

Received: 12 March 2026; Accepted: 04 May 2026

---

## Supplementary Material

This supplementary document provides additional experimental analysis supporting the proposed ROA-based hyperparameter optimization framework for Non-Local Means denoising. In particular, we report statistical summaries of the optimized parameters across the CBS68 dataset and provide additional visualizations of parameter distributions and optimization convergence behavior.

### ***A: Datasets***

To evaluate the effectiveness of the proposed Ratel Optimization Algorithm (ROA) for adaptive Non-Local Means (NLM) parameter tuning, experiments are conducted on multiple widely used benchmark datasets as well as real-world noisy images. These datasets are commonly adopted in the image denoising literature because they contain diverse structural patterns and allow both controlled and realistic evaluation scenarios.

**Urban100 Dataset.** The Urban100 dataset contains 100 high-resolution images of urban scenes characterized by strong geometric structures, repetitive patterns, and dense textures such as building facades, windows, and architectural layouts. These characteristics make the dataset particularly challenging for denoising algorithms because repetitive structures must be preserved while suppressing noise. Since the Non-Local Means (NLM) algorithm relies on patch similarity across spatially distant regions, Urban100 provides an appropriate benchmark for evaluating whether optimized parameters can effectively exploit self-similarity while avoiding structural blurring.

**CBS68 Dataset.** The CBS68 dataset consists of 68 natural images extracted from the Berkeley Segmentation Dataset (BSD). It contains diverse scenes including natural landscapes, objects, smooth regions, and textured areas. Due to its balanced mixture of image structures and moderate resolution, CBS68 has become one of the most widely used benchmark datasets for evaluating classical image denoising algorithms. In this study, CBS68 is primarily used to provide quantitative comparisons between the proposed ROA-based optimization method and baseline NLM configurations.

**Real-World Noisy Images.** In addition to synthetic benchmarks, we use the PolyU real-world noisy image dataset. The dataset consists of real photographs captured using consumer-grade cameras under varying ISO settings. Unlike synthetic experiments where noise is artificially injected, real-world images contain complex noise generated by sensor electronics, camera pipelines, and illumination variations. Evaluating the proposed framework on real noisy images allows assessment of its robustness when the noise distribution is unknown and must be estimated directly from the observed image.

### B. Per-Image ROA Optimization Results on CBS68 (Table S1)

**Table S1:** Optimal NLM hyperparameters obtained using the proposed ROA algorithm for each image in the CBS68 dataset.

Image	$h$	$p$	$d$	PSNR (dB)	SSIM
0000.png	0.0108	4	3	45.14	0.9883
0001.png	0.0100	3	3	40.24	0.9869
0002.png	0.0100	3	3	42.88	0.9852
0003.png	0.0100	3	5	40.30	0.9864
0004.png	0.0100	3	6	40.89	0.9886
0005.png	0.0100	3	3	40.69	0.9843
0006.png	0.0100	3	3	40.27	0.9930
0007.png	0.0100	3	11	40.08	0.9972
0008.png	0.0100	3	6	40.30	0.9902
0009.png	0.0100	3	3	38.96	0.9908
0010.png	0.0100	3	3	41.20	0.9873
0011.png	0.0100	3	3	40.87	0.9882
0012.png	0.0100	3	3	38.91	0.9913
0013.png	0.0100	3	3	41.08	0.9874
0014.png	0.0100	3	3	40.40	0.9878
0015.png	0.0100	3	3	40.27	0.9840
0016.png	0.0100	3	3	41.11	0.9883
0017.png	0.0100	3	3	41.04	0.9858
0018.png	0.0100	3	3	41.37	0.9929
0019.png	0.0100	3	8	40.72	0.9880
0020.png	0.0100	3	3	42.27	0.9864
0021.png	0.0100	3	3	40.59	0.9844
0022.png	0.0100	3	3	40.39	0.9936
0023.png	0.0100	3	3	42.68	0.9838
0024.png	0.0100	3	3	40.86	0.9802
0025.png	0.0100	3	3	41.63	0.9819
0026.png	0.0100	3	3	40.73	0.9940
0027.png	0.0100	3	3	40.52	0.9910
0028.png	0.0100	3	3	40.90	0.9912
0029.png	0.0100	3	3	40.64	0.9929
0030.png	0.0100	3	3	42.35	0.9900
0031.png	0.0100	3	3	40.71	0.9807
0032.png	0.0100	3	3	39.14	0.9779
0033.png	0.0100	3	8	40.27	0.9721
0034.png	0.0100	3	3	40.64	0.9813
0035.png	0.0105	3	9	40.19	0.9912
0036.png	0.0100	3	11	40.19	0.9936
0037.png	0.0100	3	4	42.16	0.9934
0038.png	0.0100	3	3	42.24	0.9860
0039.png	0.0100	3	3	41.12	0.9541
0040.png	0.0100	3	3	41.33	0.9825
0041.png	0.0100	3	3	40.49	0.9829
0042.png	0.0100	3	3	41.23	0.9912
0043.png	0.0100	3	3	41.13	0.9882
0044.png	0.0100	3	3	41.88	0.9830
0045.png	0.0100	3	3	40.89	0.9882
0046.png	0.0100	3	3	41.07	0.9875
0047.png	0.0100	3	3	43.40	0.9850
0048.png	0.0100	3	3	40.16	0.9900
0049.png	0.0100	3	9	40.23	0.9916
0050.png	0.0100	3	3	39.97	0.9792
0051.png	0.0100	3	3	39.53	0.9893
0052.png	0.0100	3	3	40.57	0.9830
0053.png	0.0100	3	3	39.87	0.9758
0054.png	0.0100	3	3	41.85	0.9748

Image	$h$	$p$	$d$	PSNR (dB)	SSIM
0055.png	0.0100	3	3	40.70	0.9894
0056.png	0.0100	3	4	40.35	0.9894
0057.png	0.0100	3	11	39.91	0.9893
0058.png	0.0100	3	4	41.75	0.9878
0059.png	0.0100	3	3	41.96	0.9767
0060.png	0.0100	3	3	42.08	0.9842
0061.png	0.0100	3	3	42.08	0.9769
0062.png	0.0100	3	3	41.87	0.9782
0063.png	0.0100	3	3	43.05	0.9902
0064.png	0.0100	3	11	40.09	0.9922
0065.png	0.0100	3	4	40.35	0.9865
0066.png	0.0100	3	3	40.58	0.9830
0067.png	0.0100	3	3	38.68	0.9938

Table S1 reports the per-image optimal NLM hyperparameters obtained by ROA on CBS68. The results indicate that ROA consistently converges to a narrow range of filtering strengths, while allowing the search radius to vary more substantially across images. This suggests that the noise-aware feasible set effectively regularizes the search for  $h$ , whereas the extent of non-local aggregation remains dependent on image structure.

### C. Statistical Summary of Optimized Parameters (Table S2)

**Table S2:** Statistical summary of optimized NLM parameters and resulting image quality metrics across the CBS68 dataset using ROA optimization.

Metric	Mean	Standard Deviation
PSNR (dB)	40.94	1.09
SSIM	0.9862	0.0066
Filtering strength $h$	0.01002	0.00011
Patch size $p$	3.01	0.12
Search radius $d$	3.97	2.25

Table S2 summarizes the distribution of the optimized parameters over the CBS68 dataset. The extremely small variance of the filtering strength  $h$  indicates stable convergence toward a common smoothing regime, while the low variance of the patch size  $p$  suggests that a small local patch is generally preferred under the evaluated conditions. In contrast, the larger variance of the search radius  $d$  implies that this parameter is more sensitive to image-specific self-similarity and texture complexity.

### D. Parameter Frequency Analysis (Table S3)

Table S3 reports the frequency of the selected discrete parameter values. The strong dominance of a single patch size confirms that ROA identifies a largely consistent local matching scale across the dataset. By comparison, the broader distribution of search-radius values shows that the amount of non-local aggregation is more adaptive and depends on the structural characteristics of individual images.

**Table S3:** Frequency of selected patch sizes and search radii across CBS68 images.

Parameter Value	Frequency	Percentage
$p = 3$	67	98.5%
$p = 4$	1	1.5%
$d = 3$	47	69.1%
$d = 4$	4	5.9%
$d = 5$	1	1.5%
$d = 6$	2	2.9%
$d = 8$	2	2.9%
$d = 9$	2	2.9%
$d = 11$	4	5.9%

### E. Parameter Frequency Analysis (Table S4)

**Table S4:** ROA hyperparameters used in the experiments.

Parameter	Value
Population size $N$	10
Maximum iterations $T$	30
Elite ratio $\eta$	not explicitly used (leader-only guidance)
Initial mutation scale $\sigma_0$	adaptive: $(1 - t/T)(0.1 + 0.9 r_i)$
Opposition horizon $\tau$	every iteration
Stagnation threshold $K$	not used in this implementation
Rescue ratio $\rho$	0.4
Lévy scale $\alpha$	implicit 1.0
Lévy exponent $\beta$	1.5
Noise bounds $(c_1, c_2)$	$h_{\min} = \max(0.01, 0.3\hat{\sigma}), h_{\max} = \min(0.2, 1.5\hat{\sigma})$
Patch candidates $\mathcal{P}$	$\{3, 4, 5, 6, 7\}$
Search-radius candidates $\mathcal{D}$	$\{3, 4, 5, 6, 7, 8, 9, 10, 11\}$

Table S4 indicates that the proposed ROA variant was configured with a compact search budget of 10 agents and 30 iterations, making it computationally feasible for repeated NLM evaluation. The algorithm relies on leader-based guidance, adaptive rank-dependent mutation, opposition learning at every iteration, and Lévy-flight exploration with exponent 1.5. In addition, the denoising strength parameter  $h$  is searched within noise-adaptive bounds derived from the estimated image noise level, while patch size and search radius are optimized over bounded discrete candidate sets. Overall, these settings show that the method emphasizes a balance between exploration, exploitation, and computational efficiency.

### D: Discussion of Supplementary Tables

Tables S1–S3 provide additional insight into the optimization behavior of the proposed ROA-based NLM framework on the CBS68 dataset. In particular, these results help clarify how the optimized parameters vary across images and which parameters are most sensitive to image content.

Table S1 reports the image-wise optimal hyperparameter configurations identified by ROA. A clear pattern is that the filtering strength  $h$  remains highly concentrated within a narrow range across most images, while the patch size  $p$  also shows strong consistency. In contrast, the search radius  $d$  exhibits greater variation from image to image. This behavior is consistent with the design of the proposed method. Since the feasible range of  $h$  is constrained using the estimated noise scale, the optimizer is encouraged to search within statistically meaningful smoothing levels. As a result, ROA tends to converge toward stable values of  $h$  that are well aligned with the underlying noise magnitude. At the same time, the variation in  $d$  indicates that the extent of non-local aggregation remains strongly dependent on image structure and self-similarity patterns.

Table S2 further confirms this trend through dataset-level statistics. The optimized filtering strength has a very small standard deviation, indicating that ROA consistently identifies a similar smoothing scale across the CBS68 images. Likewise, the patch size shows minimal variation, suggesting that small patches are generally preferred under the experimental setting considered in this study. By contrast, the search radius has a noticeably larger standard deviation, which implies that this parameter is the most image-dependent among the three NLM hyperparameters. This observation supports the view that  $h$  is primarily governed by noise scale, whereas  $d$  must adapt to the spatial redundancy and structural complexity of individual images.

The parameter frequency distribution in Table S3 provides an intuitive summary of this behavior. The patch size is overwhelmingly concentrated at a single value, demonstrating that ROA repeatedly identifies the same local patch scale as optimal for most images. However, the search radius is distributed across several values, with one dominant choice but multiple alternative selections depending on the image. This result suggests that the proposed optimizer does not simply converge to a fixed global configuration; rather, it preserves adaptivity where it is most needed. In particular, ROA appears to learn that the degree of non-local search should vary across images even when the preferred smoothing scale and patch size remain relatively stable.

Overall, the supplementary results reinforce the main conclusion of the paper: the proposed ROA framework performs structured hyperparameter optimization rather than unguided black-box search. The empirical distributions of  $h$ ,  $p$ , and  $d$  show that the noise-aware constraints successfully stabilize the search for filtering strength, while the adaptive population-based search remains flexible enough to adjust the non-local search radius according to image-specific characteristics.

### ***E: Theorem and statistical analysis***

**Theorem S1 (Feasibility Preservation).** Let  $\Omega \subset \mathbb{R}^d$  be a bounded box-constrained domain. If every updated candidate in ROA is projected onto  $\Omega$  by componentwise clipping, then all population members remain in  $\Omega$  for all iterations.

*Proof.* The initial population is sampled from  $\Omega$ . Since each updated candidate is clipped to the lower and upper bounds before evaluation, feasibility is preserved inductively for every iteration.  $\square$

**Theorem S2 (Monotonic Best-So-Far Improvement).** Let  $f_t^*$  denote the best objective value found by ROA up to iteration  $t$ . If a new candidate is accepted only when it improves the objective value of the current solution, then the sequence  $\{f_t^*\}_{t \geq 0}$  is monotonically non-increasing.

*Proof.* At each iteration, candidate replacement is performed only when the new or opposition solution yields a lower objective value. Hence no accepted update can worsen the incumbent best value, and therefore  $f_{t+1}^* \leq f_t^*$  for all  $t$ .  $\square$

**Proposition S1 (Noise-Adaptive Search Interval).** In the proposed implementation, the filtering strength parameter  $h$  is optimized over the interval

$$[\max(0.01, 0.3\hat{\sigma}), \min(0.2, 1.5\hat{\sigma})],$$

where  $\hat{\sigma}$  is the estimated image noise level. Therefore, the search range for  $h$  adapts automatically to the observed corruption level while remaining bounded.

**Proposition S2 (Exploration Capability).** Under stochastic mutation and Lévy-flight perturbation with non-degenerate sampling distributions, ROA has nonzero probability of generating candidates in any positive-measure subregion of the feasible domain. Consequently, the search process retains global exploration capability even under leader-guided updates.