



Revealing temporal sequence patterns in constrained multiobjective optimization

Miao Rong, Yunjie Wang, Chen Peng

Citation: Rong, M.; Wang, Y.; Peng, C. Revealing temporal sequence patterns in constrained multiobjective optimization. *Intell. Robot.* 2026, 6(2), 311-4. <https://dx.doi.org/10.20517/ir.2026.16>

Received: 13 May 2026
First Decision: 5 Jun 2026
Revised: 12 Jun 2026
Accepted: 12 Jun 2026
Published: 18 Jun 2026

Academic Editor:
Simon Yang
Copy Editor:
Pei-Yun Wang
Production Editor:
Pei-Yun Wang

1. INTRODUCTION

Constrained multi-objective optimization problems (CMOPs) are ubiquitous in real-world engineering applications, ranging from Internet of Things resource scheduling and autonomous vehicle path planning to analog circuit sizing and short-term crude oil scheduling. The central challenge in solving these problems lies in simultaneously balancing convergence, diversity, and feasibility - three mutually constraining objectives that must be carefully managed throughout the evolutionary process.

In recent years, constrained multi-objective evolutionary algorithms (CMOEA) have made steady progress through the adaptive design of constraint-handling techniques (CHTs) and genetic operators^[1,2]. However, most existing approaches overlook a key fact: the requirements for constraint-handling strategies change dynamically across generations. Early generations may demand strong exploration to traverse infeasible regions, whereas later generations require exploitation power to converge onto the constrained Pareto front.

In a recent study published in *IEEE Transactions on Evolutionary Computation*, Peng *et al.* propose a deep reinforcement learning framework that conceptualizes the selection of CHTs and genetic operators as a temporal sequence, offering a principled way to handle this generation-varying adaptation problem^[3].

The temporal-sequence framework also resonates with intelligent robotics. Constrained multi-objective optimization underpins robotics tasks such as autonomous path planning and safety-constrained decision making. Just as robotic systems must dynamically reconfigure strategies in changing environments, CMOEA need to adapt constraint-handling selections across generations. This synergy suggests that learning-based temporal sequence selection could inspire self-configuring optimization engines for robotic systems.



School of Mechatronic Engineering and Automation, Shanghai University, Shanghai 200444, China.

Correspondence to: Prof. Chen Peng, School of Mechatronic Engineering and Automation, Shanghai University, Shanghai 200444, China. E-mail: c.peng@shu.edu.cn

2. CORE CONTRIBUTIONS

Unlike prior adaptive CMOEAs that typically treat constraint-handling and genetic operator selection as isolated, generation-independent decisions, this work unifies both selections into a single temporal sequence, thereby viewing the evolutionary process itself as a structured sequential decision problem. The most significant conceptual contribution of this work is the formalization of the “temporal sequence of constrained handling selection”. By modeling each generation’s choice of CHT and genetic operator as a discrete time step, the authors establish a unified framework in which most existing CMOEAs^[2] - whether multiple-population, multi-stage, penalty-based, or learning-based - can be viewed as special cases with predetermined or degenerate selection sequences. This perspective elevates algorithm design from ad hoc engineering to a structured sequential decision problem.

To discover systematic patterns, the authors design a deep Q-network that learns from historical feedback, with a state representation capturing feasibility, convergence, and diversity, and an action space of nine CHT-operator combinations. A two-phase adaptive reward function balances exploration and exploitation^[3,4].

The experimental validation is thorough. The proposed CMOEA-TS is compared against nine peer algorithms on 37 benchmark instances from three test suites (MW, LIR-CMOP, and DAS-CMOP). Statistical results consistently demonstrate CMOEA-TS’s superiority across 37 benchmark instances from the MW, LIR-CMOP, and DAS-CMOP test suites with respect to the reported inverted generational distance (IGD) and hypervolume (HV) metrics^[3]. Ablation studies further verify that: (i) learning systematic patterns from the sequence significantly outperforms random selections; (ii) comprehensive selection of both CHTs and genetic operators yields better adaptability than using either component in isolation^[5]; and (iii) the proposed credit assignment function more effectively balances objective optimization and constraint satisfaction than simpler reward schemes.

3. DISCUSSION AND FUTURE DIRECTIONS

Conceptually, the temporal-sequence perspective elevates CMOEA design from ad hoc engineering to a structured sequential decision problem, revealing learnable, systematic patterns across generations and opening avenues for data-driven autonomous configuration. Nevertheless, several practical limitations may hinder its deployment in resource-constrained real-world scenarios. Despite its compelling results, several aspects warrant further investigation. The deep Q-network’s training relies on accumulated experience samples in a replay buffer. During early generations, when samples are scarce and potentially biased, Q-value estimates may exhibit high variance^[3]. Although the ϵ -greedy strategy mitigates cold-start issues, it remains unclear how the framework performs under extremely small population sizes or very short evolutionary horizons. A deeper concern is that early-generation samples are not merely scarce but often heavily biased toward random, suboptimal regions of the search space, causing the Q-network to prematurely converge on spurious correlations between state features and action values. Consequently, even ϵ -greedy exploration may struggle to recover from these early biased estimates, suggesting that warm-starting the Q-network via pre-trained models or transfer learning across similar CMOPs could be a valuable remedy.

Second, the action space is discretized into nine fixed combinations of CHTs and genetic operators. While this discretization guarantees interpretability, it also constrains the algorithm’s expressive power. Future work could explore continuous action spaces or finer-grained parameterization^[3,4] - for instance, using policy gradient methods to directly optimize crossover probabilities, mutation step sizes, or penalty coefficients as continuous variables. Such an extension, however, would introduce non-trivial challenges: the dimensionality of the action space would increase significantly, and policy-gradient estimators would suffer

from high variance when interacting with the noisy fitness landscape of evolutionary algorithms. Yet the potential payoff is substantial - continuous parameterization would eliminate the information loss inherent in manual discretization and could discover novel hybrid strategies, such as dynamically blending crossover rates or annealing penalty coefficients in ways that fixed discrete actions cannot approximate.

Third, the computational overhead of the deep Q-network merits attention in resource-constrained scenarios. The theoretical time complexity is^[4]

$$O(T_{\max}(mN^3 + ((l-1)h^2 + (r+o)h)\omega b))$$

where T_{\max} is the maximum number of generations, m is the objective dimension, N is the population size, l is the number of hidden layers, h is the average nodes per hidden layer, r and o are the input and output layer nodes, ω is training epochs per batch, and b is the batch size. which, although remaining in the same order as the evolutionary component, introduces non-negligible constants through neural network forward passes and backpropagation. Extending this framework to large-scale CMOPs with high-dimensional decision variables^[3] or dynamic environments where constraints and objectives change over time^[2] presents both practical challenges and exciting research opportunities.

All in all, this work makes a thought-provoking contribution by recasting CMOEA design as a sequential decision-making problem solvable through deep reinforcement learning^[3]. It reminds the community that evolutionary algorithm behavior across generations is not an unstructured random walk but contains learnable systematic patterns. This insight opens the door to more data-driven, autonomous algorithm configuration paradigms in evolutionary computation.

DECLARATIONS

Authors' contributions

Made substantial contributions to conception and writing of the manuscript: Rong, M.; Wang, Y.; Chen, P.
Performed manuscript revision and provided supervision: Chen, P.

Availability of data and materials

Not applicable.

AI and AI-assisted tools statement

Not applicable.

Financial support and sponsorship

This work was supported in part by the National Natural Science Foundation of China (Grant Nos. 62533016 and 62573279).

Conflicts of interest

Chen, P. is a section editor of Embodied Intelligence and Embedded Systems of the journal *Intelligence & Robotics*, but was not involved in any steps of editorial processing, notably including reviewer selection, manuscript handling, and decision making, while the other authors have declared that they have no conflicts of interest.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Copyright

© The Author(s) 2026.

REFERENCES

1. Tian, Y.; Zhang, Y.; Su, X.; Zhang, X.; Tan, K. C.; Jin, Y. Balancing objective optimization and constraint satisfaction in constrained evolutionary multiobjective optimization. *IEEE. Trans. Cybern.* **2022**, *52*, 9559-72. DOI
2. Fan, Z.; Li, W.; Cai, X.; et al. Push and pull search for solving constrained multi-objective optimization problems. *Swarm. Evol. Comput.* **2019**, *44*, 665-79. DOI
3. Peng, C.; Yan, S.; Zhong, C.; Huang, Q.; Wu, C.; Huang, H. Learning-based temporal sequence of constrained handling selection for constrained multi-objective evolutionary optimization. *IEEE. Trans. Evol. Comput.* **2026**, *30*, 1123-36. DOI
4. Zou, M.; Gong, D.; Wang, Y.; Ye, X.; Zeng, B.; Meng, F. Process knowledge-guided autonomous evolutionary optimization for constrained multiobjective problems. *IEEE. Trans. Evol. Comput.* **2024**, *28*, 193-207. DOI
5. Li, K.; Chen, R.; Fu, G.; Yao, X. Two-archive evolutionary algorithm for constrained multiobjective optimization. *IEEE. Trans. Evol. Comput.* **2019**, *23*, 303-15. DOI

Disclaimer/Publisher's Note: All statements, opinions, and data contained in this publication are solely those of the individual author(s) and contributor(s) and do not necessarily reflect those of OAE and/or the editor(s). OAE and/or the editor(s) disclaim any responsibility for harm to persons or property resulting from the use of any ideas, methods, instructions, or products mentioned in the content.



© The Author(s) 2026. Open Access This article is licensed under a Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, sharing, adaptation, distribution and reproduction in any medium or format, for any purpose, even commercially, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.