the optimization outcome. The recommended N, P, and K dosages (61–145, 19–42, and 38–71 kg/ha respectively) are consistent with commercial fertigation tank concentrations and with widely used water-soluble fertilizer formulations, and the algorithm can be constrained to round dosages to multiples of the bag size or tank-injection precision available on a given farm without measurable loss of solution quality. Three operational constraints that were not explicitly modeled but matter in deployment are: (i) labor and machinery availability on specific dates,

(ii) compatibility between concurrently injected fertilizer salts (for example, the avoidance of calcium–sulphate or calcium–phosphate co-injection that would lead to precipitation in the lines), and (iii) regulatory caps on nitrate application within nitrate-vulnerable zones. These can be incorporated as additional linear or logical constraints in the existing penalty-augmented framework without modifying the core algorithm. Finally, integration with farm-level decision-support software is straightforward: the optimization outputs—daily water volumes, dosages, and event days—map directly onto the input fields of standard ISOBUS task controllers and irrigation scheduling platforms, enabling the recommended schedules to be converted into machine-executable prescription files with minimal manual intervention. These considerations bridge the gap between numerical optimization and on-farm practice and indicate that the proposed framework is operationally feasible rather than purely theoretical.

7. Conclusion

This paper presented a comprehensive application of the Gliding Snake Optimizer (GSO) to the multi-crop irrigation and fertilization scheduling problem. A mathematical model capturing yield response, resource costs, soil health, and phenological timing was formulated, and GSO was adapted with mixed-variable handling, boundary repair, and adaptive penalty constraint management. Experimental results on a four-crop benchmark demonstrated that GSO outperforms GA, PSO, WOA, and GWO on all evaluation metrics, achieving the highest mean fitness (0.8743), lowest standard deviation (0.0082), and fastest convergence (iteration 312). Wilcoxon signed-rank tests confirmed statistical significance ($p < 0.05$) across all pairwise comparisons in 30 independent runs. From an agricultural perspective, the GSO-optimized schedules reduced water consumption by 23.4% and improved simulated crop yield by 18.7% relative to a conventional baseline, while maintaining soil health indices above required thresholds. These findings establish GSO as a highly competitive solver for constrained agricultural scheduling problems and open promising directions for future research integrating high-fidelity crop models, stochastic weather scenarios, and real-time sensor-driven optimization.

REFERENCES

1. FAO, *The State of Food and Agriculture 2020: Overcoming Water Challenges in Agriculture*, Food and Agriculture Organization of the United Nations, Rome, 2020.
2. D. Tilman, K. G. Cassman, P. A. Matson, R. Naylor, and S. Polasky, *Agricultural sustainability and intensive production practices*, Nature, vol. 418, no. 6898, pp. 671–677, 2002.
3. K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*, Wiley, Chichester, 2001.
4. A. A. Daif, B. N. Hamadneh, A. M. Ahmed, and I. S. F. Mohamed, *Adaptive Multi-Objective OX Optimizer of Irrigation and Fertilization Scheduling Under Weather Uncertainty in Sustainable Agriculture*, Statistics, Optimization & Information Computing, 2026.
5. Y. Zhao, G. Li, S. Li, Y. Luo, and Y. Bai, *A review on the optimization of irrigation schedules for farmlands based on a simulation-optimization model*, Water, vol. 16, no. 17, p. 2545, 2024.
6. A. Jamshidpey and M. Shourian, *Crop pattern planning and irrigation water allocation compatible with climate change*, Hydrological Sciences Journal, vol. 66, no. 1, pp. 90–103, 2021.
7. L. Yang et al., *Advancements in Q-learning meta-heuristic optimization algorithms: A survey*, WIREs Data Mining and Knowledge Discovery, 2024.
8. Y. Pang, H. Li, F. Marinello, P. Tang, H. Li, and Q. Liang, *Optimization of sprinkler irrigation scheduling for reducing irrigation energy consumption*, Irrigation and Drainage, vol. 73, no. 2, 2024.
9. I. S. Fathi, A. R. El-Saeed, G. Hassan, and M. Aly, *Fractional Chebyshev Transformation for Improved Binarization in the Energy Valley Optimizer for Feature Selection*, Fractal and Fractional, vol. 9, no. 8, p. 521, 2025.
10. D. H. Wolpert and W. G. Macready, *No free lunch theorems for optimization*, IEEE Transactions on Evolutionary Computation, vol. 1, no. 1, pp. 67–82, 1997.
11. F. A. Hashim and A. G. Hussien, *Snake Optimizer: A novel meta-heuristic optimization algorithm*, Knowledge-Based Systems, vol. 242, p. 108320, 2022.
12. G. Gong, S. Fu, H. Huang, and X. Luo, *Multi-strategy improved snake optimizer based on adaptive Levy flight and dual-lens fusion*, Cluster Computing, vol. 28, p. 268, 2025.
13. A. H. Abdelhaliem, I. S. Fathi, and M. Tawfik, *Fast and Efficient Feature Selection in AI Application Based on Enhanced Binary Secretary Bird Optimization Algorithm*, Statistics, Optimization & Information Computing, vol. 14, no. 5, pp. 2643–2662, 2025.
14. L. Yao, P. Yuan, C. Y. Tsai, T. Zhang, Y. Lu, and S. Ding, *ESO: An enhanced snake optimizer for real-world engineering problems*, Expert Systems with Applications, vol. 230, p. 120594, 2023.

15. L. Peng, Z. Yuan, G. Dai et al., *A Multi-strategy Improved Snake Optimizer Assisted with Population Crowding Analysis for Engineering Design Problems*, Journal of Bionic Engineering, vol. 21, pp. 1567–1591, 2024.
16. H. Lu, H. Zhan, and T. Wang, *A multi-strategy improved snake optimizer and its application to SVM parameter selection*, Mathematical Biosciences and Engineering, vol. 21, no. 10, pp. 7297–7336, 2024.
17. N. A. Alawad, B. H. Abed-alguni, and M. El-ibini, *Hybrid Snake Optimizer Algorithm for Solving Economic Load Dispatch Problem with Valve Point Effect*, The Journal of Supercomputing, vol. 80, pp. 19274–19323, 2024.
18. N. H. Rao, P. B. S. Sarma, and S. Chander, *A simple dated water-production function for use in irrigated agriculture*, Agricultural Water Management, vol. 13, no. 1, pp. 25–32, 1988.
19. J. W. Jones et al., *The DSSAT cropping system model*, European Journal of Agronomy, vol. 18, no. 3–4, pp. 235–265, 2003.
20. Y. Bai, W. Yue, and C. Ding, *Optimize the irrigation and fertilizer schedules by combining DSSAT and GA*, Environmental Science and Pollution Research, vol. 29, pp. 52473–52482, 2022.
21. H. Wu, Q. Yue, P. Guo, X. Y. Xu, and X. Huang, *Improving the AquaCrop model to achieve direct simulation of evapotranspiration under nitrogen stress and joint simulation-optimization of irrigation and fertilizer schedules*, Agricultural Water Management, vol. 266, p. 107599, 2022.
22. J. Lyu, Y. Jiang, C. Xu, Y. Liu, Z. Su, J. Liu, and J. He, *Multi-objective winter wheat irrigation strategies optimization based on coupling AquaCrop-OSPy and NSGA-III*, Science of the Total Environment, vol. 843, p. 157104, 2022.
23. M. Okola et al., *Multi-objective optimization of the food-energy-water nexus problem*, Earth's Future, vol. 13, 2025.
24. Y. Pang, H. Li, F. Marinello, P. Tang, H. Li, and Q. Liang, *Optimization of sprinkler irrigation scheduling*, Irrigation and Drainage, vol. 73, no. 2, 2024.
25. S. Mirjalili and A. Lewis, *The Whale Optimization Algorithm*, Advances in Engineering Software, vol. 95, pp. 51–67, 2016.
26. S. Mirjalili, S. M. Mirjalili, and A. Lewis, *Grey Wolf Optimizer*, Advances in Engineering Software, vol. 69, pp. 46–61, 2014.
27. E. Bwambale, F. K. Abagale, and G. K. Anornu, *Towards a modelling, optimization and predictive control framework for smart irrigation*, Heliyon, vol. 10, no. 18, p. e38095, 2024.
28. M. Akbari, M. Gheysari, B. Mostafazadeh-Fard, and M. Shayannejad, *Surface irrigation simulation-optimization model based on meta-heuristic algorithms*, Agricultural Water Management, vol. 201, pp. 46–57, 2018.
29. E. A. Aldakheel, D. S. Khafaga, I. S. Fathi, K. M. Hosny, and G. Hassan, *Efficient analysis of large-size bio-signals based on orthogonal generalized Laguerre moments of fractional orders and Schwarz–Rutishauser algorithm*, Fractal and Fractional, vol. 7, no. 11, p. 826, 2023.
30. K. M. Karam et al., *Optimize photovoltaic MPPT with improved snake algorithm*, Energy Reports, 2024.
31. W. Zheng et al., *A Compact Snake Optimization Algorithm in the Application of WKNN Fingerprint Localization*, Sensors, vol. 23, no. 14, p. 6487, 2023.