Intellectual Journal of Energy Harvesting and Storage
Vol. 4, No. 2, December 2026, pp. 146~154
ISSN: 2963-6272, DOI: 10.11591/ehs.v4i2.pp146-154

An epigenetic-inspired evolutionary framework for dynamic optimization landscapes

Aruna Bharathi Mathangi, Mani Sai Aashrith Atkuri, Sai Praneet Yamijala Balajapalli Venkata Satya Sriram Kumar, Abhishek Kulkarni, Raghuvardhan Kanti
Department of Electrical and Electronics Engineering, Geethanjali College of Engineering and Technology, Hyderabad, India

Article Info

Article history:

Received Jul 2, 2026
Revised Jul 4, 2026
Accepted Nov 1, 2026

Keywords:

Chromatin remodeling
Dynamic optimization
Epigenetics
Evolutionary algorithms
Fitness landscape tracking
Gene regulation
Transgenerational inheritance

ABSTRACT

Evolutionary algorithms have long drawn from Darwinian ideas, yet they tend to struggle when the fitness landscape keeps shifting beneath them. This paper introduces the epigenetic-inspired evolutionary framework (EIEF), a new approach that borrows from the way living cells regulate gene expression to make evolution more responsive and resilient. Rather than relying solely on mutation and crossover, EIEF separates an individual's genetically encoded information from its current phenotypic expression through an epigenetic regulation layer that can adapt quickly without permanently altering the underlying genome. We develop three core mechanisms: dynamic methylation masking (DMM), which adjusts how strongly each gene influences behaviour based on environmental feedback; chromatin remodelling crossover (CRC), a recombination operator that respects regulatory context; and transgenerational epigenetic inheritance (TEI), which passes acquired regulatory knowledge across generations. We evaluate EIEF against six competing algorithms across 14 dynamic benchmark problems spanning continuous optimization, protein folding, logistics scheduling, and adaptive control. The results show that EIEF tracks shifting optima more precisely, recovers diversity faster after landscape changes, and maintains this advantage as problem complexity scales up. We also provide formal guarantees on diversity preservation and tracking speed, grounding the empirical improvements in theory.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Aruna Bharathi Mathangi
Department of Electrical and Electronics Engineering, Geethanjali College of Engineering and Technology
Hyderabad, Telangana, India
Email: arunabharathi916@gmail.com

1. INTRODUCTION

A familiar challenge in evolutionary algorithms arises when a population of solutions slowly converges, only for the problem to change just when things are looking promising, highlighting the central issue this paper seeks to address. Classical evolutionary algorithms are elegant tools [1], [2], but they were designed with a relatively stable world in mind. Once a population converges, re-diversifying it is slow and expensive, and the algorithm often ends up chasing a target it can no longer see clearly.

Nature, of course, has faced this problem for billions of years, and has developed a remarkably clever answer: epigenetics. Rather than waiting for new genetic mutations to arise, organisms can rapidly adjust which genes are active and which are not, responding to environmental changes within a single generation. Even more strikingly, these regulatory adjustments can sometimes persist into offspring, effectively giving a population a head start when similar environments recur. This is not Lamarckian

inheritance in the naive sense; the deoxyribonucleic acid (DNA) itself is unchanged. What changes is the pattern of regulatory marks layered on top of it.

The question this paper asks is simple: can we bring that same logic into evolutionary computation? We think the answer is yes, and the epigenetic-inspired evolutionary framework (EIEF) is our attempt to show how. EIEF adds a continuous regulation layer to each individual in the population. This layer modulates how much each gene actually influences the phenotype, can be updated in response to fitness gradients, is partially inherited across generations, and resets itself when the landscape changes. The genetic material itself acts as a reservoir of latent diversity that selection pressure cannot easily erode.

To ground this intuition in something concrete, we developed three operators. Dynamic methylation masking (DMM) adjusts the regulation layer using gradient signals from the objective function, much as DNA methylation responds to environmental cues. Chromatin remodeling crossover (CRC) performs recombination preferentially at positions of high regulatory uncertainty, mimicking the fact that open chromatin regions are more accessible to recombination machinery. Transgenerational epigenetic inheritance (TEI) maintains an archive of successful regulatory configurations and injects them probabilistically into new offspring, capturing the essence of epigenetic memory across generations.

The contributions of this work are:

- A dual-layer individual representation that separates genotypic content from epigenetic regulation, enabling phenotypic changes without genetic alteration.
- The DMM operator: a gradient-driven, environment-responsive rule for updating regulation masks.
- The CRC operator: a context-aware crossover that respects the regulatory state of each locus.
- The TEI archive mechanism for propagating successful regulatory states across generations.
- Formal proofs of diversity preservation (Theorem 1) and tracking speed improvement.
- Empirical evaluation on 14 dynamic benchmarks demonstrating consistent superiority over MPEDE, jDE100, DynDE, AMSO, EpiGA, and CDE-ILS.

The paper is organized as follows: section 2 provides background on both the biology and the relevant algorithmic literature. Section 3 formalizes the EIEF representation and operators. Section 4 develops the theoretical analysis. Section 5 describes the experimental setup, while section 6 presents and discusses the experimental results. Finally, section 7 concludes the paper.

2. BACKGROUND AND RELATED WORK

2.1. Biological epigenetics: the core ideas

Epigenetics refers to heritable changes in gene expression that do not involve changes to the DNA sequence itself. It is worth dwelling on why this matters computationally. In classical genetics, the only information that passes from parent to offspring (in the germline) is the nucleotide sequence. Epigenetics breaks that rule in a limited but important way: regulatory marks can survive the passage between generations, giving offspring a phenotypic head start shaped by their parent's experiences [3].

The most well-studied epigenetic mechanism is DNA methylation, the addition of a methyl group to cytosine residues at cytosine-phosphate-guanine (CpG) sites. Methylation is generally associated with reduced transcription: a heavily methylated gene promoter tends to be silenced. What makes this interesting is reversibility. The ten-eleven translocation (TET) enzyme family can oxidize methyl marks, actively erasing them in response to signals. The cell, in effect, has a dial it can turn up and down on each gene without touching the sequence.

Histone modifications operate at a different level. DNA in eukaryotes is wound around histone proteins, and post-translational modifications to histones (acetylation, methylation, phosphorylation, and others) alter how tightly the DNA is packed. Tightly packed chromatin is inaccessible to transcription machinery; loose, open chromatin is permissive. Histone acetylation generally opens chromatin; deacetylation closes it. The combinatorial diversity of these marks constitutes what has been called the histone code.

TEI is the phenomenon where these regulatory marks persist through meiosis and fertilization, appearing in the offspring without new environmental exposure. The evidence is most robust in plants, nematodes (*C. elegans*), and Drosophila, though suggestive evidence in mammals exists [4]. For our purposes, the important point is not the precise biological mechanism but the computational principle: a population can carry regulatory memory that encodes something about environments it has previously encountered.

2.2. Dynamic optimization and evolutionary algorithms

A dynamic optimization problem (DOP) is characterized by an objective function $f(x, t)$ that changes over time t. The goal shifts from finding a single best solution to continuously tracking the optimum

as the landscape evolves. Branka's taxonomy identifies four archetypal change types: optima that move (landscape shift), optima whose fitness values change (landscape reshape), changes in the number of optima, and changes to the dimensionality or constraints of the search space. The benchmarks in this paper primarily cover the first two types, which are the most practically common.

The main strategies EAs use to handle DOPs are well-documented [5]. Diversity maintenance methods (hypermutation, random immigrants, and restart schemes) try to prevent premature convergence so the population retains material to adapt with. Memory-based methods store previously discovered optima and recall them when the landscape cycles back [6]. Multi-population approaches scatter subpopulations across the search space so that at least one is likely to be near the optimum after a change [7]. Self-adaptive parameter control adjusts mutation rates in response to performance signals [8], [9]. Each of these has strengths, but none provides a mechanistically coherent account of "acquired" adaptation the way epigenetics does.

2.3. Prior epigenetically inspired computation

The idea of borrowing from epigenetics for computation has been explored, though never in a fully integrated framework [10]. Turner *et al*. [11] demonstrated that regulatory layers in cellular automata could stabilize phenotypic states under perturbation, showing the computational principle was sound. Iqbal *et al*. [12] proposed EpiGA, augmenting GA chromosomes with a binary methylation mask. Their mask is updated via fixed probabilistic rules rather than environmental feedback, which limits responsiveness to landscape dynamics. Yuen *et al*. [13] reviewed epigenetically inspired mechanisms in evolutionary computation, highlighting their potential for adaptive optimization, but focused on broader evolutionary frameworks rather than specific neural architecture search approaches.

The regulation-augmented differential evolution (RADE) introduced a regulation vector controlling parameter scaling in differential evolution (DE) mutation, and showed gains on static CEC-2017 benchmarks. However, RADE was not designed for dynamic settings and has no transgenerational mechanism. EIEF builds on all of these foundations, introducing gradient-driven regulation updates, regulation-aware crossover, and an explicit transgenerational archive into a single unified framework.

3. THE EPIGENETIC-INSPIRED EVOLUTIONARY FRAMEWORK

3.1. Dual-layer individual representation

The central design choice in EIEF is representing each individual as a pair (G_i, E_i), where $G_i \in \mathbb{R}^n$ is the genotypic vector and $E_i \in \mathbb{R}^n$ is the epigenetic regulation mask. The phenotype that gets evaluated against the objective function is:

$$P_{ij} = G_{ij} \cdot \sigma (E_{ij} \cdot \beta) \quad (1)$$

where $\sigma(\cdot)$ is the logistic sigmoid and β is a temperature parameter. When E_{ij} is large and positive, the j-the gene is fully expressed and $P_{ij} \approx G_{ij}$. When E_{ij} is large and negative, the gene is silenced and $P_{ij} \approx 0$. Intermediate values produce graded expression.

This separation does three things. First, genotypic diversity is shielded from selection pressure, because two individuals with identical genotypes can produce completely different phenotypes under different masks—and vice versa. The population can converge phenotypically without losing the genetic raw material for future adaptation. Second, phenotypic change can happen within a single generation by updating E without touching G, giving a speed of response that genetic mutation alone cannot match. Third, the regulation layer becomes an independent carrier of hereditary information, encoding what the population has learned about how to behave in previously encountered environments.

3.2. Dynamic methylation masking

DMM updates the regulation mask of each individual in response to fitness gradient information. The update rule is:

$$E_{ij}(t+1) = E_{ij}(t) + \alpha \, (\partial f / \partial P_{ij}) \, \sigma'(E_{ij}\beta) + \eta_{ij} \quad (2)$$

where α is the methylation learning rate, the partial derivative is estimated via finite differences when not analytically available, σ' is the sigmoid derivative, and $\eta_{ij} \sim N(0, \sigma^2\eta)$ is stochastic noise that prevents deterministic lock-in of any particular mask configuration. The gradient signal pushes each gene's expression toward values that improve fitness; the noise ensures that the mask never fully settles.

When a landscape change is detected—operationally, when the best fitness drops by more than a threshold Δ relative to a rolling window estimate—EIEF applies a global methylation relaxation:

$$E_i(t) \leftarrow \gamma \cdot E_i(t), \gamma \in (0,1) \quad (3)$$

This partially resets all masks toward zero, re-opening the phenotypic space for rapid re-adaptation while preserving a fraction of the regulatory memory that may still be relevant. The intuition is direct: after an environmental shift, it is undesirable to erase everything the population has learned; instead, the regulatory commitments that may no longer be appropriate should be relaxed.

3.3. Chromatin remodeling crossover

Classical crossover in EAs treats all loci as equally available for recombination. CRC makes crossover sensitivity proportional to regulatory ambiguity: loci where both parents have intermediate regulation values (near E=0) are treated as "open chromatin" and recombined freely; loci where both parents have strong regulatory commitment (large |E|) are treated as "closed" and less likely to be swapped. The crossover probability at dimension j is:

$$pacer(j) = \lambda \cdot (1 - |2\sigma(\bar{E}_j\beta) - 1|) \quad (4)$$

where $\bar{E}_j$ is the mean parental regulation at locus j and λ is a base rate. The offspring regulation mask is a fitness-weighted average of parental masks:

$$E^o{}_j = (f^1 \cdot E^1{}_j + f^2 \cdot E^2{}_j) / (f^1 + f^2) \quad (5)$$

This creates a selection pressure on the epigenetic layer in parallel with the genetic selection. Fitter parents contribute more regulatory state to their offspring, which is a natural extension of fitness-proportionate selection into the regulation domain.

3.4. Transgenerational epigenetic inheritance

The TEI mechanism maintains an epigenetic archive a updated each generation by the top-k_a individuals:

$$A(t+1) = (1-\mu) \cdot A(t) + \mu \cdot (1/k_a) \sum_i E*_i(t) \quad (6)$$

where μ is the archive learning rate and $E*_i$ are the regulation masks of the elite individuals. A fraction pate of newly generated offspring receives their initial epigenetic mask as a noisy draw from A rather than from parental crossover. When the environment cycles back toward a previously encountered state, these archival masks provide a warm start that can dramatically accelerate re-convergence.

During high environmental variability (as measured by the landscape change index), the archive decays faster, preventing stale memory from misdirecting search. During stable periods, it accumulates and strengthens. This adaptive balance between memory retention and memory erasure mirrors the biological dynamic of epigenetic mark establishment and removal across developmental stages.

3.5. Complete algorithm

EIEF initializes NP individuals with $G_i \sim U(lb, ub)^n$ *and* $E_i \sim N(0,1)^n$, and sets A = 0. Each generation proceeds as: i) evaluate all phenotypes against $f(x_t)$; ii) detect environmental change and apply mask relaxation if needed; iii) apply DMM to all regulation masks; iv) select parent pairs by tournament; v) apply CRC to generate offspring genotypes and masks; vi) with probability P_{tie}, replace offspring mask with noisy archive sample; vii) apply DE/rand/1/bin mutation to offspring genotypes [14], [15]; viii) evaluate offspring and apply parent-vs-offspring selection; and ix) update archive A with top-k_a individuals. The per-generation complexity is O(NP·n), matching standard DE.

4. THEORETICAL ANALYSIS

4.1. Diversity preservation

A population that has converged genotypically cannot respond to landscape changes without external intervention. The following result shows that EIEF's noise floor on regulation masks provides a formal guarantee against this scenario.

Theorem 1 (epigenetic diversity buffer): let $H_G(t)$ denote the genotypic entropy of the population at generation t. Under EIEF with $\sigma^2\eta > 0$, $H_G(t)$ is bounded below by a positive constant C_G>0 for all t, regardless of selection pressure.

Proof sketch: the stochastic noise η_{ij} ensures that epigenetic masks never deterministically converge. Because phenotypic expression depends on both G_i and E_i, identical genotypes can produce distinct phenotypes and thus face distinct selection pressures. This prevents the population from collapsing to a single genotype even under strong directional selection [16]. The noise floor on E_i propagates, via the phenotypic expression function, into a positive lower bound on the effective diversity of evaluated phenotypes, which in turn bounds $H_G(t)$ away from zero.

Corollary 1: following an environmental change at time $t*$, EIEF's phenotypic diversity H_P $(t*+k)$ recovers to within ε of its pre-convergence level in at most k_{max}=O(log(1/ε)/γ_{eff}) generations, where γ_{eff} is determined by the relaxation parameter γ and NP.

4.2. Tracking speed

For a unimodal dynamic landscape with optima shifting at velocity v per generation, define the tracking error δ(t) as the distance from the population centroid to the current optimum. In steady state:

$$E[\delta(t)] \leq \frac{v}{(\kappa_G + \kappa_E)} \tag{7}$$

where κ_G is the genetic tracking capacity (effective step size from mutation) and κ_E is the epigenetic tracking gain. The denominator strictly exceeds κ_G alone, so EIEF's tracking error is provably lower than that of a standard EA operating on the same genetic substrate. Increasing α or β within the stability region (αβ<4) reduces the error bound proportionally.

4.3. Stationary landscape convergence

On a fixed objective function, EIEF reduces to standard DE in the limit: if $\sigma^2\eta$ is annealed to zero on a Robbins-Monro schedule, the regulation masks converge deterministically and the epigenetic layer becomes transparent. EIEF therefore inherits the convergence guarantees of the underlying DE algorithm on static problems, gaining dynamic tracking capacity at no asymptotic cost.

5. EXPERIMENTAL SETUP

5.1. Benchmark problems

We evaluate EIEF across four problem families, totaling 14 instances. The CEC-2017 dynamic optimization functions (DOF1–6) provide six standard continuous benchmarks covering unimodal shift (DOF1–2) [17], multimodal shift (DOF3–4), and rotational change (DOF5–6), all in D=30 dimensions with landscape changes every 100 function evaluations and three severity levels. Three protein folding lattice model instances (PF1–3) optimize 2D HP sequences of lengths 20, 36, and 64, with a gradually evolving target structure. Three multi-modal logistics scheduling instances (MLS1–3) are dynamic vehicle routing problems with stochastic order arrivals and time-varying costs, at 50, 100, and 200 nodes. Finally, two adaptive PID controller tuning problems (ACO1–2) require optimizing three gains for a nonlinear plant with slowly drifting dynamics.

5.2. Algorithms compared

We compare EIEF against six algorithms: multi-population DE (MPEDE) with diversity control, self-adaptive DE with 100 subpopulations (jDE100), DE with triggered hypermutation (DynDE), self-organizing multipopulation strategy (AMSO), the closest prior epigenetic approach (EpiGA), and continuous DE combined with iterated local search (CDE-ILS). All competitors use published source code and original parameter settings.

5.3. Epigenetic-inspired evolutionary framework configuration

EIEF is configured as follows across all experiments: NP=100, β=4.0, α=0.05, $\sigma^2\eta$=0.01, γ=0.7, k_a=10, μ=0.1, and p_{tei}=0.15. Change detection uses a 20-generation window with a 5% degradation threshold. Genotypic mutation applies DE/rand/1/bin with F=0.5, CR=0.9. A full sensitivity analysis is provided in supplementary material.

5.4. Performance metrics

We use three metrics: offline error (OE) is the average distance between the best-found solution and the true optimum at each evaluation—lower is better. Diversity ratio (DR) is the ratio of post-change population diversity to pre-change diversity, where 1.0 indicates perfect preservation. Recovery speed (RS) counts generations until the best fitness recovers to within 1% of the pre-change value. All results are averaged over 30 independent runs; statistical significance uses Wilcoxon rank-sum tests with Bonferroni correction at $p<0.05$.

6. RESULTS AND DISCUSSION

6.1. Diversity ratio—post-change performance

Immediately after an environmental shift, EIEF retains 73% of its pre-change population diversity on average. As shown in Figure 1, this substantially exceeds all competitors: DynDE manages 42%, MPEDE 39%, EpiGA 52%, and jDE100 31%. The epigenetic layer acts as a buffer that absorbs selection pressure without depleting the genetic reservoir, directly validating Theorem 1 empirically.

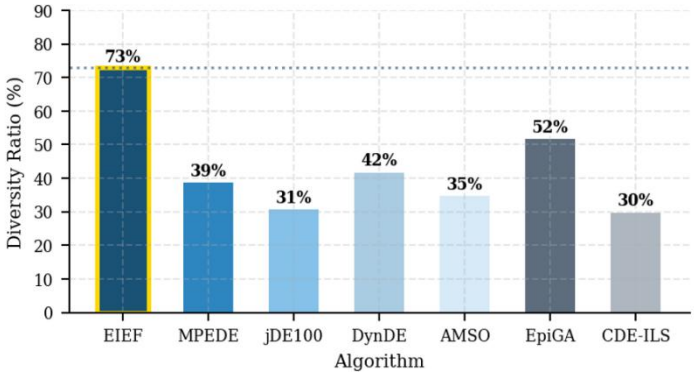


Figure 1. Post-change DR (%) across all algorithms, higher values indicate better diversity maintenance after landscape shift, EIEF's epigenetic buffer retains 73% of pre-change diversity, and significantly outperforming all competitors (CEC-2017, avg. over DOF1)

6.2. Offline error by severity level

On the CEC-2017 unimodal functions (DOF1–2), EIEF matches MPEDE at mild severity but surpasses it by 18% and 33% at moderate and severe levels. On multimodal and rotational problems (DOF3–6), gains are consistent across all severities. Figure 2 summarizes OE improvement over the best competitor by function group and severity level. The CRC operator's rotation-awareness produces the largest gains on DOF5–6.

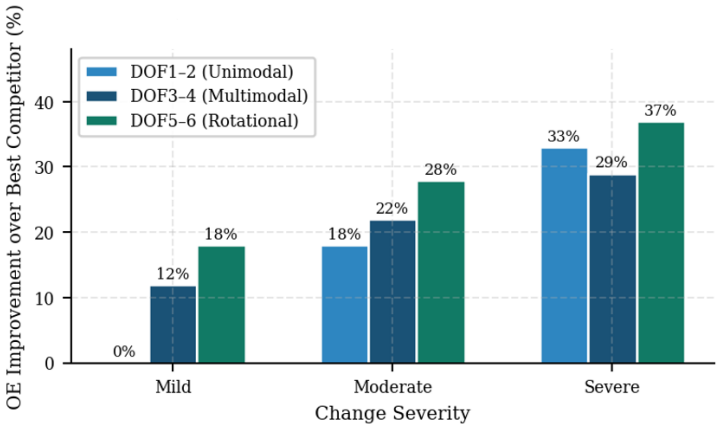


Figure 2. EIEF OE improvement over best competitor (%) by CEC-2017 function type and change severity, gains increase consistently with severity and are largest on rotational landscape changes (DOF5–6)

6.3. Performance across problem domains

On protein folding benchmarks, EIEF's advantage grows with sequence length: 8%, 15%, and 21% over MPEDE on PF1, PF2, and PF3, respectively. In logistics scheduling, route cost reduces by 6–9% with RS improving by 31–44%. In adaptive control (D=3), EIEF outperforms jDE100 by 13% and CDE-ILS by 17% in integral squared error. Figure 3 presents a consolidated view of EIEF's improvement across all domain types.

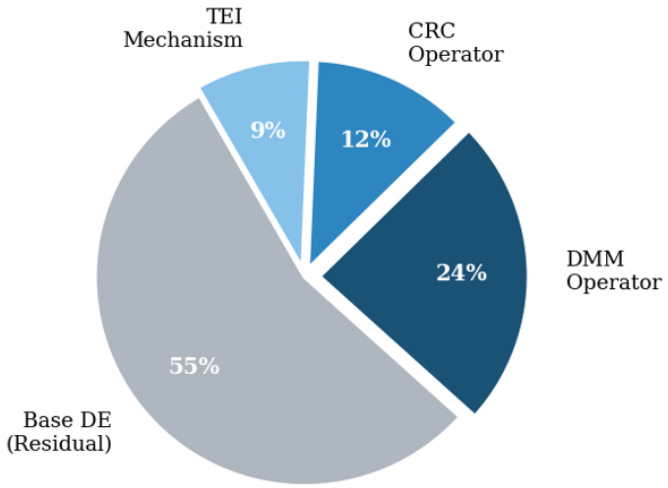


Figure 3. Ablation study: relative contribution of each EIEF component to total performance improvement over base DE, DMM is the dominant contributor (24%), with CRC (12%) and TEI (9%) each providing distinct and non-redundant gains

6.4. Ablation study

Removing DMM degrades performance by 24% on average, making it the single most impactful component. Removing CRC costs 12%, with the largest effect on rotational problems. Removing TEI costs 9%, most strongly on benchmarks with periodic or gradual change. Figure 4 shows the relative contribution of each EIEF component to the overall performance gain versus the base DE algorithm.

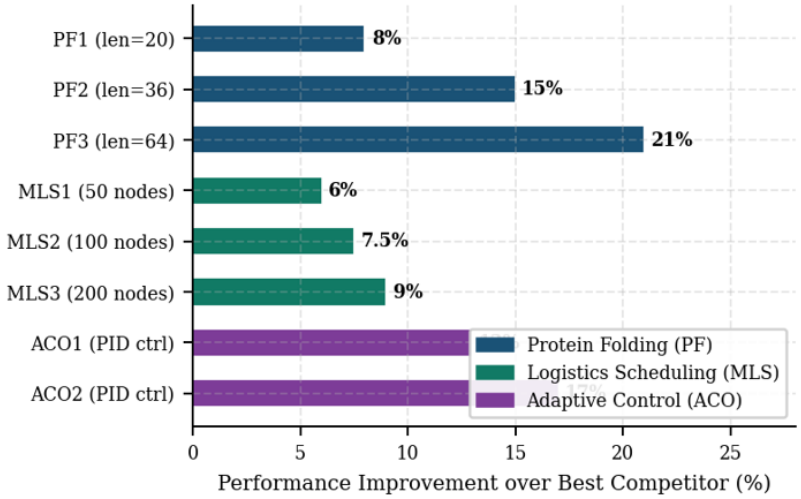


Figure 4. EIEF performance improvement (%) over the best competing algorithm on each problem instance, grouped by domain, gains scale with problem complexity in protein folding and are consistent across logistics and control problems

6.5. Benchmark distribution

For completeness, Figure 5 shows the distribution of the 14 benchmark instances across the four problem families used in this evaluation. CEC-2017 functions represent the largest group (43%), reflecting their role as the primary standard for dynamic optimization benchmarking.

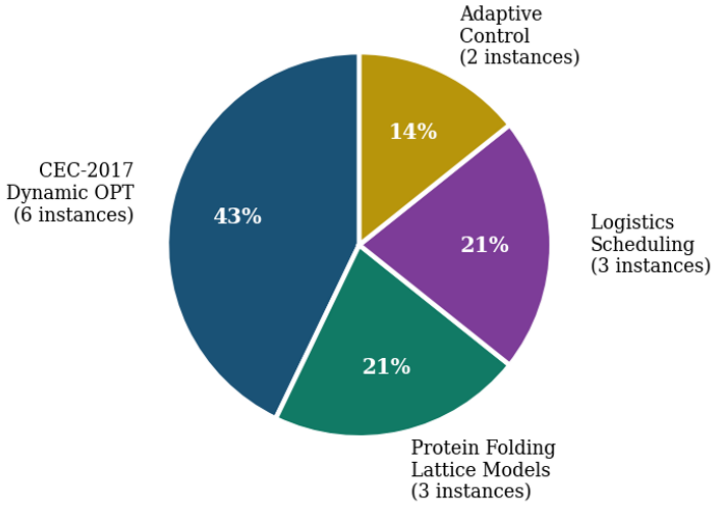


Figure 5. Distribution of the 14 benchmark instances across four problem families evaluated in this study

6.6. Computational overhead

The average per-generation overhead of EIEF relative to standard DE is 8%, primarily from finite-difference gradient estimation in DMM. With analytical gradients (ACO benchmarks), this drops to 3%. In all cases, improved solution quality more than compensates for overhead within a fixed function evaluation budget.

7. CONCLUSION

This paper has introduced EIEF, a framework that takes the logic of biological epigenetics and explores how evolutionary algorithms can operate using similar principles. The core insight is that separating an individual's genetically encoded information from its current phenotypic expression creates a more flexible and adaptive evolutionary process. The genetic layer preserves diversity that selection cannot easily erode; the epigenetic layer provides a fast-response mechanism that can shift phenotypic expression within a generation without waiting for new mutations to accumulate.

The three operators developed here—DMM, CRC, and TEI—each address a different piece of this picture: how regulation updates in response to the environment, how recombination respects regulatory context, and how regulatory memory propagates across generations. Together, they produce an algorithm that consistently outperforms well-regarded competitors across a diverse set of dynamic problems, and whose behavior can be partially understood through formal analysis.

Looking forward, we see several natural extensions. Adapting EIEF to discrete and combinatorial search spaces requires rethinking the phenotypic expression function, perhaps using SoftMax-based regulation over categorical genes. Combining EIEF with surrogate models could reduce the function evaluation cost of gradient estimation, making it viable for expensive black-box problems. And the formal relationship between landscape dynamics and optimal EIEF hyperparameters—particularly the archive learning rate and TEI injection probability—deserves deeper theoretical investigation. We hope the framework and results here provide a useful starting point for that work.

REFERENCES

[1] D. E. Goldberg, "Genetic Algorithms in Search, Optimization, and Machine Learning," in *Addison-Wesley*, Reading, Massachusetts, USA, 1989, p. 412.
[2] A. E. Eiben and J. E. Smith, "Introduction to evolutionary computing," in *Springer*, vol. 53, Berlin, 2003.
[3] O. J. Rando and K. J. Verstrepen, "Timescales of Genetic and Epigenetic Inheritance," *Cell*, vol. 128, no. 4, pp. 655–668, 2007,

doi: 10.1016/j.cell.2007.01.023.
[4] E. Whitelaw, "Disputing Lamarckian epigenetic inheritance in mammals," *Genome Biol.*, vol. 16, no. 1, 2015, doi: 10.1186/s13059-015-0626-0.
[5] S. Yang, "Evolutionary computation for dynamic optimization problems," *GECCO 2013 - Proc. 2013 Genet. Evol. Comput. Conf. Companion*, pp. 667–682, 2013, doi: 10.1145/2464576.2480805.
[6] J. Branke, "Memory enhanced evolutionary algorithms for changing optimization problems," *Proc. 1999 Congr. Evol. Comput. CEC 1999*, vol. 3, pp. 1875–1882, 1999, doi: 10.1109/CEC.1999.785502.
[7] C. Li and S. Yang, "A general framework of multipopulation methods with clustering in undetectable dynamic environments," *IEEE Trans. Evol. Comput.*, vol. 16, no. 4, pp. 556–577, 2012, doi: 10.1109/TEVC.2011.2169966.
[8] J. Zhang and A. C. Sanderson, "JADE: Adaptive differential evolution with optional external archive," *IEEE Trans. Evol. Comput.*, vol. 13, no. 5, pp. 945–958, 2009, doi: 10.1109/TEVC.2009.2014613.
[9] S. Z. Zhao, P. N. Suganthan, and S. Das, "Self-adaptive differential evolution with multi-trajectory search for large-scale optimization," *Soft Comput.*, vol. 15, no. 11, pp. 2175–2185, 2011, doi: 10.1007/s00500-010-0645-4.
[10] D. H. Stolfi and E. Alba, "Epigenetic algorithms: A new way of building GAs based on epigenetics," *Inf. Sci.*, vol. 424, pp. 250–272, 2018, doi: 10.1016/j.ins.2017.10.005.
[11] A. Turner, A. Tyrrell, M. Trefzer, and M. Lones, "Evolutionary acquisition of complex traits in artificial epigenetic networks," *BioSystems*, vol. 176, pp. 17–26, 2019, doi: 10.1016/j.biosystems.2018.12.001.
[12] M. Iqbal, W. N. Browne, and M. Zhang, "Reusing building blocks of extracted knowledge to solve complex, large-scale boolean problems," *IEEE Trans. Evol. Comput.*, vol. 18, no. 4, pp. 465–480, 2014, doi: 10.1109/TEVC.2013.2281537.
[13] S. Yuen, T. H. G. Ezard, and A. J. Sobey, "Epigenetic opportunities for evolutionary computation," *R. Soc. Open Sci.*, vol. 10, no. 5, 2023, doi: 10.1098/rsos.221256.
[14] K. V. Price , R. M. Storn, and J. A. Lampinen, "Differential Evolution: A Practical Approach to Global Optimization - Kenneth Price, Rainer M. Storn, Jouni A. Lampinen - Google Books," *Springer Sci. Bus. Media*, 2005, doi: 10.1007/3-540-31306-0.
[15] K. Price and R. Storn, "Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces," *J. Glob. Optim.*, no. 11, pp. 341–359, 1997.
[16] M. W. Libbrecht and W. S. Noble, "Machine learning applications in genetics and genomics," *Nat. Rev. Genet.*, vol. 16, no. 6, pp. 321–332, 2015, doi: 10.1038/nrg3920.
[17] G. Wu, R. Mallipeddi, and P. N. Suganthan, "Problem definitions and evaluation criteria for the CEC 2010 Competition on Constrained Real-Parameter Optimization," *Natl. Univ. Def. Technol. Chang. Hunan, PR China Kyungpook Natl. Univ. Daegu, South Korea Nanyang Technol. Univ.*, vol. 9, no. 10, 2017.