Proceedings of the 8th International Scientific and Practical Conference «Progressive Science and Achievements»
(September 6-8, 2026).
Dallas, USA

No
308

RADIO ENGINEERING, ELECTRONICS AND ELECTRICAL ENGINEERING

How Much Fitness Information is Enough: Fitness Quantization in Metaheuristic Optimization

Cinar Ridvan Firat[1], Lale Timur[2]

[1] Assistant Professor in Engineering Faculty,
Department of Computer Engineering
Batman University; Türkiye

[2] Assistant Professor in Engineering Faculty,
Department of Electrical and Electronics Engineering
Batman University; Türkiye

Abstract. We study how many bits of objective-function information population-based metaheuristics require. A decision-level interface separates the true objective from the fitness visible to the optimizer. Differential Evolution (DE), Particle Swarm Optimization (PSO), and a real-coded Genetic Algorithm (GA) are evaluated on five functions in 3,600 paired runs. Under dynamic comparison-relative quantization, Minimum Sufficient Fitness Resolution (MSFR) ranges from 1 to 6 bits, but one bit is worse than full precision for every algorithm-function pair. Fixed-reference controls perform substantially worse, showing that bit depth and range allocation cannot be separated. Low-resolution outcomes also depend on tie policy.
Keywords: *fitness quantization, metaheuristic optimization, differential evolution, particle swarm optimization, fitness resolution.*

I. INTRODUCTION

Metaheuristics usually receive full numerical objective values even when their decisions may need only coarse distinctions. Reducing this information could simplify sensing, communication, or hardware interfaces, but it may also create plateaus and alter selection pressure. Because optimization performance is problem-dependent [1], the required precision should be measured rather than assumed.

We ask how many fitness bits preserve useful performance. The study contributes: (i) a strict separation between true objective values and decision information; (ii) paired experiments across DE [2], PSO [3], and a real-coded GA grounded in evolutionary search [4], [5]; and (iii) an empirical MSFR defined by a precommitted practical tolerance and universal information-theoretic constant is not claimed in this paper.

Quantization changes not only numerical precision but also the comparisons available to an optimizer. It can create ties, artificial plateaus, and altered selection pressure. Moreover, the effective information provided by a bit depth depends on how its levels are allocated over the objective range [6]. Therefore, the quantizer should be treated as part of the optimization interface rather than merely as a storage format.

II. FITNESS INFORMATION QUANTIZATION

For candidate x, $f(x)$ is the true objective and $q_b(x) = Q_b\big(f(x)\big)$ is the value visible to the optimizer. Within every active comparison pool, true values are normalized as

$$z_i = \frac{(f_i - f_{min})}{(f_{max} - f_{min} + \varepsilon)} \quad (1)$$

where $\varepsilon = 10^{-12}$ prevents division by zero. The normalized value is mapped to at most 2^b levels:

$$q_i^{(b)} = \frac{round\big((2^b - 1)\cdot z_i\big)}{(2^b - 1)} \quad (2)$$

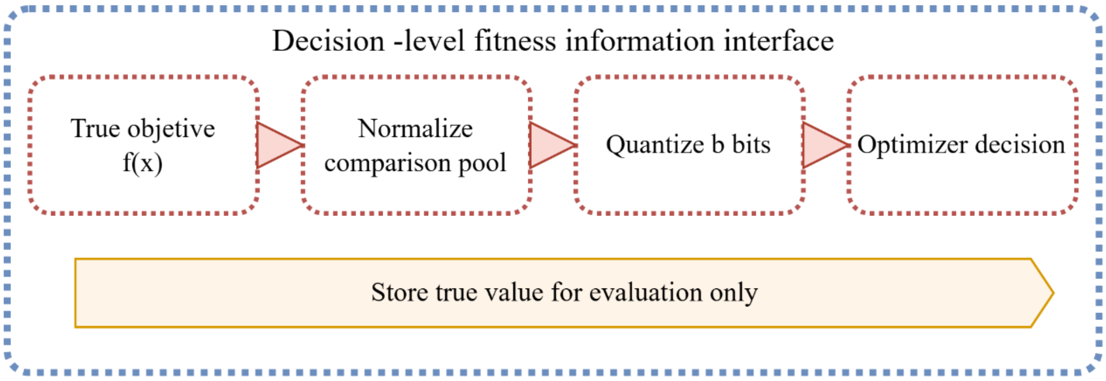

Figure 1
True values drive evaluation;
only visible fitness drives optimizer decisions

Dynamic Quantization (DQ) [6] recomputes the scale for each comparison pool. Historical quantized scores are never compared across iterations: DE jointly quantizes parents and trials; PSO jointly quantizes the incumbent global best and personal-best candidates. Equal visible values retain the incumbent and never use hidden true ordering.

III. EXPERIMENTAL DESIGN

We minimize Sphere, Rosenbrock, Rastrigin, Ackley, and Griewank [7] at dimension 10. Each condition uses population 30, 500 iterations, and exactly 15,030 function evaluations. Full, 16, 8, 6, 4, 3, 2, and 1-bit conditions share 30 seeds. DE [2] uses $F = 0.5$ and $CR = 0.9$. PSO [3] uses inertia 0.7298 and $c_1 = c_2 = 1.49618$. The GA [4] uses binary tournament selection, arithmetic crossover, and Gaussian mutation. Eleven automated tests verify reproducibility, equal budgets, level capacity, Full-mode equivalence, and absence of hidden-fitness leakage.

Table I

Experimental Configuration

Item	Setting	Item	Setting
Algorithms	DE, PSO, real-coded GA	**Functions**	Five; D=10
Population	30	**Budget**	15,030 NFEs
Runs	30 paired seeds	**Information**	Full; 16, 8, 6, 4, 3, 2, 1 bit
Primary model	Dynamic pool-relative	**Tie policy**	Retain incumbent

The primary outcome is $E = |f(x_{returned}) - f*|$, using the true value of the solution returned through visible information. The external best true value ever evaluated is stored only as a diagnostic. Cross-function analysis uses $L = \log_{10}(E + 10^{-12})$. A depth is practically sufficient when it and every tested higher depth have median $E \leq 10^{-8}$, or median L no more than 0.5 above Full; the smallest is b_{MSFR}.

Reduced conditions are paired with Full by seed. Wilcoxon signed-rank [8] tests follow recommended nonparametric practice for evolutionary algorithms; seven depths within each algorithm-function pair receive Holm correction [9]. Paired rank-biserial correlation reports effect size.

IV. RESULTS

The response to precision is neither universal nor strictly linear. One bit is worse than Full for all 15 algorithm-function pairs after correction, with rank-biserial effects from 0.96 to 1.00. Across 105 reduced-versus-Full comparisons, 64 are significant: 63 favor Full. DE-Ackley at

8 bits is the sole significant improvement (median $\Delta L = -0.318$, adjusted $p = 0.0326$, $r_{rb} = -0.544$), an isolated effect rather than evidence for general regularization.

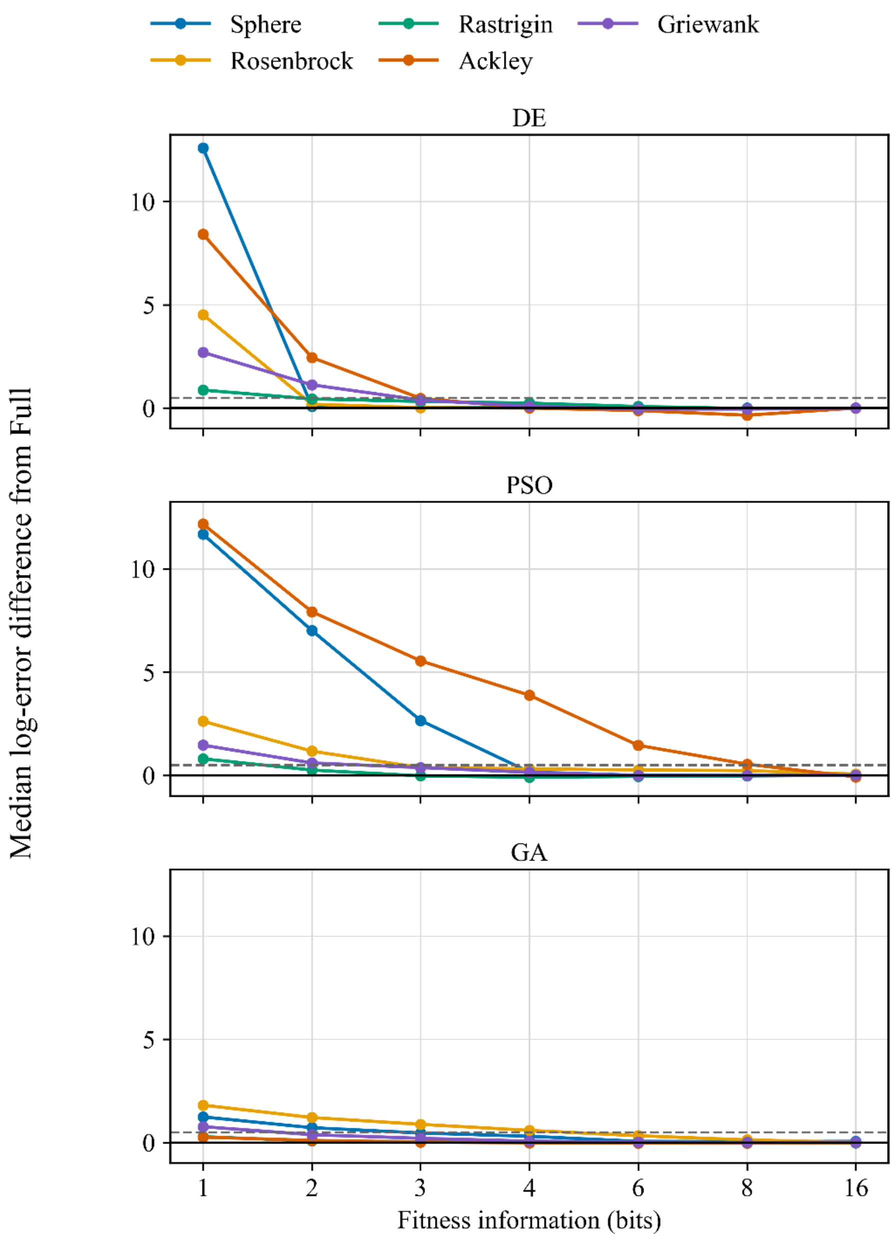

Figure 2
Median log-error difference from Full. Dashed line: 0.5-log practical tolerance; lower is better

A. CONVERGENCE AND INFORMATION SATURATION

Figure 3 shows representative Rastrigin trajectories. Low-bit conditions initially reduce true objective values but plateau earlier as visible ties become frequent. Higher resolutions generally track Full more closely; moderate quantization can still change the trajectory enough to produce pair-specific differences.

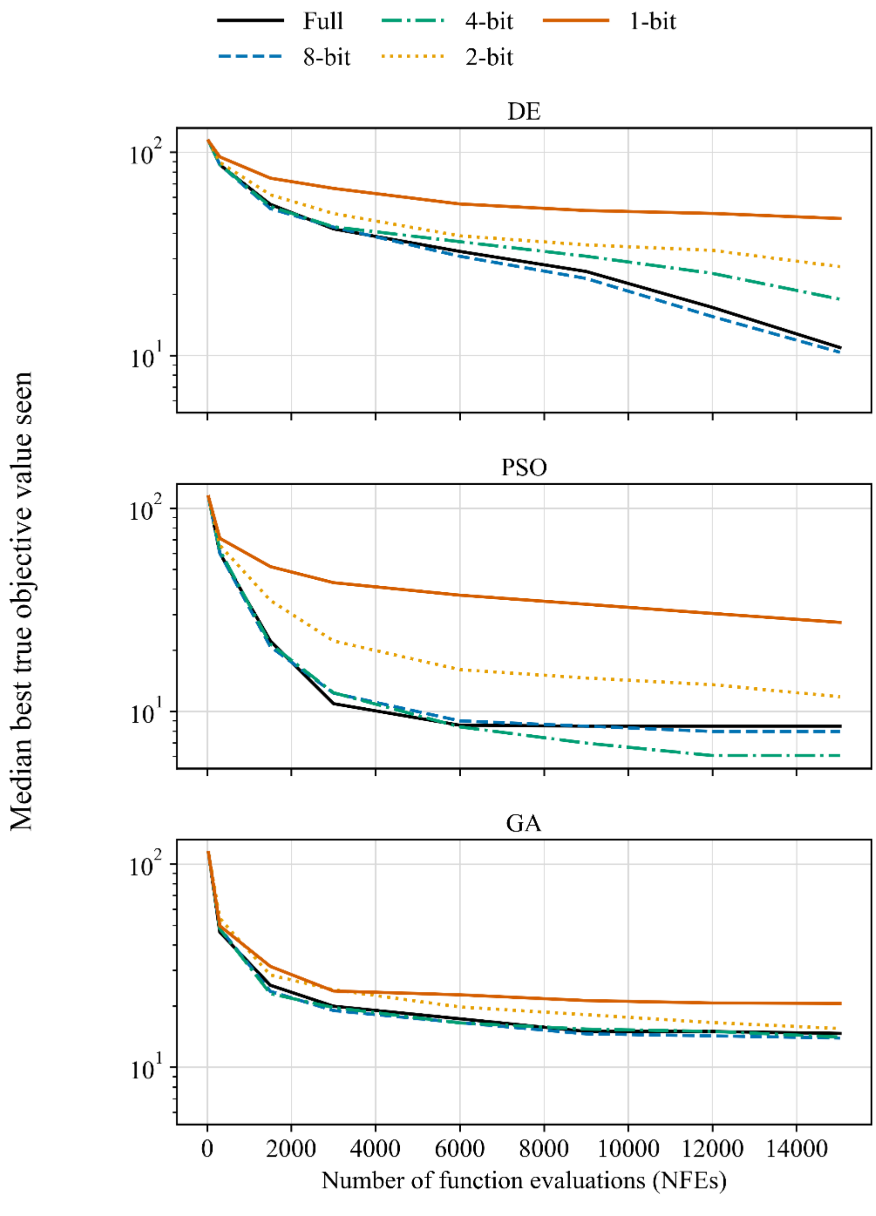

Figure 3
Median best true objective seen versus NFEs on Rastrigin. This oracle diagnostic is not used for optimizer decisions

RADIO ENGINEERING, ELECTRONICS AND ELECTRICAL ENGINEERING

Median one-bit tie rates span 0.368-0.742 for DE, 0.384-0.841 for PSO, and 0.544-0.853 for GA. The returned-solution error can exceed the oracle best-seen error, especially for the non-elitist GA and coarse fitness. Reporting the oracle minimum as the primary outcome would therefore overstate what the optimizer can identify from its available information.

B. MINIMUM SUFFICIENT FITNESS RESOLUTION

MSFR ranges from 1 to 6 bits (Fig. 4). DE needs 2-3 bits across the five functions; PSO needs 2-6; GA needs 1-6. The largest requirements occur for PSO-Ackley and GA-Rosenbrock. GA needs only one bit on Rastrigin and Ackley under the practical rule because its Full baseline is comparatively coarse; this is equivalence to that algorithm's baseline, not high absolute accuracy.

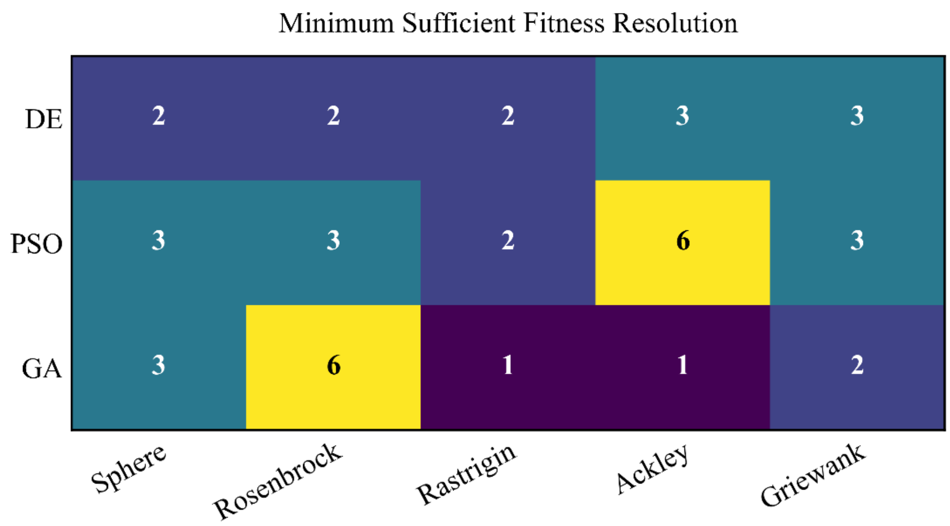

Figure 4
Empirical MSFR under DQ, the stated budget, and the precommitted practical rule

Thus, MSFR is an operational measure of the coarsest fitness resolution that preserves performance relative to each algorithm’s own Full-precision baseline. It does not imply equal absolute accuracy across algorithms or functions, particularly when the baseline itself converges poorly. For this reason, MSFR values should be interpreted together with the absolute returned-solution errors reported in Table II.

Table II

Median True Error: Full / 4-bit / 1-bit

Function	DE	PSO	GA
Sphere	9e-23 / 3.2e-14 / 3.9	9.8e-25 / 4.4e-13 / 0.47	0.12 / 0.25 / 2.2
Rosenbrock	6.3 / 6.2 / 2.1e+05	3.1 / 6.4 / 1.3e+03	39 / 1.5e+02 / 2.5e+03
Rastrigin	11 / 19 / 81	8.5 / 6.8 / 53	41 / 41 / 80
Ackley	1.3e-10 / 1.3e-10 / 0.034	9.2e-12 / 7.7e-08 / 15	4.6 / 4.5 / 8.5
Griewank	0.03 / 0.037 / 15	0.086 / 0.12 / 2.5	1.4 / 1.7 / 8.5

Table II complements the practical MSFR threshold with absolute returned-solution errors. In particular, GA's low MSFR on some functions reflects closeness to its own Full baseline rather than uniformly high accuracy.

C. PRACTICAL AND STATISTICAL INTERPRETATION

MSFR and hypothesis testing answer different questions. MSFR identifies the coarsest depth that remains within a predeclared practical envelope and requires every tested higher depth to satisfy the same rule. The paired tests instead ask whether seed-level log errors differ detectably from Full. A condition can therefore be practically sufficient while remaining statistically distinguishable, particularly on Sphere where both conditions may attain negligible absolute error but differ by several orders of magnitude.

The accuracy floor prevents such near-zero differences from dominating the practical classification. Conversely, MSFR does not imply that all lower depths fail by the same mechanism. At one bit, large tie rates and early plateaus dominate; at intermediate depths, the remaining loss depends on landscape geometry and algorithm memory. Reporting Fig. 4 together with Table II preserves both relative sufficiency and absolute solution quality.

D. QUANTIZER AND TIE-POLICY CONTROLS

Fixed-Reference Quantization (FQ) uses the known optimum as lower bound and the maximum of 10,000 deterministic uniform domain samples as upper bound. The 360-run FQ subset disagrees substantially with DQ: median disadvantages reach 13.3 log units on Sphere. A fixed global range spends levels on early large values and loses distinctions near the optimum. Nominal bit depth therefore does not uniquely specify effective decision information.

A 120-run sensitivity replaces incumbent retention by seeded unbiased tie decisions at 1-2 bits on Rosenbrock and Rastrigin. Effects depend on the algorithm. Random ties are especially harmful for PSO-Rosenbrock, increasing median log-error by 1.10 at one bit and 1.47 at two bits (both paired $p = 0.00195$, $r_{rb} = 1.00$). Tie handling is a material part of the low-resolution condition.

V. DISCUSSION

The central result is conditional saturation. With DQ and incumbent retention, several bits often preserve practically useful performance, but the threshold changes with both the search mechanism and landscape. DE is relatively consistent at 2-3 bits. PSO requires finer resolution on Ackley, where its memory dynamics benefit from continued discrimination near the optimum. GA appears tolerant on some multimodal functions but needs six bits on Rosenbrock.

The controls sharpen the interpretation. First, a bit count is incomplete without a range-allocation rule: DQ continually reallocates levels to the active population, whereas the tested FQ expends most levels far from late-stage candidates. Second, low-bit ties change algorithm dynamics. Incumbent retention creates conservative plateaus; random replacement injects drift and can destabilize PSO memory. Thus a statement such as "four bits are enough" is meaningful only together with the quantizer, tie rule, algorithm, landscape, budget, and practical tolerance.

The isolated DE-Ackley improvement is consistent with altered selection pressure, but one significant benefit among 105 comparisons is insufficient to label quantization a regularizer. The appropriate conclusion is narrower: moderate quantization sometimes changes search without measurable practical loss, while very coarse information reliably degrades the tested optimizers.

A. PRACTICAL IMPLICATIONS

For constrained interfaces, the results suggest separating numerical transport from optimization logic. An adaptive quantizer can provide useful distinctions with few nominal bits, but its reference population and tie behavior must be part of the interface specification. A hardware or distributed implementation that advertises only bit depth is underspecified: two four-bit interfaces can expose very different effective resolution near the current search region.

B. LIMITATIONS

The study uses deterministic analytic functions at D=10 and one NFE budget. DQ is population-relative and therefore transmits ordinal distinctions through an adaptive scale; FQ represents only one deterministic calibration. MSFR depends on the 10^{-8} floor and 0.5-log tolerance. Correction is applied within each algorithm-function family rather than globally over all 105 tests. Finally, the simple non-elitist GA is a mechanism probe, not a claim about all genetic algorithms.

VI. CONCLUSION

This study shows that the fitness resolution required by metaheuristic optimizers is not universal. Under dynamic population-relative quantization, the empirical MSFR ranged from 1 to 6 bits across the tested algorithms and functions. Although one-bit fitness was statistically worse than Full precision in every pair, it remained practically sufficient for two GA cases, highlighting the difference between statistical and practical significance. Fixed-reference controls performed substantially worse, while alternative tie handling changed low-resolution behavior. Therefore, bit depth cannot be interpreted independently of the quantization range, tie policy, algorithm, landscape, and evaluation budget. Overall, moderate adaptive quantization can preserve useful performance with limited fitness information, whereas excessively coarse quantization generally introduces ties and premature plateaus. Future work should test noisy objectives, higher dimensions, and alternative adaptive quantizers.

References:

[1] S. Boyd and L. Vandenberghe, Convex Optimization. Cambridge University Press, 2004. doi: 10.1017/CBO9780511804441.
[2] R. Storn and K. Price, "Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces," Journal of Global Optimization, vol. 11, no. 4, pp. 341–359, Dec. 1997, doi: 10.1023/A:1008202821328.
[3] R. Poli, J. Kennedy, and T. Blackwell, "Particle swarm optimization," Swarm Intelligence, vol. 1, no. 1, pp. 33–57, Oct. 2007, doi: 10.1007/s11721-007-0002-0.
[4] M. Srinivas and L. M. Patnaik, "Genetic algorithms: a survey," Computer (Long. Beach. Calif)., vol. 27, no. 6, pp. 17–26, Jun. 1994, doi: 10.1109/2.294849.

[5] P. Liashchynskyi and P. Liashchynskyi, "Grid Search, Random Search, Genetic Algorithm: A Big Comparison for NAS," Dec. 2019, Accessed: Sep. 09, 2025. [Online]. Available: http://arxiv.org/abs/1912.06059
[6] R. M. Gray and D. L. Neuhoff, "Quantization," IEEE Trans. Inf. Theory, vol. 44, no. 6, pp. 2325-2383, 1998, doi: 10.1109/18.720541.
[7] M. Jamil and X. S. Yang, "A literature survey of benchmark functions for global optimisation problems," International Journal of Mathematical Modelling and Numerical Optimisation, vol. 4, no. 2, p. 150-194, 2013, doi: 10.1504/IJMMNO.2013.055204.
[8] F. Wilcoxon, "Individual Comparisons by Ranking Methods," Biometrics Bulletin, vol. 1, no. 6, p. 80, Dec. 1945, doi: 10.2307/3001968.
[9] S. Holm, "A Simple Sequentially Rejective Multiple Test Procedure," Scandinavian Journal of Statistics, vol. 6, no. 2, pp. 65-70, 1979, [Online]. Available: http://www.jstor.org/stable/4615733