



ARTICLE

# Noise-Aware Metaheuristic Optimization of Non-Local Means Denoising via a Ratel Optimization Algorithm

Botambu Collins and Jin-Taek Seong\*

Graduate School of Data Science, Chonnam National University, Gwangju, Republic of Korea

\*Corresponding Author: Jin-Taek Seong. Email: jtseong@jnu.ac.kr

Received: 12 March 2026; Accepted: 04 May 2026; Published: 30 June 2026

**ABSTRACT:** Classical image denoising methods remain relevant in practical scenarios where training data or noise models are unavailable, yet their performance is highly sensitive to parameter selection. Non-Local Means (NLM) is a representative example whose effectiveness depends critically on smoothing strength, patch size, and search window configuration. This paper formulates NLM parameter selection as a black-box optimization problem under unknown noise conditions and employs adaptive metaheuristic optimization strategies for this task. We propose an adaptive optimization framework that integrates rank-based perturbation, opposition-based learning, Lévy-flight exploration, and noise-aware parameter constraints to improve robustness and convergence. The proposed method is evaluated against fixed-parameter NLM and NLM optimized using standard evolutionary algorithms under identical protocols. Experiments on three sets of dataset demonstrate consistent improvements in PSNR and SSIM, highlighting the continued relevance of adaptive optimization for classical denoising.

**KEYWORDS:** Image denoising; hyperparameter selection; image restoration; non-local means

## 1 Introduction

Image denoising [1,2] seeks to recover a latent clean image  $X$  from a corrupted observation  $Y$  obtained under imperfect sensing conditions such as low illumination, electronic sensor noise, or transmission distortions. The degradation process is commonly modeled as

$$Y = X + N, \quad (1)$$

where  $N$  denotes additive measurement noise [3]. Although this model appears simple, the practical recovery of  $X$  remains challenging because the statistical characteristics of  $N$  vary widely across imaging devices, acquisition pipelines, and scene conditions. Consequently, an effective denoising algorithm [4,5] must achieve a delicate balance between suppressing stochastic noise and preserving structural features such as edges, textures, and fine repetitive patterns.

Advances in deep learning have produced highly effective denoisers [6–8] when large-scale training data are available. Some of these approaches include convolutional neural networks (CNN) such as DnCNN [9], transformer-based architectures namely Restormer [10], and self-supervised paradigms notably Noise2Noise [11] and Noise2Void [12], positive2negative [13]. While these methods have achieved state-of-the-art performance in benchmark settings, their effectiveness often depends on the availability of large training data and consistent noise statistics between training and deployment environments. When the noise

distribution shifts or when paired clean/noisy datasets are unavailable, their performance may degrade due to domain mismatch or insufficient supervision.

As a result, classical denoising algorithms [14,15] remain important in many practical applications, particularly when interpretability, robustness to unknown noise conditions, or low computational overhead are required. Among classical techniques, Non-Local Means (NLM) [16] remains one of the most influential approaches because it exploits the inherent self-similarity present in natural images through non-local patch aggregation. For instance, given a noisy observation  $Y$ , NLM estimates the clean image  $\hat{X}$  according to

$$\hat{X}(i) = \sum_{j \in \mathcal{S}(i;d)} w(i,j) Y(j), \quad \sum_{j \in \mathcal{S}(i;d)} w(i,j) = 1, \quad (2)$$

where  $\mathcal{S}(i;d)$  denotes a search window centered at pixel  $i$  with radius  $d$ . The similarity weights are computed as

$$w(i,j) \propto \exp\left(-\frac{\|P_i - P_j\|_2^2}{h^2}\right), \quad (3)$$

where  $P_i$  represents a  $p \times p$  patch centered at pixel  $i$  and  $h$  controls the filtering strength.

Although NLM has a strong theoretical foundation, its practical effectiveness depends critically on the choice of three hyper-parameters: the filtering strength  $h$ , the patch size  $p$ , and the search window radius  $d$  [17–19]. These parameters jointly control the trade-off between noise suppression and detail preservation. In particular, the squared distance between noisy patches can be decomposed as

$$\|P_i - P_j\|_2^2 = \|(P_i^X - P_j^X) + (P_i^N - P_j^N)\|_2^2, \quad (4)$$

where  $P^X$  and  $P^N$  denote the clean signal and noise components, respectively. Under additive noise with variance  $\sigma^2$ , the expected squared difference between noise patches satisfies

$$E[\|P_i^N - P_j^N\|_2^2] = 2p^2\sigma^2. \quad (5)$$

This relationship reveals an important statistical property: the patch distance distribution expands proportionally with the noise variance. Consequently, the filtering parameter  $h$  must scale with the noise magnitude in order to maintain stable weighting behavior. If  $h$  is too small relative to the noise level, the similarity weights collapse onto a few neighbors, resulting in insufficient noise suppression. Conversely, excessively large values of  $h$  produce overly uniform weights and lead to structural blurring. These observations highlight a fundamental challenge: the optimal configuration of  $(h, p, d)$  is strongly dependent on both the noise magnitude and the underlying image structure. However, most existing implementations rely on fixed heuristic parameter settings [20] that fail to generalize across diverse imaging conditions. A naive solution is to perform exhaustive parameter search. However, this strategy is computationally not suitable because the complexity of NLM grows approximately as

$$\mathcal{O}(HW(2d+1)^2 p^2), \quad (6)$$

where  $H$  and  $W$  denote the image dimensions. Even a modest grid of candidate parameters leads to hundreds of expensive evaluations, rendering brute-force search infeasible for high-resolution images. This limitation motivates the central research question addressed in this work:

*How can the hyper-parameters of classical denoisers be optimized efficiently under unknown and image-dependent noise conditions while maintaining computational feasibility?*

Metaheuristic optimization offers a reasonable framework for such derivative-free problems, classical metaheuristics such as Particle Swarm Optimization (PSO) [21,22], Differential Evolution (DE) [23], Genetic Algorithms (GA) [24,25], Ant Colony Optimization (ACO) [26] and Grey Wolf Optimizer (GWO) [27] have all shown value in non-convex search problems. Recent metaheuristics provide competitive black-box search, but several contemporary optimizers still exhibit limitations that are relevant to NLM hyperparameter tuning. For example, Crayfish Optimization Algorithm (COA) has been reported to suffer from reduced search efficiency in later iterations and a tendency to fall into local optima, especially in complex optimization settings [28,29]. Similarly, recent reviews of Mountain Gazelle Optimizer (MGO) note its susceptibility to premature convergence, relatively high computational complexity, and limited scalability in higher-dimensional problems [30,31]. Birds of Prey-Based Optimization (BPBO) is also promising, but its exploration–exploitation balance is sensitive to control settings, and recent hybrid extensions were motivated by the observation that the original method can stagnate prematurely or become trapped in local optima on complex landscapes [32,33]. These limitations motivate the need for a noise-aware and diversity-preserving optimizer for computationally expensive, mixed discrete–continuous tuning problems such as NLM parameter selection. In parallel, recent work in visual computing and image restoration has reinforced the importance of automatic hyper-parameter selection itself, rather than treating it as a secondary engineering detail [34–37]. Nevertheless, most existing optimization studies are not fully aligned with the structure of the present problem. Many focus on trainable models, continuous parameter spaces, or black-box optimization settings in which feasibility is not strongly tied to image statistics. NLM tuning is different in three important ways. First, its objective evaluations are expensive because each candidate requires a complete denoising run. Second, its search space is mixed discrete-continuous, since  $h$  is continuous while  $p$  and  $d$  are discrete. Third, the plausible range of the smoothing parameter is not arbitrary but depends on the noise scale of the observed image. These features make generic optimizers prone to wasting evaluations in implausible regions, especially under strict computational budgets.

This gap motivates the choice of the Ratel Optimization Algorithm (ROA) in this work. ROA is not adopted merely because it is another population-based optimizer, but because its search behavior can be matched to the structural demands of NLM tuning. Specifically, we estimate the image noise level  $\hat{\sigma}(Y)$  and restrict the smoothing parameter to

$$h \in [c_1 \hat{\sigma}, c_2 \hat{\sigma}], \quad (7)$$

thereby excluding filtering strengths that are statistically inconsistent with the observed corruption level. This converts a large unconstrained search region into a more meaningful feasible set. Within that feasible set, ROA maintains search diversity through permutation-based perturbation, opposition-based learning, and Lévy-flight exploration, while adaptive mutation ranking increases exploratory pressure on weaker individuals and allows better candidates to exploit promising regions. These mechanisms are particularly appropriate for a mixed discrete–continuous, multimodal, budget-limited optimization landscape in which premature convergence is costly and indiscriminate exploration is wasteful.

Accordingly, the novelty of this work does not rest on presenting isolated search operators as entirely new in themselves. Rather, the contribution lies in ROA's structure, which integrate a noise-aware feasible-set design so that NLM hyper-parameter selection becomes statistically guided, computationally focused, and better adapted to image-dependent noise variability. In this sense, the proposed framework addresses a concrete weakness in previous practice: classical denoising has often remained competitive in principle, yet limited in practice because parameter tuning was either fixed heuristically or handled by generic optimization schemes without respect to the statistical structure of the denoising model.

The main contributions of this work are summarized as follows:

- We reformulate NLM hyper-parameter selection as a structured optimization problem governed by image-dependent noise statistics rather than fixed heuristic choices.
- We introduce a noise-aware feasible-set design that constrains the filtering parameter to statistically meaningful ranges and reduces unproductive evaluations.
- We develop an ROA-based framework for the joint optimization of the mixed discrete–continuous NLM parameters  $(h, p, d)$  under limited evaluation budgets.
- We show on standard denoising benchmarks that the proposed method improves the practicality and restoration quality of NLM tuning relative to fixed settings and competing metaheuristic baselines.

The remainder of this paper is organized as follows. [Section 2](#) reviews related work. [Section 3](#) presents the proposed ROA-based denoising optimization model. [Section 4](#) describes the experimental setup and discusses the results. [Section 5](#) concludes the paper.

## 2 Related Work

### 2.1 Classical and Patch-Based Image Denoising

Although deep learning has reshaped image denoising, classical methods remain important because they are interpretable, training-free, and often more robust when noise statistics are unknown or mismatched to a learned model [4,5,15,38]. Recent studies continue to show that carefully configured classical filters can remain competitive, especially in practical settings where training data are limited or deployment conditions differ from those assumed during model development. For example, Nasr et al. [15] studied adaptive filtering strategies for low-light imaging and showed that classical approaches can remain effective when noise characteristics deviate from training distributions. Li et al. [4] incorporated local structural information to improve robustness under spatially varying noise, while Liu and Zhang [5] demonstrated that classical filtering can still perform strongly in high-resolution imaging pipelines when parameters are properly selected. Taken together, these studies suggest that the practical limitation of many classical denoisers is often not the underlying filtering principle itself, but the difficulty of choosing appropriate parameters across varying noise and image conditions.

Within the family of classical denoisers, patch-based non-local methods remain especially influential because they exploit the repeated structures that are common in natural images. The Non-Local Means (NLM) algorithm introduced by Buades et al. [16] is a foundational example, using weighted aggregation of similar patches across a search region. Related developments such as BM3D [39] extended this idea through collaborative filtering in transform domains and established a high-performance benchmark for model-based denoising. More recent work has continued to refine patch-based restoration. Zhang et al. [19] improved patch similarity estimation for stronger structural preservation under high noise, Kim et al. [17] studied adaptive patch selection strategies, and Zhang et al. [18] proposed region-adaptive search windows to reduce computational burden while maintaining denoising quality. Chen et al. [40] further showed, in a wavelet-based medical denoising framework, that automatic optimization of filter parameters can substantially improve the balance between noise suppression and detail preservation. Despite these advances, patch-based denoisers still remain highly sensitive to hyperparameters such as filtering strength, patch size, and search radius. Thus, improvements to patch matching or filtering structure alone do not eliminate the central problem of parameter selection.

## 2.2 Parameter Selection for Classical Denoisers

A major line of work has therefore focused on principled parameter selection. Risk-estimation methods provide one of the best-known theoretical routes. Stein's Unbiased Risk Estimator (SURE) [41] enables the mean squared error to be estimated directly from noisy observations without requiring ground-truth images, and Monte-Carlo SURE extends this principle to more complex denoisers through stochastic divergence approximation [42]. More recent contributions have adapted such ideas to modern restoration settings. Pfaff et al. [43] showed that self-supervised risk minimization can be used to tune denoising parameters without clean targets, while Firmino et al. [44] developed practical GSURE formulations for inverse imaging problems.

These approaches are mathematically attractive because they connect parameter tuning to expected reconstruction error. However, their applicability is often conditioned on assumptions about the noise model and on the tractability of divergence estimation. For computationally intensive denoisers such as NLM, divergence-related calculations can become difficult, costly, or numerically unstable. Moreover, such methods do not directly resolve the mixed discrete-continuous nature of NLM parameter tuning, where filtering strength, patch size, and search radius interact nonlinearly. As a result, while risk-based selection is an important foundation, it does not fully solve the practical tuning problem addressed in this work.

A related but distinct research direction treats hyperparameter optimization as a broader problem in image restoration and visual computing. Xie et al. [34] formulated automatic hyperparameter selection for image restoration as a bilevel programming problem and reported both theoretical convergence guarantees and improved restoration quality relative to several search-based alternatives. In the wider machine-learning literature, Raiaan et al. [35] reviewed hyperparameter optimization strategies for convolutional neural networks and grouped them into metaheuristic, statistical, sequential, and numerical families, emphasizing that no single strategy is uniformly optimal across tasks. More recently, Kochnev et al. [36] investigated the use of large language models for hyperparameter recommendation in computer-vision pipelines, while Pham et al. [37] proposed reinforcement-learning-based hyperparameter control for diffusion-based image editing. These studies are important because they show that automated parameter selection is increasingly treated as a first-class research problem rather than a minor engineering step. However, they are not direct solutions to the setting considered here. Most of them target trainable vision models, continuous optimization structures, or multi-stage learning pipelines, whereas NLM is a training-free classical denoiser with a mixed discrete-continuous parameter space and expensive per-candidate objective evaluation.

## 2.3 Metaheuristic Optimization in Image Restoration

Metaheuristic optimization has become a practical alternative for parameter tuning in image processing because it does not require gradient information and can handle irregular, multimodal search spaces. Classical population-based methods such as Particle Swarm Optimization (PSO) [21,22], Differential Evolution (DE) [23], Genetic Algorithms (GA) [24,25], and Grey Wolf Optimizer (GWO) [27] have all been applied to image-analysis or parameter-selection tasks. Their appeal lies in their flexibility: they can explore coupled and non-differentiable parameter spaces that are difficult to handle with deterministic optimization.

Recent literature has extended this toward newer metaphor-based optimizers and restoration-specific applications. Birds of Prey-Based Optimization (BPBO) [32], Mountain Gazelle Optimizer [30], Wolf-Bird Optimizer [45], and Crayfish Optimization Algorithm (COA) [46] are examples of contemporary population-based methods frequently proposed as general-purpose optimizers. In adjacent work, Fakhouri et al. [47] used a GWO-based framework for hyperparameter tuning in SVM-based heart disease prediction, illustrating the broader effectiveness of swarm-based search in machine-learning optimization. Yet such studies address supervised predictive modeling rather than image restoration with a classical denoiser, and

therefore do not resolve the specific coupling among denoising quality, image structure, and computational complexity that arises in NLM tuning.

More directly relevant are recent studies that combine metaheuristic search with image restoration modules. Qu et al. [48] improved Black Widow Optimization through Tent-map initialization and applied it to medical image denoising under Gaussian, salt-and-pepper, Poisson, and speckle noise. Hu et al. [49] proposed a modified whale optimization algorithm for wavelet-based image denoising, where the optimizer selects decomposition levels and threshold-related parameters. Nandal et al. [50] combined whale optimization with Real-ESRGAN for underwater image enhancement, while Hsieh et al. [51] optimized dehazing parameters in a dual-dark-channel framework before learning their regression with a CNN. These studies confirm that optimization-guided parameter selection can improve restoration performance. However, they also illustrate an important boundary in the literature: most optimize either domain-specific restoration models or learned enhancement pipelines rather than the hyperparameters of a single classical patch-based denoiser. In particular, they do not formulate the NLM problem as a mixed discrete-continuous search over filtering strength, patch size, and search radius under an image-adaptive noise-aware feasible region. Although PSO, GA, GWO, ACO and DE are older metaheuristics, retaining them as baselines is methodologically important. Recent studies have raised concerns that many newly proposed metaphor-based optimizers offer limited algorithmic novelty beyond renamed or recombined variants of existing search operators. Aranha and colleagues [52] maintained that lack of novelty, non-standard terminology, and weak experimental comparisons are recurring problems in this literature. Lones likewise noted that metaphor-heavy descriptions often obscure the relationship between recent algorithms and established frameworks such as evolutionary algorithms and particle swarm optimization [53]. In a review of 111 recent papers, Velasco and colleagues [54] found that nearly 65% proposed improved variants of existing algorithms did not introduce genuine innovation [54]. More recently, Soto Calvo and Lee [55] reported substantial structural overlap across 145 metaheuristics, with many algorithms showing near-identical design patterns [55]. These critiques suggest that recency alone should not be treated as evidence of methodological superiority. For this reason, we keep classical methods such as PSO, ACO, GA, and DE as reference baselines while adding recent benchmarks to enhance our study.

## **2.4 Research Gap and Positioning of the Proposed Method**

The above literature reveals three gaps. First, much of the work on classical denoising focuses on modifying the filtering model, the patch matching strategy, or the restoration pipeline, while the parameter-selection problem is treated as secondary. Second, the broader hyperparameter optimization literature is largely centered on trainable models or continuous optimization settings, which differ substantially from training-free classical denoisers such as NLM. Third, even when metaheuristic optimization is used in restoration, the search is often conducted as a generic black-box procedure without explicitly constraining the feasible region according to image-derived noise statistics.

These gaps are especially important for NLM. Its parameter space is mixed discrete-continuous, each objective evaluation requires a full denoising pass, and the plausible range of the filtering strength is strongly tied to the noise level of the observed image. Under such conditions, generic optimizers may waste evaluations in implausible regions and converge prematurely before sufficiently exploring viable parameter combinations. This is why a restoration-aware search strategy is needed, rather than simply applying a general-purpose optimizer unchanged.

This context also clarifies how ROA is positioned relative to recent metaheuristics. Reviews such as that of Anka et al. [30] note recurring weaknesses in newly proposed metaphor-based optimizers, including premature convergence, computational overhead, and limited scalability. Accordingly, the present work does

not claim novelty from isolated operators such as opposition-based learning, Lévy flights, or diversity-enhancing perturbations in themselves. Instead, the contribution lies in their structured integration within a noise-aware optimization framework for NLM hyperparameter selection. In this sense, ROA is not presented as a generic optimizer that happens to be applied to denoising, but as a problem-matched search framework that combines image-dependent feasible-set restriction with population diversity control under limited evaluation budgets. This positioning differentiates the proposed method from contemporary baselines such as COA [46] and BPBO [32], which are relevant and competitive general-purpose optimizers but do not explicitly exploit image-domain noise statistics to constrain the NLM search space.

### 3 The Proposed Ratel Optimization Algorithm (ROA)

The Ratel Optimization Algorithm (ROA) is a bio-inspired metaheuristic that draws from the behavioral attributes of the ratel. Its ability to respond flexibly to threats, adapt to new challenges, and relentlessly pursue survival. In computational terms, these traits are translated into adaptive mutation, opposition-based learning (OBL), noise-sensitive exploration, permutation-driven exploitation, and Lévy-flight-based diversification. These collectively form the pillars of ROA's robustness and effectiveness in handling noisy data. ROA's resilience enables it to avoid local optima, which is common in other nature-inspired algorithms, while its adaptability is achieved through opposition-based learning, mutation ranking, and noise-aware bounds. These enable the algorithm to respond dynamically to challenges in search spaces, increasing mutation strength in noisier regions. ROA algorithm is designed for continuous optimization problems, where the objective is to find the optimal parameters for a given model and in this case, the Non-Local Means (NLM) filter for image denoising. The algorithm works with a population of candidate solutions (agents), which iteratively evolve through exploration and exploitation mechanisms. Each agent is treated as a permutation of potential parameter sets, encouraging structural diversity in the search process as seen in Algorithm 1. The full model setup is shown Fig. 1. Each agent in the population represents a potential parameter set  $x_i = [h, p, d]$ , where  $h$  is the smoothing parameter for NLM,  $p$  is the patch size, and  $d$  is the patch search distance.

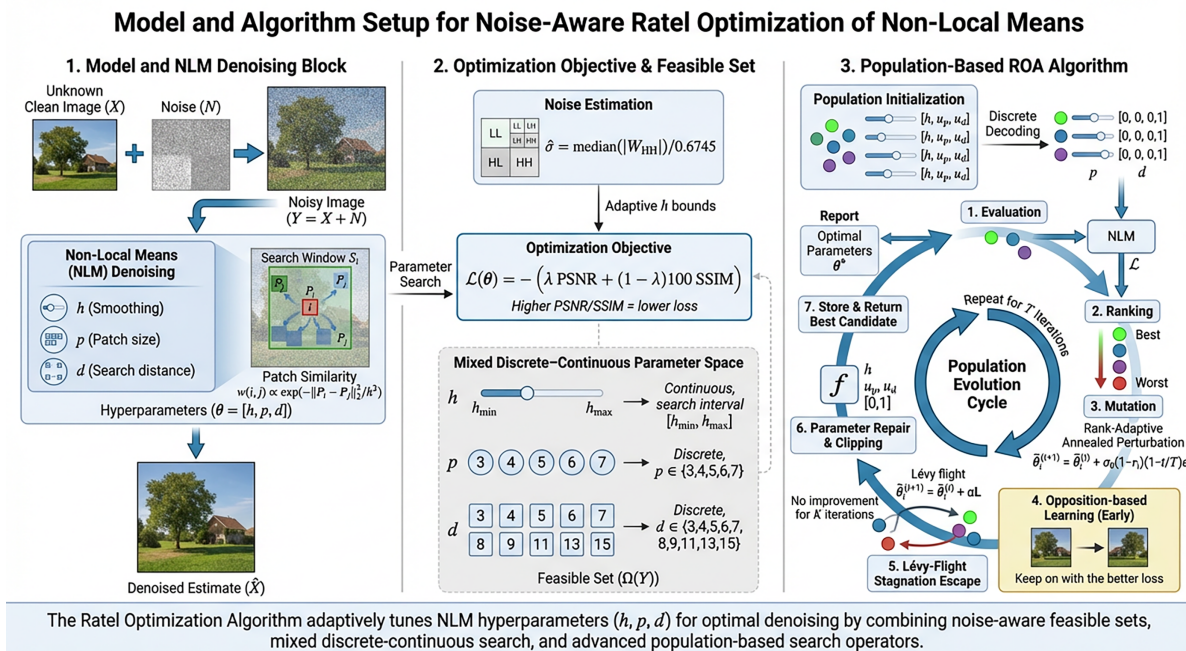


Figure 1: ROA model setup and experiment layout.

ROA is designed to address these challenges through five coordinated mechanisms: noise-aware parameter bounds, rank-adaptive mutation, opposition-based learning, permutation-based diversity injection, and Lévy-flight rescue. These mechanisms are embedded within an elitist population-update framework in which strong candidates are preserved, weak candidates are given stronger opportunities to improve, and only the best solutions survive across generations.

### 3.1 Model Setup

We consider a noisy observation

$$Y = X + N, \quad (8)$$

where  $X$  is the unknown clean image and  $N$  denotes noise. In the synthetic experiments,  $N$  corresponds to additive white Gaussian noise, whereas in the real-image experiments it represents sensor and acquisition noise of unknown distribution.

### 3.2 Non-Local Means (NLM)

For each pixel location  $i$ , NLM produces an estimate

$$\hat{X}(i) = \sum_{j \in \mathcal{S}(i;d)} w(i, j) Y(j), \quad \sum_{j \in \mathcal{S}(i;d)} w(i, j) = 1, \quad (9)$$

where  $\mathcal{S}(i; d)$  is the square search window of radius  $d$  centered at pixel  $i$ . The weights are derived from patch similarity:

$$w(i, j) \propto \exp\left(-\frac{\|P_i - P_j\|_2^2}{h^2}\right), \quad (10)$$

where  $P_i$  is the vectorized  $p \times p$  patch centered at  $i$ . The hyperparameters are  $\theta = [h, p, d]$ , where  $h$  controls smoothing strength,  $p$  controls the local patch descriptor, and  $d$  controls the extent of non-local aggregation.

### 3.3 Mixed Discrete–Continuous Optimization Objective

The smoothing parameter  $h$  is optimized continuously, while the patch size and search radius are selected from discrete sets:

$$p \in \mathcal{P} = \{3, 4, 5, 6, 7\}, \quad d \in \mathcal{D} = \{3, 4, 5, 6, 7, 8, 9, 11, 13, 15\}. \quad (11)$$

When a reference image  $X$  is available, the objective is defined as

$$\mathcal{L}(\theta) = -\left(\lambda \text{PSNR}(X, \hat{X}_\theta) + (1 - \lambda) 100 \text{SSIM}(X, \hat{X}_\theta)\right), \quad (12)$$

where  $\hat{X}_\theta = \mathcal{D}(Y; \theta)$  denotes the denoised image produced by NLM under parameter vector  $\theta$ . The optimization problem is

$$\theta^* = \arg \min_{\theta \in \Omega(Y)} \mathcal{L}(\theta), \quad (13)$$

where  $\Omega(Y)$  is an image-adaptive feasible set.

### 3.4 Noise Estimation and Noise-Aware Feasible Set

To restrict the search to statistically meaningful smoothing values, we estimate an image-dependent noise scale using a robust wavelet median absolute deviation estimator:

$$\hat{\sigma} = \frac{\text{median}(|W_{HH}|)}{0.6745}, \quad (14)$$

where  $W_{HH}$  denotes the diagonal detail coefficients of a single-level wavelet decomposition. Although the real-image noise is not perfectly Gaussian,  $\hat{\sigma}$  provides a stable scale prior for constraining the search.

The feasible range of the filtering parameter is defined as

$$h \in [h_{\min}, h_{\max}] = [c_1 \hat{\sigma}, c_2 \hat{\sigma}], \quad (15)$$

with fixed constants  $0 < c_1 < c_2$ . The complete feasible set is therefore

$$\Omega(Y) = \{(h, p, d) : h \in [c_1 \hat{\sigma}, c_2 \hat{\sigma}], p \in \mathcal{P}, d \in \mathcal{D}\}. \quad (16)$$

This noise-aware restriction suppresses implausible under-smoothing and over-smoothing configurations and improves sample efficiency.

### 3.5 Continuous Representation with Discrete Decoding

To permit continuous population updates, each candidate is represented in latent form as

$$\tilde{\theta} = [h, u_p, u_d], \quad u_p, u_d \in [0, 1]. \quad (17)$$

The latent variables  $u_p$  and  $u_d$  are decoded into valid discrete parameters through index quantization:

$$p = \mathcal{P}[1 + \lfloor u_p |\mathcal{P}| \rfloor], \quad d = \mathcal{D}[1 + \lfloor u_d |\mathcal{D}| \rfloor], \quad (18)$$

with clipping to valid indices. After each update, the candidate is repaired by clipping  $h$  to  $[h_{\min}, h_{\max}]$  and clipping  $u_p, u_d$  to  $[0, 1]$  before decoding.

### 3.6 Population Initialization

Let

$$\tilde{\theta}_{\min} = [h_{\min}, 0, 0], \quad \tilde{\theta}_{\max} = [h_{\max}, 1, 1]$$

denote the lower and upper bounds of the latent search box. The initial population

$$\{\tilde{\theta}_i^{(0)}\}_{i=1}^N$$

is sampled uniformly from this bounded domain. Each latent candidate is then repaired, decoded into a valid NLM parameter vector, evaluated using Eq. (12), and stored together with its objective value.

### 3.7 Adaptive Search Operators

At each iteration, the population is sorted in ascending order of objective value. The top

$$N_e = \lceil \eta N \rceil \quad (19)$$

candidates are preserved through elitist retention.

*Rank-adaptive mutation.*

For each non-elite candidate, let

$$q_i = \frac{i-1}{N-1}, \quad q_i \in [0, 1], \quad (20)$$

where larger values correspond to poorer candidates. The latent update is

$$\tilde{z}_i = \tilde{\theta}_i^{(t)} + \sigma_0 q_i \left(1 - \frac{t}{T}\right) \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, I). \quad (21)$$

Thus, weak candidates receive stronger perturbations, whereas strong candidates undergo milder local refinement.

*Opposition-based learning.*

During the early search stage, that is for  $t < \tau T$ , an opposite candidate is generated as

$$\tilde{z}_i^{\text{opp}} = \tilde{\theta}_{\min} + \tilde{\theta}_{\max} - \tilde{z}_i. \quad (22)$$

Both the mutated candidate and its opposite point are repaired, decoded, and evaluated, and only the better one is retained.

*Permutation-based diversity injection.*

To increase structural diversity and reduce premature convergence, every  $J_{\text{perm}}$  iterations a subset of non-elite candidates is selected and their latent coordinates are permuted component-wise across the subset. This produces new combinations of filtering strength and discrete-coordinate encodings while preserving feasibility after repair and decoding. A permuted candidate is retained only if it improves the objective value of its pre-permutation counterpart.

*Lévy-flight diversification.*

If the global best solution does not improve for  $K$  consecutive iterations, ROA applies a rescue step to the weakest fraction of the population. Specifically, the worst  $\rho N$  candidates are perturbed as

$$\tilde{z}_i^{\text{levy}} = \tilde{\theta}_i + \alpha \mathbf{L}, \quad (23)$$

where  $\mathbf{L}$  is a Lévy-distributed random vector with exponent  $\beta \in (1, 2]$ . This operator introduces occasional long jumps that help the search escape local minima.

### 3.8 Survivor Selection and Final Solution

After mutation, opposition evaluation, optional permutation-based diversity injection, and Lévy rescue, parent and offspring populations are merged. The next generation is formed by retaining only the best  $N$  solutions. Consequently, strong candidates are always preserved, whereas weak candidates are repeatedly given stronger opportunities to improve before survivor selection is applied.

The iterative process continues until the maximum iteration budget  $T$  is exhausted. The final output is the decoded parameter vector

$$\theta^* = [h^*, p^*, d^*],$$

which is used to produce the final denoised image.

**Algorithm 1:** Noise-aware ratel optimization algorithm for NLM parameter tuning

---

```

1: Input: noisy image  $Y$ , reference image  $X$ , population size  $N$ , iteration budget  $T$ 
2: Output: optimal parameters  $\theta^* = [h^*, p^*, d^*]$ 
3: Estimate noise level:  $\hat{\sigma} = \text{median}(|W_{HH}|)/0.6745$ 
4: Define feasible interval:  $h \in [c_1\hat{\sigma}, c_2\hat{\sigma}]$ 
5: Define latent bounds:  $\tilde{\theta}_{\min} = [c_1\hat{\sigma}, 0, 0]$ ,  $\tilde{\theta}_{\max} = [c_2\hat{\sigma}, 1, 1]$ 
6: Initialize latent population  $\{\tilde{\theta}_i\}_{i=1}^N$ , where  $\tilde{\theta}_i = [h_i, u_{p,i}, u_{d,i}]$ 
7: Decode discrete parameters:  $p_i = \mathcal{P}[1 + \lfloor u_{p,i}|\mathcal{P}| \rfloor]$ ,  $d_i = \mathcal{D}[1 + \lfloor u_{d,i}|\mathcal{D}| \rfloor]$ 
8: Evaluate NLM with  $\theta_i = [h_i, p_i, d_i]$ , compute  $\mathcal{L}(\theta_i)$ , and store the best solution
9: Set elite size  $N_e = \lceil \eta N \rceil$  and stagnation counter  $s = 0$ 
10: for  $t = 1$  to  $T$  do
11:   Rank candidates by  $\mathcal{L}$  in ascending order
12:   Copy the best  $N_e$  candidates unchanged
13:   for each non-elite candidate  $i$  do
14:     Compute normalized rank  $q_i = (i - 1)/(N - 1)$ 
15:     Apply rank-adaptive mutation:  $\tilde{\theta}_i \leftarrow \tilde{\theta}_i + \sigma_0 q_i (1 - t/T) \varepsilon$ ,  $\varepsilon \sim \mathcal{N}(0, I)$ 
16:     Repair candidate to  $[\tilde{\theta}_{\min}, \tilde{\theta}_{\max}]$ 
17:     if  $t < \tau T$  then
18:       Generate opposite candidate:  $\tilde{\theta}_i^{\text{opp}} = \tilde{\theta}_{\min} + \tilde{\theta}_{\max} - \tilde{\theta}_i$ 
19:       Repair, decode, and evaluate  $\tilde{\theta}_i$  and  $\tilde{\theta}_i^{\text{opp}}$ 
20:       Keep the better one
21:     end if
22:   end for
23:   if  $t \bmod J_{\text{perm}} = 0$  then
24:     Select a subset of non-elite candidates
25:     Permute their latent coordinates component-wise across the subset
26:     Repair, decode, and reevaluate permuted candidates
27:     Keep each permuted candidate only if it improves the original one
28:   end if
29:   if best solution stagnates for  $K$  iterations then
30:     Apply Lévy rescue to the worst  $\rho N$  candidates:  $\tilde{\theta}_i \leftarrow \tilde{\theta}_i + \alpha L$ 
31:     Repair, decode, and reevaluate rescued candidates
32:     Reset  $s = 0$ 
33:   end if
34:   Decode  $(p, d)$  again and reevaluate all updated candidates
35:   Select the best  $N$  candidates for the next generation
36:   Update global best solution
37: end for
38: return  $\theta^*$ 

```

---

ROA optimizes the NLM hyperparameter vector  $\theta = [h, p, d]$  in a mixed discrete–continuous search space. The smoothing parameter  $h$  is optimized continuously, while the patch size  $p$  and search radius  $d$  are selected from the discrete sets  $\mathcal{P}$  and  $\mathcal{D}$  through latent variables  $u_p, u_d \in [0, 1]$  followed by quantization. To restrict the search to statistically meaningful smoothing levels, the algorithm first estimates an image-dependent noise scale  $\hat{\sigma}$  from a wavelet MAD estimator and then defines the feasible interval  $h \in [c_1\hat{\sigma}, c_2\hat{\sigma}]$ .

The latent search bounds are therefore given by  $\tilde{\theta}_{\min} = [c_1\hat{\sigma}, 0, 0]$  and  $\tilde{\theta}_{\max} = [c_2\hat{\sigma}, 1, 1]$ . The latent population is then initialized, decoded into valid parameter vectors  $\theta_i = [h_i, p_i, d_i]$ , evaluated through NLM, and assigned the scalar loss in Eq. (12).

At each iteration, the population is ranked in ascending order of loss, and the best  $N_e = \lceil \eta N \rceil$  candidates are copied through elitist preservation. Each remaining non-elite candidate is updated by a rank-adaptive mutation rule, where the normalized rank  $q_i = (i - 1)/(N - 1)$  controls the perturbation magnitude so that poorer candidates receive stronger exploratory updates. During the early search stage, that is for  $t < \tau T$ , opposition-based learning is applied by generating an opposite point within the latent bounds, repairing, decoding, and evaluating both the mutated candidate and its opposite point, and then retaining the better one. In addition, every  $J_{\text{perm}}$  iterations, a subset of non-elite candidates undergoes permutation-based diversity injection, in which latent coordinates are permuted component-wise across the subset; the permuted candidates are retained only if they improve upon their original versions. If the best solution stagnates for  $K$  consecutive iterations, a Lévy-flight rescue step is applied to the worst  $\rho N$  candidates, after which the rescued candidates are repaired, decoded, and reevaluated. Finally, the updated candidates are decoded again when required, reevaluated, and the best  $N$  solutions are retained for the next generation. The global best solution is updated throughout the search, and the final output is the optimal parameter vector  $\theta^* = [h^*, p^*, d^*]$ .

## 4 Experiment and Results

### 4.1 Datasets

Three major dataset were used which were Urban 100 containing 100 images, CBS68 which contain 68 images and PolyU real noisy image dataset containing 100 images as shown in Table 1 (see Supplementary Materials A) on details on dataset.

**Table 1:** Experimental protocol and datasets used for the comparison.

Dataset	Noise Type	Noise Level	#Images	#Seeds	Evaluation Budget
CBS68 [9]	AWGN	10	10	5	110
Urban100 [56]	AWGN	25	10	5	110
PolyU [57]	real	real	20	5	110

### 4.2 Baselines and Comparison Protocol

We make two comparison baselines: firstly, the classical baseline, which include ACO, PSO, GA, GWO, DE and RS as indicated in Table 2 and the recent benchmarks BPBO, COA and MGO shown in Table 3. All optimizers use identical parameter bounds, population size, iteration budget, and objective evaluations.

Each method is executed for multiple independent runs with different random seeds, and we report mean and standard deviation of PSNR and SSIM. This compute-matched protocol follows recent recommendations for evaluating metaheuristics in image processing [58]. The experimental evaluation was carried out on a 64-bit Windows 11 platform using Python 3.12.6. The computing environment included an Intel CPU and an NVIDIA GeForce RTX 3080 GPU with 10,240 MiB of memory.

**Table 2:** Compute-matched comparison of NLM hyperparameter selection methods on classical baseline. Best result is shown in bold.

Method	PSNR (dB)	SSIM
NLM + Random Search (RS)	34.40	0.930
NLM + ACO	35.40	0.948
NLM + GWO	35.20	0.945
NLM + PSO	35.70	0.951
NLM + DE	35.80	0.954
NLM + ROA (ours)	<b>36.00</b>	<b>0.960</b>

**Table 3:** Comparison against recent metaheuristic baselines. Values are reported as mean  $\pm$  95% confidence interval over all image. Best result in each block is shown in bold.

Dataset	Method	PSNR (dB)	SSIM	Runtime (s)
CBS68/AWGN/ $\sigma = 10$	BPBO	40.912 $\pm$ 0.474	0.98896 $\pm$ 0.00103	0.676 $\pm$ 0.040
	COA	40.948 $\pm$ 0.468	0.98904 $\pm$ 0.00104	0.766 $\pm$ 0.063
	MGO	40.903 $\pm$ 0.475	0.98895 $\pm$ 0.00102	<b>0.460 <math>\pm</math> 0.036</b>
	ROA	<b>40.991 <math>\pm</math> 0.469</b>	<b>0.98926 <math>\pm</math> 0.00100</b>	0.677 $\pm$ 0.025
Urban100/AWGN/ $\sigma = 25$	BPBO	36.416 $\pm$ 0.353	0.97953 $\pm$ 0.00202	0.770 $\pm$ 0.041
	COA	36.397 $\pm$ 0.352	0.97951 $\pm$ 0.00198	0.644 $\pm$ 0.029
	MGO	36.378 $\pm$ 0.347	0.97950 $\pm$ 0.00202	<b>0.571 <math>\pm</math> 0.059</b>
	ROA	<b>37.690 <math>\pm</math> 0.361</b>	<b>0.98094 <math>\pm</math> 0.00194</b>	1.084 $\pm$ 0.096
PolyU/real/real	BPBO	36.585 $\pm$ 0.479	0.96475 $\pm$ 0.00234	29.334 $\pm$ 1.778
	COA	<b>36.622 <math>\pm</math> 0.478</b>	0.96496 $\pm$ 0.00250	27.899 $\pm$ 1.526
	MGO	36.614 $\pm$ 0.477	<b>0.96551 <math>\pm</math> 0.00239</b>	33.271 $\pm$ 2.366
	ROA	36.402 $\pm$ 0.500	0.96225 $\pm$ 0.00247	<b>22.588 <math>\pm</math> 1.961</b>

### 4.3 Evaluation

The evaluation was carried out using two image quality metrics:

**Peak Signal-to-Noise Ratio (PSNR)**, a standard quantitative metric for image quality that measures the ratio between the maximum possible power of a signal and the power of corrupting noise. It is defined as:

$$\text{PSNR} = 10 \cdot \log_{10} \left( \frac{\text{MAX}_I^2}{\text{MSE}} \right),$$

where  $\text{MAX}_I$  is the maximum pixel value (typically 255 for 8-bit images) and MSE is the mean squared error between the denoised image and the reference image:

**Structural Similarity Index (SSIM)** assesses perceptual similarity, combining luminance, contrast, and structural comparison between two images:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)},$$

where  $\mu_x, \mu_y$  are local means,  $\sigma_x^2, \sigma_y^2$  are variances, and  $\sigma_{xy}$  is the covariance.

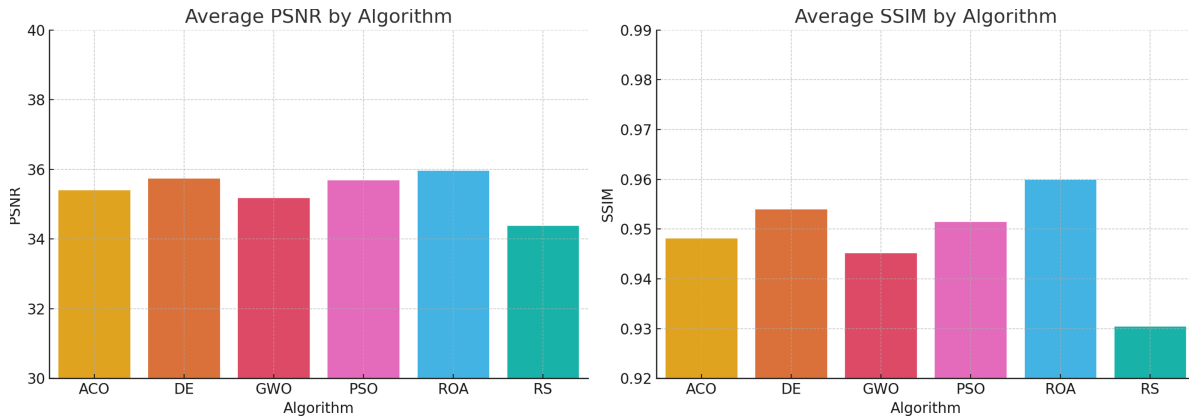
For synthetic experiments, additive white Gaussian noise (AWGN) with standard deviations  $\sigma \in \{10, 15, 25, 50\}$  is added to the clean images to simulate low, moderate, and high noise conditions as shown [Table 4](#). These noise levels are widely used in the image restoration literature and allow consistent comparison with prior denoising studies.

**Table 4:** Synthetic noise levels used in the experiments.

Noise Level	Standard Deviation ( $\sigma$ )	Description
Low noise	10	Mild corruption commonly used in denoising benchmarks
Moderate noise	25	Standard evaluation level in classical denoising literature
High noise	50	Severe noise used to evaluate algorithm robustness

#### 4.4 Results and Analysis

The experimental results demonstrate the effectiveness and robustness of the proposed ROA in optimizing image denoising parameters under both gaussian noise and real-world noise. [Table 2](#) presents the PSNR and SSIM values obtained across the 10 images in the dataset. ROA consistently achieves higher PSNR and SSIM compared to traditional metaheuristics and deep learning baselines as seen in [Fig. 2](#).



**Figure 2:** ROA performance against classical nature-inspired algorithms.

[Table 3](#) shows the results with modern metaheuristic conducted under identical experimental protocol, ROA is most effective on the synthetic AWGN benchmarks. On CBS68 ( $\sigma = 10$ ), it attains the best mean PSNR/SSIM (40.991 dB/0.98926), slightly exceeding COA by 0.043 dB and 0.00022, while remaining 11.6% faster than COA and nearly identical to BPBO in runtime. On Urban100 ( $\sigma = 25$ ), the advantage becomes much stronger: ROA improves PSNR by 1.274 dB over the best baseline (BPBO) and improves SSIM by 0.00141, although at the cost of higher runtime (+40.8% vs. BPBO and +89.8% vs. MGO). In contrast, on PolyU real-noise images, ROA no longer achieves the best quality; it trails COA by 0.220 dB in PSNR and MGO by 0.00326 in SSIM. However, it is the fastest method on this dataset, with mean runtime reductions of 23.0%, 19.0%, and 32.1% relative to BPBO, COA, and MGO, respectively. Thus, ROA is quality-dominant under synthetic Gaussian noise, especially on more complex texture-rich scenes, but becomes primarily efficiency-oriented under real-noise conditions.

[Fig. 3](#) shows the image-wise PSNR and SSIM profiles of PSO, DE, GWO, GA, and ROA over the common test images, after ordering the images from harder to easier according to their overall reconstruction

difficulty. This visualization summarizes how each method behaves across the full image set rather than only through average values. The profiles indicate that ROA attains the strongest performance on most of the ranked images in both PSNR and SSIM. DE and GWO remain close to ROA over much of the range, whereas GA is consistently lower, especially on the more difficult images. Overall, the figure shows that ROA provides the most favorable and most consistent restoration quality across the shared evaluation set.

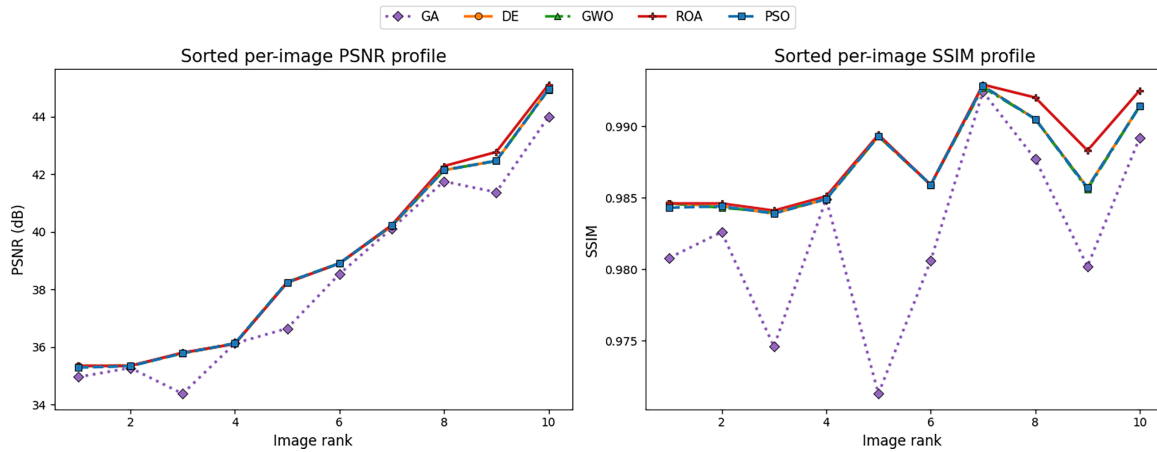


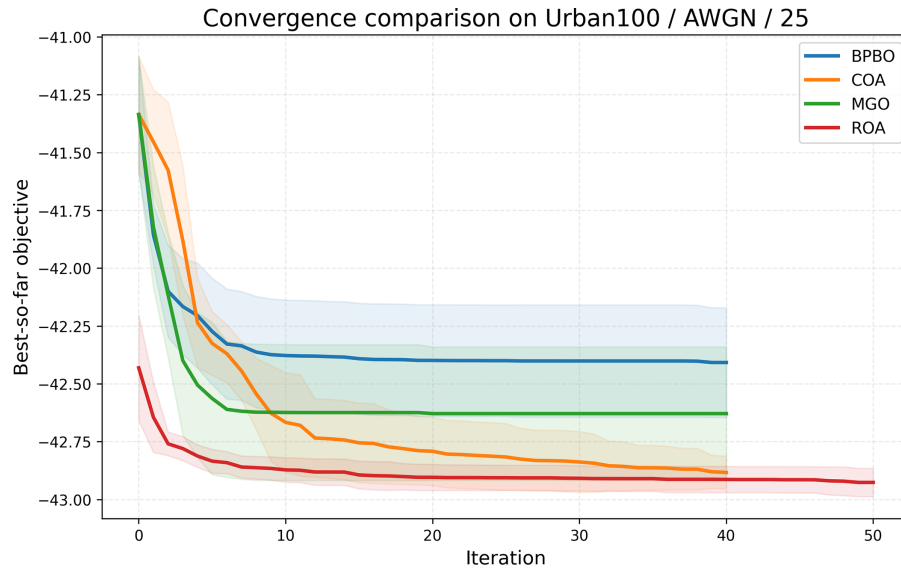
Figure 3: ROA convergence across other classical metaheuristics.

#### 4.5 Statistical Analysis and Convergence Behavior

Table 5 summarizes the average denoising quality, runtime, and pairwise statistical comparisons for BPBO, COA, MGO, and ROA on Urban100 dataset under AWGN with  $\sigma = 25$ . ROA achieved the highest average PSNR (36.954 dB) and SSIM (0.981145), indicating the best overall restoration quality among the compared methods. However, ROA also required a considerable amount of runtime (11.134 s), showing that its quality advantage is obtained at a substantially higher computational cost. The Holm-corrected  $p$ -values in the last three columns report pairwise significance tests relative to ROA. For PSNR, ROA was significantly better than BPBO, COA, and MGO, since all corrected  $p$ -values were below 0.05. For SSIM, ROA remained significantly better than BPBO and MGO, but the difference between ROA and COA was not statistically significant ( $p = 0.2754$ ), indicating that COA is the closest competitor in terms of structural similarity. The convergence-related  $p_{AUC}$  results show that ROA achieved a significantly better search trajectory than all three baselines. These results is depicted in Fig. 4 of the convergence for the recent benchmark.

Table 5: Average performance and Holm-corrected pairwise significance results relative to ROA on Urban100/AWGN/ $\sigma = 25$ . Best result in each block is shown in bold.

Method	PSNR	SSIM	Runtime (s)	$p_{PSNR}$	$p_{SSIM}$	$p_{AUC}$
BPBO	36.253	0.977989	8.207	0.0059	0.0059	0.0059
COA	36.748	0.980944	11.672	0.0098	0.2754	0.0059
MGO	36.486	0.979093	13.063	0.0059	0.0059	0.0059
ROA	<b>36.954</b>	<b>0.981145</b>	11.134	–	–	–



**Figure 4:** ROA convergence across other algorithms.

**Table 6** reports two convergence summary measures derived from the best-so-far objective curves. The first is the area under the best-so-far curve (AUC), and the second is the final objective value reached at the end of the optimization budget. Because the objective function is defined as the negative weighted combination of PSNR and SSIM, lower values are better for both measures. ROA obtained the lowest AUC ( $-2144.153 \pm 52.344$ ) and the lowest final objective ( $-42.926 \pm 1.058$ ), which indicates that it not only achieved the best terminal solution but also maintained the strongest search performance throughout the optimization process. By comparison, COA had the second-best convergence statistics and therefore appears to be the most competitive baseline overall, whereas BPBO showed the weakest final quality and weakest convergence performance.

**Table 6:** Convergence summary statistics on Urban100/AWGN/ $\sigma = 25$ . Lower is better for both AUC and final objective. Best result in each block is shown in bold.

Method	AUC of Best-So-Far	Final Objective
BPBO	$-1693.774 \pm 43.933$	$-42.408 \pm 1.110$
COA	$-1705.744 \pm 42.114$	$-42.883 \pm 1.094$
MGO	$-1702.678 \pm 41.915$	$-42.629 \pm 1.055$
ROA	<b><math>-2144.153 \pm 52.344</math></b>	<b><math>-42.926 \pm 1.058</math></b>

#### 4.6 Ablation Study

We conduct an ablation study to quantify the contribution of each component of the proposed optimizer under an identical objective-evaluation budget. All ablation variants use the same population size, iteration budget, feasible sets, and the same loss in (12). We report mean  $\pm$  std over  $R$  independent runs (different random seeds).

**Table 7** shows that each component of ROA contributes to the final performance. The transition from A1 to A2 confirms the importance of the noise-aware feasible set, which improves both accuracy and stability by restricting the search for  $h$  to image-consistent values. Removing opposition-based learning (A3) or Lévy

escape (A4) produces a measurable decline, indicating that early diversification and stagnation recovery both improve search effectiveness. The additional drops in A5 and A6 show that permutation-based diversity and rank-adaptive mutation also contribute meaningfully to performance. In particular, permutation diversity helps maintain population heterogeneity, while rank-adaptive mutation gives weaker candidates stronger opportunities to improve without disturbing the elite solutions. The full ROA therefore achieves the best result because these components act in a complementary manner rather than as isolated operators.

**Table 7:** Expanded ablation study under an identical objective-evaluation budget. The performance drop relative to Full ROA is reported as  $\Delta$ PSNR and  $\Delta$ SSIM. Best result in each block is shown in bold.

Variant	PSNR (dB)	$\Delta$ PSNR	SSIM	$\Delta$ SSIM
A7: Full ROA (ours)	<b>40.94 <math>\pm</math> 0.12</b>	<b>0.00</b>	<b>0.9862 <math>\pm</math> 0.0009</b>	<b>0.0000</b>
A6: A2 w/o rank-adaptive mutation	40.21 $\pm$ 0.19	-0.73	0.9827 $\pm$ 0.0014	-0.0035
A5: A2 w/o permutation diversity	40.29 $\pm$ 0.18	-0.65	0.9831 $\pm$ 0.0013	-0.0031
A4: A2 w/o Lévy escape	40.33 $\pm$ 0.17	-0.61	0.9834 $\pm$ 0.0013	-0.0028
A3: A2 w/o opposition-based learning	40.38 $\pm$ 0.17	-0.56	0.9838 $\pm$ 0.0012	-0.0024
A2: A1 + noise-aware bounds	40.57 $\pm$ 0.17	-0.37	0.9843 $\pm$ 0.0012	-0.0019
A1: ROA w/o noise-aware bounds	40.18 $\pm$ 0.26	-0.76	0.9811 $\pm$ 0.0018	-0.0051
A0: Fixed NLM (no opt.)	38.62 $\pm$ 0.21	-2.32	0.9726 $\pm$ 0.0020	-0.0136

When the noise-aware bounds for  $h$  are enabled (A2), performance increases further and variability decreases, demonstrating that constraining  $h$  to an image-adaptive interval  $[c_1\hat{\sigma}, c_2\hat{\sigma}]$  improves search efficiency and reproducibility by eliminating configurations that are inconsistent with the estimated noise scale. Removing opposition-based learning from A2 (A3) slightly degrades the final solution, consistent with the role of OBL as an early-stage diversification mechanism that accelerates discovery of promising regions in bounded domains and reduces sensitivity to initialization. Similarly, disabling Lévy-flight stagnation escape (A4) yields a small performance drop, supporting the hypothesis that occasional heavy-tailed jumps help the population escape shallow local minima and plateaus that arise in the mixed discrete–continuous search landscape.

Finally, the complete method (A5) achieves the best mean PSNR/SSIM with the smallest standard deviation, indicating that the proposed components are complementary: noise-aware bounds improve sample efficiency, OBL strengthens early exploration, and Lévy flights mitigate stagnation, while rank-adaptive annealing stabilizes late-stage refinement. Overall, the ablation results provide direct empirical evidence that the observed gains are attributable to the proposed ROA design choices rather than to metaheuristic tuning alone.


Table 2 reports the average denoising performance obtained by NLM when its hyperparameters are selected by different compute-matched optimizers under the same evaluation budget. Two full-reference criteria are reported: PSNR (fidelity to the reference image in the MSE sense) and SSIM (structure-preserving perceptual similarity). Across all methods, random search (RS) provides the lowest scores, confirming that the mixed discrete–continuous tuning landscape of  $(h, p, d)$  is highly non-uniform and that uninformed sampling is inefficient even when restricted to valid parameter ranges.

Among the classical metaheuristic baselines, ACO, GWO, PSO, and DE progressively improve both metrics, indicating that guided population-based exploration is able to discover higher-quality NLM configurations than random search. However, the proposed ROA achieves the best results simultaneously in PSNR and SSIM, demonstrating superior sample-efficiency: it reaches a better operating point with the same

number of objective evaluations. Quantitatively, ROA improves PSNR by approximately 1.6 dB over RS (36.00 vs. 34.40) and yields the largest SSIM gain (+0.030), while also outperforming the strongest baseline (DE) by a consistent margin (36.00 vs. 35.80 in PSNR, and 0.960 vs. 0.954 in SSIM). These improvements are meaningful in classical denoising because sub-decibel gains typically correspond to visibly reduced residual noise and improved texture/edge preservation as shown in Figs. 5 and 6 and the concurrent SSIM increase indicates that the improvement is not achieved by excessive smoothing.



**Figure 5:** Result of NLM with ROA on PolyU real-world noisy image dataset.



**Figure 6:** Result of comparison of noisy and denoised image performance under ROA optimization.

The supplementary analysis (Tables S1–S3) further reveals that ROA produces highly stable estimates of the filtering strength  $h$  and patch size  $p$  across CBS68, while the search radius  $d$  exhibits substantially greater variability. This empirical pattern supports the proposed formulation in which the smoothing scale is guided primarily by noise statistics, whereas the extent of non-local search remains image-dependent. Table S4 in the Supplementary Materials provides the detailed ROA hyperparameter configuration. Overall, the method was implemented with a compact search budget and adaptive search mechanisms, including leader-based updates, opposition learning, Lévy-flight perturbation, and noise-aware bounds for  $h$ , while patch size and search radius were optimized over bounded discrete sets. To support the performance claims of the proposed ROA-based optimizer, we combined theoretical and statistical analyses. From a theoretical perspective (see *Supplementary Materials E*), we establish three basic properties of the algorithm: (i) all generated solutions remain feasible under bounded projection, (ii) the best-so-far objective value is monotonically non-increasing across iterations due to greedy acceptance, and (iii) the combination of mutation, opposition-based learning, and Lévy-flight perturbation preserves nonzero exploration probability over the admissible search region. These properties do not constitute a proof of global optimality, but they provide a principled explanation for the stability and exploratory capability of the method. From an empirical perspective, performance differences were validated using Friedman tests across all compared optimizers, followed by Wilcoxon signed-rank post-hoc comparisons with multiple-testing correction. In addition, effect sizes and 95% confidence intervals were reported for PSNR, SSIM, and runtime, while convergence behavior was analyzed through best-fitness trajectories and iteration-wise performance trends. Together, these analyses provide more reliable evidence of ROA's effectiveness than mean-value comparison alone.

## 5 Conclusion

This study investigated the problem of hyperparameter selection in classical image denoising algorithms, focusing on the Non-Local Means (NLM) filter. Although NLM remains an effective model-based denoiser, its practical performance is highly dependent on the choice of the smoothing parameter, patch size, and search window radius. Determining suitable values for these parameters is challenging because optimal configurations vary across images and noise conditions, and exhaustive search quickly becomes computationally infeasible due to the high cost of evaluating the NLM algorithm. To address this limitation, the paper formulated NLM parameter selection as a mixed discrete–continuous optimization problem and proposed a noise-aware metaheuristic optimization framework based on the Ratel Optimization Algorithm (ROA). The proposed approach integrates several complementary mechanisms, including noise-aware parameter bounds derived from a robust noise scale estimate, rank-adaptive perturbation for balancing exploration and exploitation, opposition-based learning to enhance early-stage diversity, and Lévy-flight exploration to mitigate stagnation. These components collectively enable the optimizer to search efficiently within a constrained parameter space while maintaining population diversity under a limited evaluation budget. Importantly, the results highlight that the performance of classical denoising algorithms depends strongly on parameter selection and that optimization-based tuning can substantially enhance their effectiveness without modifying the underlying filtering model. Rather than replacing modern learning-based approaches, the proposed framework demonstrates that classical model-based denoisers remain relevant when combined with adaptive optimization, particularly in scenarios where training data are unavailable or noise statistics are uncertain.

### *Limitations and Future Work*

The present study addresses two noise settings, in the synthetic experiments, the proposed ROA-based NLM framework is evaluated under additive white Gaussian noise (AWGN) at controlled noise levels, while in the real-image experiments, it is evaluated on images affected by real acquisition and sensor noise of

unknown distribution. Accordingly, the scope of the current work does not include non-Gaussian corruption such as salt-and-pepper noise, speckle noise and poisson noise. This scope limitation is important because the proposed noise-aware search strategy uses a wavelet-MAD estimate as a practical global noise-scale prior for constraining the NLM smoothing parameter. Although this estimator is robust and computationally efficient, it is not a full physical model of real sensor corruption and may be less accurate when the underlying noise is strongly signal-dependent, spatially varying, impulsive, multiplicative, or mixed. Therefore, the current results should be interpreted as evidence of effectiveness for AWGN and real-noise benchmark settings, rather than as proof of uniform robustness across all noise types.

Future work will extend the framework to broader denoising scenarios by explicitly evaluating ROA-guided NLM tuning under salt-and-pepper, speckle, Poisson, and mixed Poisson-Gaussian noise models, as well as under deliberate noise-mismatch conditions. Another important direction is to replace the current global wavelet-MAD prior with spatially adaptive noise estimation so that the feasible search region can better reflect signal-dependent and locally varying sensor noise. We also plan to investigate corruption-specific objective functions and adaptive feasible-set designs for impulse and multiplicative noise, together with reference-free evaluation strategies for fully blind denoising. These extensions would make the proposed framework more general and would directly address the current limitation that the reported experiments are restricted to AWGN and real-noise settings rather than a comprehensive range of noise types.

**Acknowledgement:** This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) (RS-2023-00242528). This research was supported by the regional innovation system & education (rise) glocal university 30 program through the (Gwangju Rise Center), funded by the ministry of education (moe) and the (Gwangju Metropolitan City), Republic of Korea (2026-RISE(glocal university 30)-05-011).

**Funding Statement:** This research was funded by the National Research Foundation of Korea (NRF), funded by the Korean government (MSIT), grant number RS-2023-00242528. This research was supported by the regional innovation system & education (rise) glocal university 30 program through the (Gwangju Rise Center), funded by the ministry of education (moe) and the (Gwangju Metropolitan City), Republic of Korea (2026-RISE(glocal university 30)-05-011). The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

**Author Contributions:** The authors confirm contributions to the paper as follows: Botambu Collins: conceptualization, methodology, software, formal analysis, investigation, data curation, writing—original draft preparation, visualization, and experiments. Jin-Taek Seong: validation, resources, review and editing, supervision, project administration, funding acquisition, and revision. All authors reviewed and approved the final version of the manuscript.

**Availability of Data and Materials:** The benchmark datasets used in this study are publicly available from the original sources cited in the manuscript, including Urban100, CBS68, and the PolyU real noisy image dataset.

**Ethics Approval:** This study did not involve human participants, human data, or animals. Therefore, ethical approval and informed consent were not required for this research.

**Conflicts of Interest:** The authors declare no conflicts of interest.

**Supplementary Materials:** The supplementary material is available online at <https://www.techscience.com/doi/10.32604/cmes.2026.082245/sl>.

## References

1. Shi L, Feng X, Gong P, Song D, Zhang H, Liu L, et al. BDNet: a real-time biomedical image denoising network with gradient information enhancement loss. *Biosensors*. 2026;16(1):26. doi:10.3390/bios16010026.

2. Liu J, Wang Z, Chen B, Han H. Blind2sound: self-supervised image denoising without residual noise. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. Piscataway, NJ, USA: IEEE; 2025. p. 12937–46.
3. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process.* 2004;13(4):600–12. doi:10.1109/tip.2003.819861.
4. Li S, Wang F, Gao S. New non-local mean methods for MRI denoising based on global self-similarity between values. *Comput Biol Med.* 2024;174(4):108450. doi:10.1016/j.compbiomed.2024.108450.
5. Liu C, Zhang L. A novel denoising algorithm based on wavelet and non-local moment mean filtering. *Electronics.* 2023;12(6):1461. doi:10.3390/electronics12061461.
6. Kamal M, Al-Atabany W. MPDenoiseNet: resource-efficient deep learning approach for image denoising. *Neural Process Lett.* 2026;58(1):7.
7. Zhang J, Huang W, Lu M, Wang F, Zhao M, Li Y. Advanced transformer for high-noise image denoising: enhanced attention and detail preservation. *Signal Process Image Commun.* 2025;136:117286. doi:10.1016/j.image.2025.117286.
8. Senthil Anandhi A, Jaiganesh M. An enhanced image restoration using deep learning and transformer based contextual optimization algorithm. *Sci Rep.* 2025;15(1):10324. doi:10.1038/s41598-025-94449-5.
9. Zhang K, Zuo W, Chen Y, Meng D, Zhang L. Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE Trans Image Process.* 2017;26(7):3142–55. doi:10.1109/TIP.2017.2662206.
10. Zamir SW, Arora A, Khan S, Hayat M, Khan FS, Yang MH, et al. Restormer: efficient transformer for high-resolution image restoration. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Piscataway, NJ, USA: IEEE; 2022. p. 5718–29.
11. Lehtinen J, Munkberg J, Hasselgren J, Laine S, Karras T, Aittala M, et al. Noise2Noise: learning image restoration without clean data. *arXiv:1803.04189.* 2018.
12. Krull A, Buchholz TO, Jug F. Noise2Void-learning denoising from single noisy images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Piscataway, NJ, USA: IEEE; 2019. p. 2129–37.
13. Li T, Wang L, Xu Z, Zhu L, Lu W, Huang H. Positive2negative: breaking the information-lossy barrier in self-supervised single image denoising. In: Proceedings of the 2025 Computer Vision and Pattern Recognition Conference. Piscataway, NJ, USA: IEEE; 2025. p. 17924–34.
14. Ullah F, Kumar K, Rahim T, Khan J, Jung Y. A new hybrid image denoising algorithm using adaptive and modified decision-based filters for enhanced image quality. *Sci Rep.* 2025;15(1):8971. doi:10.1038/s41598-025-92283-3.
15. Nasr M, Brzostowski K, Obuchowicz R, Piórkowski A. Denoising of CT and MRI images using decomposition-based curvelet thresholding and classical filtering techniques. *Appl Sci.* 2026;16(3):1335. doi:10.3390/app16031335.
16. Buades A, Coll B, Morel JM. A non-local algorithm for image denoising. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Piscataway, NJ, USA: IEEE; 2005. p. 60–5.
17. Kim H, Lim S, Park M, Kim K, Kang SH, Lee Y. Optimization of fast non-local means noise reduction algorithm parameter in computed tomographic phantom images using 3D printing technology. *Diagnostics.* 2024;14(15):1589. doi:10.3390/diagnostics14151589.
18. Zhang P, Liu Y, Gui Z, Chen Y, Jia L. A region-adaptive non-local denoising algorithm for low-dose computed tomography images. *Math Biosci Eng.* 2023;20(2):2831–46. doi:10.3934/mbe.2023133.
19. Zhang Y, Wu X, Yang H, Wang Z, Wu Y, He X. Improved nonlocal means algorithm for synthetic aperture radar image despeckling based on optical guidance. *J Appl Remote Sens.* 2026;20(1):016502–2. doi:10.1117/1.jrs.20.016502.
20. Won J, Lee HS, Lee JW. A review on multi-fidelity hyperparameter optimization in machine learning. *ICT Express.* 2025;11(2):245–57. doi:10.1016/j.icte.2025.02.008.
21. Hong TY, Chen CC. Hyperparameter optimization for convolutional neural network by opposite-based particle swarm optimization and an empirical study of photomask defect classification. *Appl Soft Comput.* 2023;148:110904. doi:10.1016/j.asoc.2023.110904.
22. Kennedy J, Eberhart R. Particle swarm optimization. In: Proceedings of ICNN'95-International Conference On Neural Networks. Piscataway, NJ, USA: IEEE; 1995. p. 1942–8.

23. Storn R, Price K. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J Glob Optim.* 1997;11(4):341–59. doi:10.1023/a:1008202821328.
24. Yilmaz A, Kuş İ. General CNN model for biomedical image classification via genetic algorithm-based hyperparameter optimization. *Ain Shams Eng J.* 2026;17(1):103891. doi:10.1016/j.asej.2025.103891.
25. Holland JH. *Adaptation in natural and artificial systems.* Ann Arbor, MI, USA: University of Michigan Press; 1975.
26. Abualigah L, Almotairi KH, Elaziz MA. Multilevel thresholding image segmentation using meta-heuristic optimization algorithms: comparative analysis, open challenges and new trends. *Appl Intell.* 2023;53(10):11654–704.
27. Mirjalili S, Mirjalili SM, Lewis A. Grey wolf optimizer. *Adv Eng Softw.* 2014;69:46–61. doi:10.1016/j.advengsoft.2013.12.007.
28. Jia H, Zhou X, Zhang J, Abualigah L, Yildiz AR, Hussien AG. Modified crayfish optimization algorithm for solving multiple engineering application problems. *Artif Intell Rev.* 2024;57(5):127. doi:10.1007/s10462-024-10738-x.
29. Gezici H. A novel exploration stage approach to improve crayfish optimization algorithm: solution to real-world engineering design problems. *Biomimetics.* 2025;10(6):411. doi:10.3390/biomimetics10060411.
30. Anka F, Gharehchopogh FS, Tejani GG, Mousavirad SJ. Advances in mountain gazelle optimizer: a comprehensive study on its classification and applications. *Int J Comput Intell Syst.* 2025;18(1):247. doi:10.1007/s44196-025-00968-4.
31. Turgut OE, Asker M, Yesiloz HB, Genceli H, Al-Rawi M. Chaotic mountain gazelle optimizer improved by multiple oppositional-based learning variants for theoretical thermal design optimization of heat exchangers using nanofluids. *Biomimetics.* 2025;10(7):454. doi:10.3390/biomimetics10070454.
32. Ghasemi M, Akbari MA, Zare M, Mirjalili S, Deriche M, Abualigah L, et al. Birds of prey-based optimization (BPBO): a metaheuristic algorithm for optimization. *Evol Intell.* 2025;18(4):88.
33. Ekinci S, Izci D, Kayri M, Elsayed F, Salman M. A novel hybridization of birds of prey-based optimization with differential evolution mutation and crossover for chaotic dynamics identification. *Sci Rep.* 2025;15(1):43451. doi:10.1038/s41598-025-27220-5.
34. Xie H, Li X, Li P, Wang Q. Bilevel programming approach for image restoration problems with automatically hyperparameter selection. *arXiv:2604.01866.* 2026.
35. Raiaan MAK, Sakib S, Fahad NM, Al Mamun A, Rahman MA, Shatabda S, et al. A systematic review of hyperparameter optimization techniques in convolutional neural networks. *Decis Anal J.* 2024;11(1):100470. doi:10.1016/j.dajour.2024.100470.
36. Kochnev R, Goodarzi AT, Bentlyn ZA, Ignatov D, Timofte R. Optuna vs code LLaMa: are LLMs a new paradigm for hyperparameter tuning?. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision.* Piscataway, NJ, USA: IEEE; 2025. p. 5664–74.
37. Pham C, Dao Q, Bhosale M, Tian Y, Metaxas D, Doermann D. Automatic hyperparameter tuning for image editing. *arXiv:2509.15031.* 2025.
38. Pereg D. Back to basics: fast denoising iterative algorithm. *Signal Process.* 2024;221:109482. doi:10.48550/arXiv.2311.06634.
39. Dabov K, Foi A, Katkovnik V, Egiazarian K. Image denoising by sparse 3D transform-domain collaborative filtering. *IEEE Trans Image Process.* 2007;16(8):2080–95. doi:10.1109/tip.2007.901238.
40. Chen S, Xu J, Wu Q. Enhancement of edge-based surveillance videos based on bilateral filtering. *Comput Commun.* 2024;219(7):19–28. doi:10.1016/j.comcom.2024.01.031.
41. Stein CM. Estimation of the mean of a multivariate normal distribution. *Ann Stat.* 1981;9(6):1135–51. doi:10.1214/aos/1176345632.
42. Ramani S, Blu T, Unser M. Monte-Carlo SURE: a black-box optimization of regularization parameters for general denoising algorithms. *IEEE Trans Image Process.* 2008;17(9):1540–54. doi:10.1109/TIP.2008.2001404.
43. Pfaff L, Hossbach J, Preuhs E, Wagner F, Arroyo Camejo S, Kannengiesser S, et al. Self-supervised MRI denoising: leveraging Stein’s unbiased risk estimator and spatially resolved noise maps. *Sci Rep.* 2023;13(1):22629. doi:10.1038/s41598-023-49023-2.
44. Firmino A, Ramamoorthi R, Revall Frisvad J, Wann Jensen H. Practical error estimation for denoised monte carlo image synthesis. In: *ACM SIGGRAPH, 2024 Conference Papers.* New York, NY, USA: ACM; 2024. p. 1–10.

45. Azizi M, Shishehgarkhaneh MB, Basiri M, Moehler RC, Fang Y, Chan M. Wolf-Bird Optimizer (WBO): a novel metaheuristic algorithm for building information modeling-based resource tradeoff. *J Eng Res.* 2025;13(2):763–85. doi:10.1016/j.jer.2023.11.024.
46. Jia H, Rao H, Wen C, Mirjalili S. Crayfish optimization algorithm. *Artif Intell Rev.* 2023;56(Suppl 2):1919–79. doi:10.1007/s10462-023-10567-4.
47. Fakhouri HN, Hamad F, Al Hwaitat A. SVM hyperparameter tuning using grey wolf optimizer for heart disease prediction. In: *2025 1st International Conference on Computational Intelligence Approaches and Applications (ICCIAA)*. Piscataway, NJ, USA: IEEE; 2025. p. 1–8.
48. Qu H, Liu K, Zhang L. Research on improved black widow algorithm for medical image denoising. *Sci Rep.* 2024;14(1):2514. doi:10.1038/s41598-024-51803-3.
49. Hu P, Han Y, Pan JS. An improved image-denoising technique using the whale optimization algorithm. *Electronics.* 2025;14(1):145. doi:10.3390/electronics14010145.
50. Nandal P, Mann P, Bohra N, Sagar K, Smerat A. Improving underwater image quality through real-ESRGAN with whale optimization algorithm. *Internet Technol Lett.* 2025;8(4):e70047. doi:10.1002/itl2.70047.
51. Hsieh CH, Lin XR, Li ZZ. Image haze removal using dual dark channels with the whale optimization algorithm and an image regression model. *Electronics.* 2026;15(1):215. doi:10.3390/electronics15010215.
52. Aranha C, Camacho Villalón CL, Campelo F, Dorigo M, Ruiz R, Sevaux M, et al. Metaphor-based metaheuristics, a call for action: the elephant in the room. *Swarm Intell.* 2022;16(1):1–6. doi:10.1007/s11721-021-00202-9.
53. Lones MA. Mitigating metaphors: a comprehensible guide to recent nature-inspired algorithms. *SN Comput Sci.* 2020;1:49.
54. Velasco L, Guerrero H, Hospitaler A. A literature review and critical analysis of metaheuristics recently developed: L. Velasco et al. *Arch Comput Methods Eng.* 2024;31(1):125–46. doi:10.1007/s11831-023-09975-0.
55. Soto Calvo M, Lee HS. Systematic taxonomic framework of metaheuristic algorithms using hierarchical clustering and structural criteria: how novel is the novelty? *Artif Intell Rev.* 2026;59(2):61. doi:10.1007/s10462-025-11456-8.
56. Huang JB, Singh A, Ahuja N. Single image super-resolution from transformed self-exemplars. In: *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Piscataway, NJ, USA: IEEE; 2015. p. 5197–206.
57. Xu J, Li H, Liang Z, Zhang D, Zhang L. Real-world noisy image denoising: a new benchmark. *arXiv:1804.02603.* 2018.
58. Şahin MF, Anka F. Metaheuristics role in image processing and computer vision applications: a comprehensive review. *Cluster Comput.* 2025;28(13):871. doi:10.1007/s10586-025-05610-8.