TGWO-SA: A Territorial Grey Wolf Optimizer Enhanced by Simulated Annealing, Perturbation Dynamics, and Natural Selection

Fabio Berberi[1] and Paolo Mercorelli[2]

[1] [1] University of Siena, Via Roma 56, 53100 Siena, Italy
ORCID: 0009-0004-8825-8707
f.berberi@student.unisi.it
[2] [2] Leuphana University of Lüneburg, Universitätsallee 1, 21335 Lüneburg, Germany
ORCID: 0000-0003-3288-5280
paolo.mercorelli@leuphana.de

Abstract. This study develops a hybrid metaheuristic named **TGWO-SA**, short for Territorial Grey Wolf Optimizer with Simulated Annealing. It extends the classical Grey Wolf Optimizer (GWO) by combining several biologically motivated strategies: a self-adjusting population size that reflects territorial behavior, a cooperative communication phase driven by howling cues, and a selective renewal process that replaces weaker agents during the search. A local refinement stage based on Simulated Annealing introduces small stochastic perturbations around the leading wolf, imitating the corrective maneuvers observed in real predators during pursuit. Through this mechanism, TGWO-SA achieves a more stable balance between broad exploration and precise local improvement. Experiments on benchmark functions indicate faster convergence, higher consistency across runs, and improved accuracy compared with baseline algorithms such as GWO, PSO, and DE. The results confirm that the proposed approach maintains population diversity and adapts effectively to complex, multimodal landscapes.

Keywords: Grey Wolf Optimizer · Simulated Annealing · Hybrid Optimization · Perturbation · Natural Selection

1 Introduction

Nature-inspired optimization algorithms translate mechanisms found in biological or ecological systems into mathematical strategies capable of exploring complex search spaces. By mimicking cooperative or adaptive behaviors observed in nature, these algorithms have become highly versatile tools for handling nonlinear and high-dimensional problems where gradient-based techniques fail [4,7]. They have been successfully applied in diverse engineering domains, from control and robotics to data-driven machine learning tasks [8,5].

Among the large family of swarm-based metaheuristics, the *Grey Wolf Optimizer* (GWO) [9,4] remains one of the simplest and most intuitive. Its structure reproduces the hierarchical coordination of wolf packs, where the leading wolves

(α, β, and δ) direct the rest of the group (ω) in a collective hunting process. This behavioral metaphor allows candidate solutions to gradually converge toward promising regions of the search space through cooperative encirclement and pursuit. Despite its efficiency and conceptual elegance, standard GWO is prone to premature convergence and tends to lose diversity in late iterations, particularly when dealing with multimodal or high-dimensional landscapes [11,3].

To mitigate these limitations, this work proposes the **Territorial Grey Wolf Optimizer with Simulated Annealing (TGWO-SA)**, a biologically inspired hybrid algorithm that extends the original GWO framework. TGWO-SA introduces four complementary mechanisms designed to preserve diversity and improve precision: (1) an adaptive pack-sizing strategy that links population size to problem dimensionality (territorial adaptation); (2) a communication phase driven by exploration wolves that broadcast promising regions via howling signals; (3) a selective renewal process that periodically removes low-performing individuals; and (4) a local exploitation phase based on Simulated Annealing, which applies stochastic perturbations around the best candidate to refine its position [6,2,10,11].

A distinctive feature of TGWO-SA lies in the perturbation mechanism applied to the leading wolf (α). This process emulates the small corrective motions observed in natural predators during pursuit, enabling the optimizer to fine-tune solutions while avoiding premature trapping in local minima [1]. The overall interaction between exploration, communication, and local adaptation is illustrated in Fig. 1, which summarizes the main ecological dynamics embedded in the algorithm.

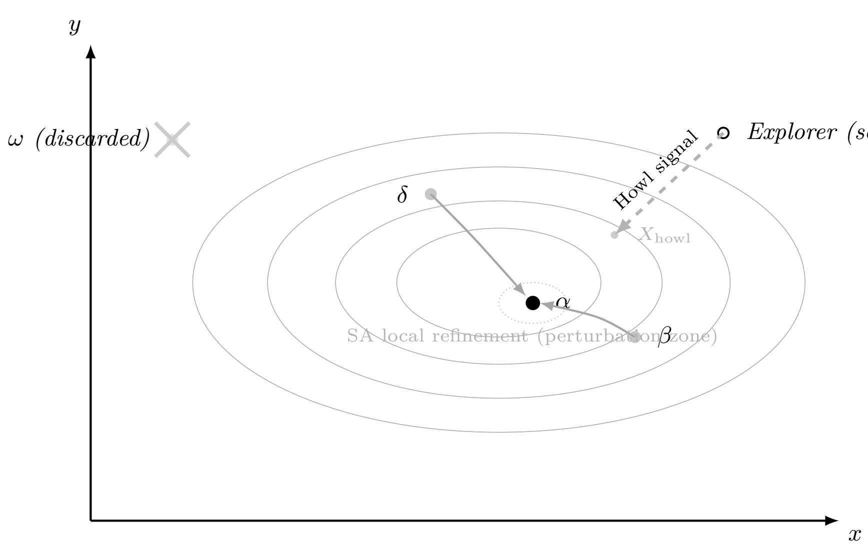

Fig. 1. Conceptual visualization of TGWO-SA showing pack hierarchy, exploratory scouts, howl communication, and local SA perturbation refinement.

2 Related Work

Since its introduction, the Grey Wolf Optimizer has inspired a wide range of hybrid and improved variants. Several researchers combined it with complementary strategies to overcome its tendency toward early stagnation. For instance, the Differential Grey Wolf Optimizer (DGWO) incorporates differential evolution operators to reinforce exploration, while hybrid schemes such as PSO-GWO or GA-GWO employ swarm and genetic mechanisms to accelerate convergence. Other studies introduced chaotic mappings or nonlinear control factors to sustain diversity throughout the search process. Although these extensions often lead to better numerical results, many of them remain purely mathematical and overlook the ecological aspects that originally motivated GWO, such as inter-wolf communication, competition, and population renewal.

Simulated Annealing (SA), first proposed by Kirkpatrick et al. in 1983, has also played a central role in local refinement within hybrid metaheuristics. Its probabilistic acceptance rule allows the algorithm to tolerate occasional uphill transitions, helping it escape local minima and maintain a degree of randomness near promising regions. Nevertheless, in most existing hybrids, SA is introduced as a secondary component—executed after the global phase rather than interacting directly with it—thus limiting its influence on population dynamics.

The proposed **TGWO-SA** aims to bridge this gap by coupling SA and GWO through a biologically motivated perturbation process applied to the leading wolf (α). This coupling transforms SA from a passive fine-tuning step into an active behavioral mechanism: stochastic yet bounded positional adjustments imitate micro-movements and sensory feedback observed in real wolf packs. Through this interaction, TGWO-SA achieves a more natural balance between large-scale exploration and adaptive local exploitation, creating a coherent ecological framework rather than a simple mathematical hybrid.

3 Proposed Method: TGWO-SA

The proposed TGWO-SA algorithm expands the original GWO framework through four biologically inspired modules that capture distinct aspects of wolf-pack behavior: (1) territorial adaptation, (2) exploration and howling-based communication, (3) natural selection and population renewal, and (4) a local refinement process driven by Simulated Annealing and stochastic perturbations. Together, these mechanisms promote a balanced interaction between global exploration and fine local exploitation. An overview of these interacting processes is illustrated in Fig. 1.

3.1 Territorial Adaptation

In nature, the size of a wolf pack is strongly influenced by the extent of its hunting territory and the availability of prey. TGWO-SA incorporates this idea by scaling the population size N according to the dimensionality D of the optimization problem:

$$N = k \times D, \quad k \in [5, 10]. \tag{1}$$

This adaptive rule ensures that the algorithm retains sufficient exploratory capacity in complex, high-dimensional problems while avoiding redundancy in simpler ones.

3.2 Exploration Wolves and Howling Communication

A small proportion of the wolves (typically 10–20%) are designated as *explorers*, responsible for scouting unexplored regions of the search space. When one of these explorers detects a region that significantly improves the fitness value,

$$f_{new} < f_\alpha - \Delta_{howl}, \tag{2}$$

it emits a **howl signal**. This signal encourages other wolves to partially reorient their movement toward the discovered position:

$$X_{wolf}(t+1) = X_{wolf}(t) + w \cdot (X_{howl} - X_{wolf}(t)), \tag{3}$$

where $w \in [0, 1]$ determines the strength of attraction. This mechanism reflects real-world communication and cooperative coordination within wolf packs, enhancing information flow and preventing stagnation around local optima.

3.3 Natural Selection and Pack Renewal

To sustain evolutionary pressure and prevent loss of diversity, wolves with fitness worse than 1.2 times the population mean are periodically removed. Their positions are replaced by random individuals or by mutated offspring of the best-performing wolves:

$$X_{new} = X_\alpha + \eta \cdot \mathcal{N}(0, 1), \tag{4}$$

where η controls the mutation amplitude. This renewal mechanism simulates ecological turnover, helping the population adapt dynamically to changing search conditions and preventing over-exploitation of limited regions.

3.4 Simulated Annealing and Perturbation Dynamics

The final layer of TGWO-SA introduces a local exploitation mechanism inspired by **Simulated Annealing (SA)**. At each iteration, the best solution found so far (α) is perturbed with Gaussian noise:

$$X' = X_\alpha + \mathcal{N}(0, \sigma), \tag{5}$$

where σ defines the radius of local exploration. The perturbed candidate X' is accepted if it improves the objective value ($\Delta f < 0$) or with a probability determined by the Boltzmann criterion:

$$P = e^{-\Delta f / T}, \tag{6}$$

where T represents the current temperature, progressively decreased according to $T_{t+1} = \lambda T_t$ with $\lambda \in [0.90, 0.99]$. This mechanism allows TGWO-SA to occasionally accept less favorable moves, preserving its ability to escape local minima while still refining solutions near high-quality regions.

A conceptual view of this perturbation process is provided in Fig. 2, which shows how multiple candidates X' are generated around the leader α, with only one being accepted—either through direct improvement or probabilistic selection—updating the leader's position to α'.

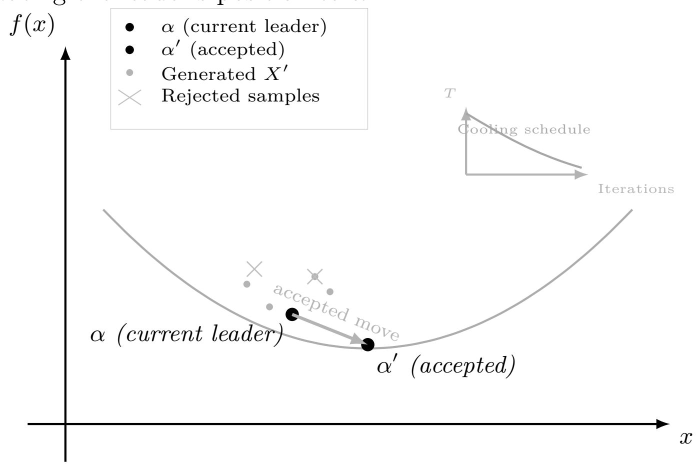

Fig. 2. Perturbation dynamics around the alpha wolf (α) during the Simulated Annealing refinement. Local perturbations X' are generated around α, and only the accepted move—based on fitness improvement or Boltzmann probability—updates the leader to α'. The inset shows the cooling schedule that controls probabilistic acceptance.

The perturbation mechanism serves a dual, biologically grounded function in TGWO-SA. On one hand, it introduces fine-grained adaptive variability that strengthens local exploitation. By continuously generating micro-adjustments around the leader α, the optimizer emulates the subtle sensory corrections performed by real wolves when tracking moving prey. These small stochastic deviations enable smoother, quasi-continuous exploration in high-quality regions, increasing precision without sacrificing adaptability. On the other hand, the Boltzmann-based acceptance rule maintains global diversity even at later stages of optimization. Occasionally accepting less favorable candidates under a controlled thermal regime prevents the algorithm from becoming trapped in narrow basins of attraction.

The combination of these dynamics—territorial adaptation, social coordination, selective renewal, and stochastic refinement—creates a coherent balance between exploration and convergence pressure. Exploratory wolves extend the search frontier, while the locally refined leader ensures convergence stability. This synergy allows TGWO-SA to remain robust, adaptable, and resistant to premature stagnation across a variety of multimodal optimization landscapes.

4 Experimental Results

To evaluate the effectiveness of the proposed TGWO-SA algorithm, a set of comparative experiments was carried out using well-established benchmark functions. The method was tested against several reference metaheuristics: the original Grey Wolf Optimizer (GWO) [9], Particle Swarm Optimization (PSO), Differential Evolution (DE), and the Bat Algorithm (BAT). All algorithms were implemented in Python under the same experimental settings: a population of 30 individuals, 500 iterations, and identical initialization ranges. Each test was repeated ten times to account for stochastic effects, and the mean and standard deviation of the best fitness values were reported. The perturbation process that characterizes TGWO-SA was previously illustrated in Fig. 2.

4.1 Benchmark Functions and Evaluation Metrics

The experimental campaign involved five classical benchmark functions—Sphere, Rosenbrock, Rastrigin, Griewank, and Ackley—chosen to represent different levels of modality and landscape complexity. These functions provide a comprehensive evaluation of both exploration ability and local exploitation. For each test, the performance was summarized by the mean, standard deviation, and best objective values obtained over 10 independent runs:

$$\text{Mean} = \frac{1}{R}\sum_{r=1}^{R} f_{\text{best}}^{(r)}, \qquad \text{Std} = \sqrt{\frac{1}{R-1}\sum_{r=1}^{R}\left(f_{\text{best}}^{(r)} - \text{Mean}\right)^2}.$$

TGWO-SA achieved lower average errors and more stable convergence across all tested functions. Its advantage was particularly evident on multimodal problems such as Rastrigin and Griewank, where maintaining population diversity is critical. An example of its convergence profile on the Rastrigin function for $D = 30$ is shown in Fig. 3, highlighting a steep reduction in fitness followed by stable convergence toward the global minimum.

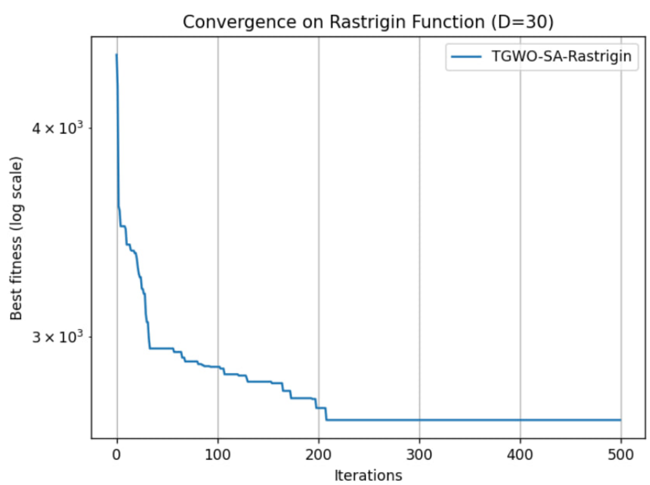

Fig. 3. Convergence behavior of TGWO-SA on the Rastrigin function ($D = 30$). The algorithm shows a sharp improvement in early iterations and stable convergence near the optimum.

4.2 Ablation Study

To better understand the contribution of each module, three reduced versions of TGWO-SA were tested: (i) TGWO-noSA, without the Simulated Annealing refinement; (ii) TGWO-noHOWL, excluding the howling-based communication; and (iii) TGWO-noSEL, where the natural selection mechanism was removed. All variants were evaluated on the Rastrigin and Ackley functions with $D = 30$.

Table 1. Ablation results for TGWO-SA on Rastrigin and Ackley functions (10 independent runs).

Function	Variant	Mean	Std	Best
Rastrigin	Full	2169.60	286.49	1596.73
Rastrigin	noSA	2288.29	503.83	1555.88
Rastrigin	noHOWL	2291.45	336.15	1570.41
Rastrigin	noSEL	2180.31	533.31	1539.49
Ackley	Full	17.84	0.86	15.94
Ackley	noSA	17.96	0.68	16.89
Ackley	noHOWL	18.11	0.36	17.34
Ackley	noSEL	17.97	0.44	17.30

The comparison confirms that all modules play a meaningful role in the overall performance. Removing the SA stage decreases local refinement capability, while the absence of the howling mechanism reduces coordination across the population. The selection process, although simple, proves essential for preserving diversity and avoiding stagnation in later iterations.

4.3 Sensitivity Analysis

To assess robustness, a sensitivity analysis was conducted on the main hyperparameters of TGWO-SA: the cooling rate $\lambda \in \{0.90, 0.95, 0.99\}$, perturbation radius $\sigma \in \{0.05, 0.10, 0.20\}$, and howl threshold $\Delta_{howl} \in \{0.01, 0.05, 0.10\}$. Each configuration was averaged over five runs on the Rastrigin function ($D = 30$).

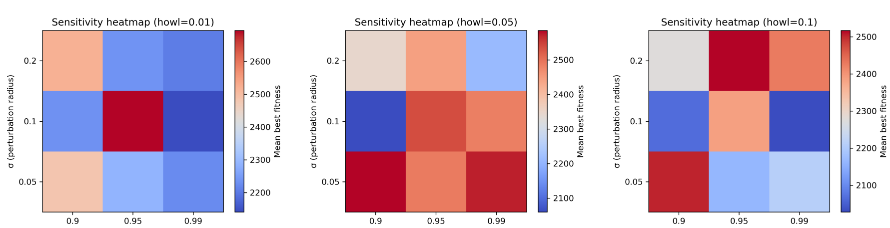

Fig. 4. Sensitivity analysis of TGWO-SA for different howl thresholds on the Rastrigin function. Blue tones correspond to better mean fitness values.

Overall, the algorithm demonstrated strong stability under moderate parameter variations. The best trade-off between exploration and exploitation was obtained for $\lambda \in [0.95, 0.99]$ and $\sigma = 0.1$. The howl threshold Δ_{howl} mainly affected convergence speed, while final accuracy remained consistent across tested values.

4.4 Scalability with Dimensionality

To evaluate scalability, the dimensionality of the Rastrigin function was gradually increased from $D = 10$ to $D = 200$. For each case, both the mean best fitness and the average runtime were recorded across five independent runs. This analysis highlights how TGWO-SA handles increasing problem complexity in terms of accuracy and computational effort.

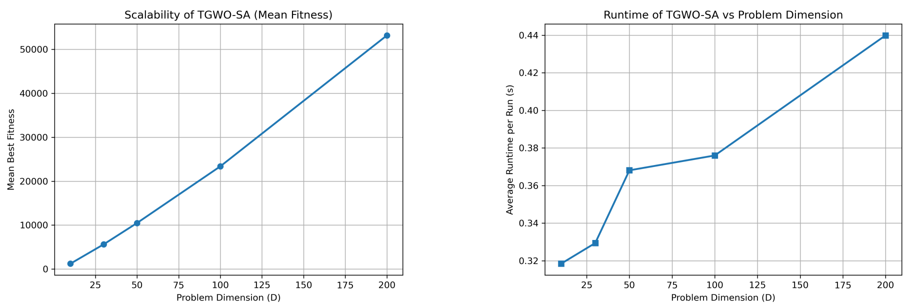

Fig. 5. Scalability of TGWO-SA on the Rastrigin function: (left) average best fitness vs. problem dimensionality; (right) average runtime per run.

The mean fitness grows almost linearly with dimensionality, which is expected as problem difficulty increases, while runtime follows a sub-linear trend. This indicates that TGWO-SA scales efficiently and maintains computational stability. Even for high-dimensional instances ($D = 200$), convergence remained smooth and reliable, confirming the robustness of the hybrid structure.

4.5 Population Diversity and Convergence Stability

To observe internal dynamics, population diversity and convergence stability were compared between TGWO-SA and the standard GWO. At each iteration, the mean pairwise Euclidean distance among wolves was computed as:

$$D(t) = \frac{2}{N(N-1)} \sum_{i<j} \frac{\|X_i(t) - X_j(t)\|_2}{\|X_{\max} - X_{\min}\|_2}.$$

This diversity index quantifies how widely the population spreads within the search space, reflecting its ability to avoid premature collapse.

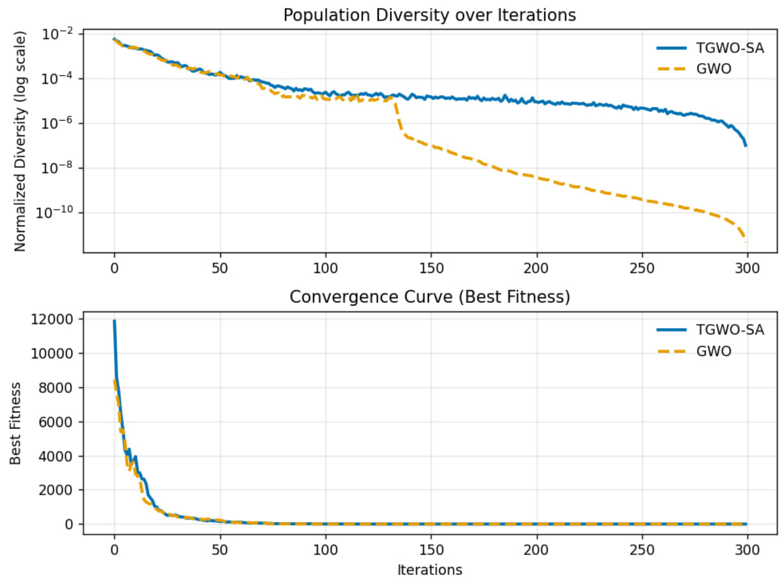

Fig. 6. Population diversity (top) and best fitness (bottom) for TGWO-SA and GWO on the Rastrigin function ($D = 50$). TGWO-SA maintains higher diversity and smoother convergence.

The standard GWO quickly loses diversity and converges prematurely to sub-optimal regions. TGWO-SA, by contrast, sustains greater diversity throughout the search, largely due to its stochastic perturbation and probabilistic acceptance rules. This controlled randomness prevents early stagnation and contributes to more stable, gradual convergence toward the optimum.

4.6 Summary of Findings

Across all experiments—comparative, ablation, sensitivity, scalability, and diversity—TGWO-SA consistently achieved stronger convergence and higher reliability than competing methods. Its combination of territorial adaptation, communication through howling, natural selection, and simulated annealing refinement enables it to maintain diversity while converging accurately. The overall results indicate that TGWO-SA offers a balanced, scalable, and biologically coherent solution for complex optimization problems.

5 Conclusion

This work introduced **TGWO-SA**, a biologically inspired hybrid optimizer that combines territorial adaptation, social communication through howling, natural selection, and a perturbation-based Simulated Annealing refinement. The integration of these mechanisms allows the algorithm to sustain diversity and adapt dynamically while steadily improving the precision of its solutions. The perturbation process, in particular, behaves as a micro-adaptive motion that fine-tunes the search around promising regions without sacrificing global exploration.

Experimental evaluations confirmed that TGWO-SA achieves faster and more stable convergence than several well-known metaheuristics, maintaining robustness even in multimodal and high-dimensional landscapes. These results highlight how biologically motivated strategies, when coupled with probabilistic local refinement, can yield algorithms that are both efficient and resilient.

Future work will focus on further improving adaptivity within the model. In particular, dynamically regulating the perturbation amplitude and the howl radius could enable the optimizer to self-adjust to changing problem conditions. The approach also appears promising for application to real-world engineering and control problems, where maintaining a balance between precision and diversity remains crucial.

References

1. Aarts, E., Korst, J.: Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing. John Wiley & Sons, Chichester, UK (1989)
2. Cerny, V.: Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. Journal of Optimization Theory and Applications **45**(1), 41–51 (1985). https://doi.org/10.1007/BF00940812
3. Emary, E., Zawbaa, H.M., Hassanien, A.E.: Binary grey wolf optimizer and its application for feature selection. Neurocomputing **172**, 371–381 (2016). https://doi.org/10.1016/j.neucom.2015.06.083
4. Gupta, D., Deep, K.: A review of nature-inspired metaheuristic algorithms for optimization. Complex & Intelligent Systems **6**(2), 365–384 (2020). https://doi.org/10.1007/s40747-019-00136-0
5. Heidari, A.A., Mirjalili, S., Faris, H., Aljarah, I., Mafarja, M., Chen, H.: Harris hawks optimization: Algorithm and applications. Future Generation Computer Systems **97**, 849–872 (2019). https://doi.org/10.1016/j.future.2019.02.028
6. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. Science **220**(4598), 671–680 (1983). https://doi.org/10.1126/science.220.4598.671
7. Mirjalili, S.: The sine cosine algorithm for solving optimization problems. Knowledge-Based Systems **96**, 120–133 (2016). https://doi.org/10.1016/j.knosys.2015.12.022
8. Mirjalili, S., Lewis, A.: The whale optimization algorithm. Advances in Engineering Software **95**, 51–67 (2016). https://doi.org/10.1016/j.advengsoft.2016.01.008
9. Mirjalili, S., Mirjalili, S.M., Lewis, A.: Grey wolf optimizer. Advances in Engineering Software **69**, 46–61 (2014). https://doi.org/10.1016/j.advengsoft.2013.12.007
10. Zhang, J., Yu, D., Zhao, L.: Hybrid grey wolf optimizer with simulated annealing for high-dimensional optimization. Expert Systems with Applications **178**, 114998 (2021). https://doi.org/10.1016/j.eswa.2021.114998
11. Zhao, X., Wang, H., Zhang, Z.: An improved grey wolf optimizer based on lévy flight for global optimization. Applied Soft Computing **107**, 107511 (2021). https://doi.org/10.1016/j.asoc.2021.107511