International Journal of Multidisciplinary Research And Studies

Available online at www.ijmras.com
INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH AND STUDIES

E ISSN: 2640 -7272
Volume:09: Issue:02 (2026)

Page no. 9-14

"An Empirical Study on Energy Efficiency and Solution Effectiveness of Evolutionary Algorithms"

Mihel T Den

Department of Computer Science, Technical University of Denmark, Lyngby, Denmark

ABSTRACT

The rapid growth of computational intelligence and large-scale optimization has raised concerns regarding the energy consumption and environmental impact of algorithmic processes. Evolutionary Algorithms (EAs), while widely recognized for their robustness and flexibility in solving complex optimization problems, often require extensive computational resources, which directly translate into increased energy usage. This study presents an empirical comparison of several widely used evolutionary algorithms by jointly analyzing their energy consumption and the quality of solutions they produce. Through standardized experimental setups and benchmark optimization problems, the trade-offs between solution optimality and energy efficiency are investigated. Energy metrics are analyzed alongside convergence behavior and final fitness values to provide a holistic assessment of algorithmic performance. The findings highlight that energy-efficient algorithm design and implementation choices can significantly influence sustainability without severely compromising solution quality. This study contributes to the emerging field of energy-aware evolutionary computation by offering experimental evidence and practical insights for researchers and practitioners.

Keywords: Evolutionary Algorithms, Energy Consumption, Optimization, Solution Quality, Sustainable Computing

INTRODUCTION

The Pyrenean Capercaillie (*Tetrao urogallus aquitanicus*) represents a Nutrition is a Evolutionary Algorithms (EAs) have established themselves as a fundamental pillar of modern metaheuristic optimization. Their utility stems from their ability to navigate complex search spaces that are often nonlinear, multimodal, and high-dimensional—environments where traditional gradient-based methods frequently struggle. By simulating the principles of Darwinian biological evolution, EAs maintain a population of candidate solutions that undergo iterative refinement through stochastic operators: selection, recombination (crossover), and mutation [2–4]. This robust framework has led to widespread adoption across high-impact sectors, including structural engineering design, smart grid energy management, and the orchestration of resource allocation in cloud computing environment.

However, the proliferation of EAs has coincided with a period of intense scrutiny regarding the environmental sustainability of large-scale computation. As optimization tasks grow in complexity, the "computational intensity" of these algorithms—often requiring thousands of generations and millions of fitness evaluations—translates into significant electrical demand. When executed on modern heterogeneous platforms comprising multi-core CPUs and specialized GPUs, this energy consumption becomes a critical non-functional requirement [1,6]. The emerging field of "Green AI" now posits that the carbon footprint of an algorithm is a metric of equal importance to its predictive accuracy or convergence speed .

Despite this shift in perspective, a disjoint persists in current literature. While theoretical studies have optimized the convergence rates of EAs [3,5] and some research has profiled hardware-level energy draw [6], few empirical studies provide a dual-objective analysis of energy expenditure versus solution quality. In energy-sensitive applications—such as the real-time control of HVAC systems or the management of smart home IoT networks—an algorithm that achieves a 1% increase in solution accuracy at the cost of a 50% increase in power consumption may be practically unviable.

This study addresses this critical gap by conducting a rigorous experimental comparison of multiple evolutionary paradigms. We evaluate their performance through a unified framework that measures both optimization fitness and total energy consumption (measured in Joules) across standardized benchmarks. By establishing these benchmarks, this work provides a roadmap for "energy-aware" algorithm selection, ensuring that the next generation of intelligent systems is both high-performing and environmentally responsible.

2. Materials and Methods

2.1 Evolutionary Algorithms Considered

The experimental analysis focuses on a diverse set of representative evolutionary algorithms (EAs) that form the backbone of modern metaheuristic research: Genetic Algorithms (GA), Differential Evolution (DE), and Evolutionary Programming (EP). These were selected not only for their widespread adoption but for their distinct approaches to navigating search landscapes .

Genetic Algorithms are utilized as the primary baseline, employing binary or real-valued representations and relying heavily on crossover (recombination) to explore the search space. In contrast, Differential Evolution is prioritized for its efficacy in continuous domains; it utilizes vector-based mutation strategies that adapt to the objective function's topology, often leading to superior convergence speeds in high-dimensional problems [5]. Finally, Evolutionary Programming is included for its focus on phenotypic evolution and self-adaptive mutation rates, which have historically proven effective in complex energy-related optimization tasks where the relationship between parameters is non-linear .

2.2 Benchmark Problems and Experimental Setup

To ensure a rigorous and reproducible evaluation, we utilized a suite of standardized benchmark functions. These functions were selected to represent varying degrees of difficulty, including unimodal landscapes (to test convergence speed) and multimodal landscapes (to test the algorithm's ability to escape local optima). Challenges such as the Rosenbrock, Rastrigin, and Ackley functions were implemented to simulate different dimensionality and landscape complexities .

The experimental environment was strictly controlled to eliminate external variables. All algorithms were initialized with a population size ($N = 100$) and terminated after a fixed number of fitness evaluations ($FE_{max} = 10,000$) or upon reaching a predefined tolerance.

To mitigate the impact of stochastic noise, each experiment was repeated for 30 independent trials using consistent random seed control. Hardware parity was maintained throughout, using a dedicated workstation with a fixed CPU frequency to prevent dynamic scaling from biasing energy readings .

2.3 Energy Consumption Measurement

Energy consumption was quantified using a system-level methodology that integrates software-based power estimation with hardware registers. We utilized Running Average Power Limit (RAPL) interfaces to capture high-resolution data on CPU and DRAM energy usage. This aligns with the "Green AI" reporting standards, ensuring that we account for both the computational logic of the EA and the memory overhead of maintaining large populations .

Our measurement protocol focused on two temporal scales: Total Energy Consumed (E_{total}) per run and Average Energy per Iteration (E_{iter}). By decoupling these metrics, we can determine if an algorithm is "expensive" due to its complex operators (high E_{iter}) or simply due to a slow convergence rate requiring more cycles to find a solution.

2.4 Evaluation Metrics

The algorithms are evaluated through a bi-objective lens to identify the most sustainable configurations:

1. Solution Quality: This is quantified by the final fitness value reached and the Convergence Rate, which tracks how quickly an algorithm approaches the global optimum.

2. Energy Efficiency: Measured in Joules (J), this metric evaluates the total electrical cost of the optimization process. We also introduce the Energy-Success Ratio, which calculates the energy spent per unit of improvement in fitness value.

This dual-metric approach allows for a "Pareto-style" analysis, identifying algorithms that may provide a 95% optimal solution while consuming only 50% of the energy of more exhaustive methods.

3. Results

The experimental results reveal significant disparities in the intersection of energy consumption and solution quality across the evaluated evolutionary paradigms. While all three algorithms—Genetic Algorithms (GA), Differential Evolution (DE), and Evolutionary Programming (EP)—were capable of reaching the global optima in simpler unimodal landscapes, their efficiency profiles diverged sharply as problem complexity increased.

Differential Evolution emerged as the most high-performing candidate in terms of solution quality. Across high-dimensional benchmarks, DE consistently achieved lower final fitness values and exhibited faster convergence rates, corroborating its reputation for robustness in continuous optimization [5,15]. Interestingly, while DE is often considered computationally intensive due to its vector-based mutation, its ability to "strike" the optimum in fewer generations resulted in the lowest Total Energy Consumed (E_{total}). This suggests that in evolutionary computation, algorithmic "intelligence" (reducing the number of iterations) is a more effective green strategy than merely simplifying the operators.

In contrast, Genetic Algorithms demonstrated a more taxing energy profile. Although GA provided moderate solution quality, it frequently suffered from "stagnation phases" where the population would cease significant improvement but continue to consume power through repetitive crossover and mutation cycles. This prolonged execution time, required to meet the termination criteria, led to a significantly higher carbon footprint per successful optimization run. This observation aligns with earlier findings suggesting that the "tail end" of an optimization process is often the most energy-inefficient phase [6,11].

Evolutionary Programming displayed a unique "stable-but-slow" profile. Its energy consumption per iteration was consistently lower than DE, likely due to the absence of complex recombination operators. However, EP was more prone to becoming trapped in local optima within highly multimodal landscapes, such as the Rastrigin function. Consequently, while EP might be considered "energy-efficient" on a per-second basis, its Energy-Success Ratio was poor in complex scenarios because it consumed power without reliably reaching the desired solution quality [8].

Finally, the results underscored the impact of implementation-level factors. We observed that energy usage was not solely a function of the algorithm's logic but was influenced by system-level behaviors such as memory management and "hysteresis effects"—where the hardware's thermal state from a previous run influenced the power draw of the next [9,10]. This confirms that achieving "Green EA" requires a holistic view of both algorithmic efficiency and software-hardware synergy.

To expand your Discussion into a robust 500-word academic analysis, we will synthesize your findings into broader implications for "Green Computing," explore the integration of hybrid systems, and propose a roadmap for the next generation of energy-conscious EAs.

4. Discussion

The empirical evidence gathered in this study confirms that the "performance" of an Evolutionary Algorithm can no longer be defined through a single lens of solution optimality. Instead, our results reveal a nuanced, multi-dimensional trade-off where energy consumption acts as a significant constraint on algorithmic utility. While Differential Evolution (DE) demonstrated superior precision, its energy efficiency is not an inherent trait but rather a product of its convergence velocity. When implementation overhead or hardware-level inefficiencies are introduced, the "cost-per-fitness-gain" can fluctuate dramatically [5,9]. This suggests that the selection of an EA must move beyond purely mathematical benchmarks to include the operational environment as a primary variable.

4.1 Redefining Performance in Sustainable Computing

As the global research community pivots toward sustainable computing, there is a burgeoning necessity to recalibrate our traditional metrics. Traditionally, EAs are judged on their ability to reach a global optimum within a set number of generations. However, in the context of large-scale infrastructure systems—such as smart grids or decentralized cloud clusters—the "best" solution is the one that minimizes the total resource footprint of the search process [14–16].

We propose that "Energy-Aware Evolutionary Algorithms" should be categorized not just by their convergence accuracy, but by their Energy-to-Accuracy (ETA) ratio. This perspective is particularly critical for battery-operated edge devices in smart environments, where an exhaustive search for a 0.1% improvement in fitness could result in a catastrophic depletion of local power reserves.

4.2 Interdisciplinary Synergy and Sustainability

The discussion of energy efficiency naturally extends into the realm of interdisciplinary integration. Recent advancements have shown that combining EAs with fuzzy logic and intelligent decision-support systems can significantly dampen computational intensity. In applications such as water resource management and environmental monitoring, fuzzy-based selection mechanisms can prune the search space more effectively than stochastic mutation alone [17–21]. By integrating these "soft computing" techniques, researchers can achieve a "best-of-both-worlds" scenario: the robust global search of an EA paired with the resource-saving heuristics of fuzzy systems. This hybrid approach represents a vital pathway for balancing high-level computational performance with the constraints of environmental sustainability.

4.3 Towards Adaptive, Energy-Sensitive Frameworks

Looking forward, the evolution of these algorithms must move toward dynamic adaptivity. Current EAs operate with fixed parameters (population size, mutation rate) that do not account for the energy state of the host system. Future research should investigate mechanisms that dynamically balance exploration and exploitation based on real-time energy telemetry.

One radical yet promising direction involves incorporating energy consumption directly into the fitness function. By treating Joules as a "cost" within a multi-objective optimization framework (e.g., using an NSGA-II approach), algorithms could evolve solutions that are both technically effective and energy-efficient. Such "Green Fitness Functions" would represent a paradigm shift in algorithm design, ensuring that the next generation of artificial intelligence is not only intelligent but also environmentally responsible.

5. Conclusions

This study has provided a rigorous empirical investigation into the intersection of evolutionary computation and environmental sustainability. By jointly evaluating energy consumption and solution quality, we have challenged the traditional paradigm that views algorithmic performance solely through the lens of fitness optimization. Our findings demonstrate that the relationship between computational accuracy and electrical cost is non-linear and highly sensitive to the chosen evolutionary framework.

The experimental results underscore a fundamental truth in metaheuristic research: no single algorithm universally outperforms others across all metrics. While Differential Evolution (DE) proved to be the most "energy-efficient" in terms of reaching high-quality solutions quickly, its advantage is contingent upon the complexity of the search landscape. In contrast, Genetic Algorithms (GA) and Evolutionary Programming (EP), while occasionally more stable in their per-iteration energy draw, demonstrated that prolonged convergence times can lead to a disproportionately high total carbon footprint. This reinforces the necessity for "context-aware" algorithm selection, where the choice of a solver is dictated as much by the energy constraints of the hardware as by the mathematical requirements of the problem.

Furthermore, this work has highlighted that "greenness" in optimization is not merely a software trait. The influence of implementation choices—ranging from the selection of programming languages to the management of system-level resources—plays a nontrivial role in an algorithm's energy profile. By documenting these variations, this study contributes a foundational dataset to the growing field of energy-aware evolutionary computation. We have shown that by prioritizing algorithms with superior "Energy-to-Accuracy" ratios, practitioners can significantly reduce the environmental impact of intelligent systems without sacrificing technical rigor.

Looking ahead, this research serves as a call to action for the optimization community. The transition toward sustainable AI requires a shift from "brute-force" search methods toward adaptive, energy-sensitive frameworks. Future efforts should focus on developing "Green Fitness Functions" that internalize energy costs as an optimization objective, ensuring that the evolutionary process itself remains environmentally responsible. As optimization tasks continue to scale across smart grids, cloud infrastructures, and edge computing, the integration of sustainability into algorithm design will be the defining hallmark of the next generation of computational intelligence.

References

1. Henderson, P.; Hu, J.; Romoff, J.; Brunskill, E.; Jurafsky, D.; Pineau, J. Towards the systematic reporting of the energy and carbon footprints of machine learning. J. Mach. Learn. Res. 2020, 21, 10039–10081.

2. Slowik, A.; Kwasnicka, H. Evolutionary algorithms and their applications to engineering problems. Neural Comput. Appl. 2020.

3. Doerr, B.; Neumann, F. A survey on recent progress in the theory of evolutionary algorithms for discrete optimization. ACM Trans. Evol. Learn. Optim. 2021.

4. Bartz-Beielstein, T.; Branke, J.; Mehnen, J.; Mersmann, O. Evolutionary algorithms. Wiley Interdiscip. Rev. Data Min. Knowl. Discov. 2014 .

5. Mayer, D.G.; Kinghorn, B.; Archer, A.A. Differential evolution—an easy and efficient evolutionary algorithm for model optimisation. Agric. Syst. 2005.

6. de Vega, F.F.; et al. A Cross-Platform Assessment of Energy Consumption in Evolutionary Algorithms. In PPSN; Springer, 2016.

7. Aquino-Brítez, S.; et al. Towards an Energy Consumption Index for Deep Learning Models. Sensors 2025.

8. Maryam, K.; Sardaraz, M.; Tahir, M. Evolutionary algorithms in cloud computing from the perspective of energy consumption. ICET 2018.

9. Merelo-Guervós, J.J.; García-Valdez, M. How evolutionary algorithms consume energy depending on the language. LION 2024.

10. Cotta, C.; Martínez-Cruz, J. Evaluating implementation choices on energy consumption in EAs. EvoStar.

11. Merelo, J.; et al. Benchmarking languages for evolutionary algorithms. EvoApplications 2016.

12. Merelo-Guervós, J.J.; et al. Measuring Energy Consumption of BBOB Fitness Functions. EvoStar.

13. Fong, K.F.; et al. HVAC system optimization by evolutionary programming. Energy Build. 2006.

14. Lee, W.S.; et al. Optimal chiller loading by differential evolution. Energy Build. 2011.

15. Essiet, I.O.; et al. Optimized energy consumption model for smart home. Energy 2019.