A Heterogeneous Ant Colony Framework with Semantic Manifold Learning for Dynamic Ore-Flow Blending

Hao Tang
tanghao.ec@gmail.com
School of Automation, China University of Geosciences
Wuhan, China
Hubei Key Laboratory of Advanced Control and Intelligent Automation for Complex Systems
Wuhan, China
Engineering Research Center of Intelligent Technology for Geo-Exploration, Ministry of Education
Wuhan, China

Changhe Li*, Guoyu Chen, Shoufei Han
changhe.lw@gmail.com
chenguoyumail@163.com
hanshoufei@gmail.com
State Key Laboratory of Digital Intelligent Technology for Unmanned Coal Mining, Anhui University of Science & Technology
Huainan, China
School of Artificial Intelligence, Anhui University of Science & Technology
Hefei, China

Michalis Mavrovouniotis
michalis.mavrovouniotis@eratosthenes.org.cy
ERATOSTHENES Centre of Excellence, Cyprus University of Technology
Limassol, Cyprus

Miqing Li
m.li.8@bham.ac.uk
School of Computer Science, University of Birmingham
Birmingham, U.K

Abstract

Ore-flow blending in open pit mining is a complex real-time scheduling decision problem constrained by high dimensional spatiotemporal dynamics and stochastic uncertainties. Traditional methods often struggle to reconcile the need for long term global stability with the requirement for sub second response times to dynamic events. To address this, we propose a heterogeneous ant colony optimization framework with semantic manifold learning. The framework bridges the gap between static offline learning and dynamic online adaptation. In the offline phase, we introduce a semantic manifold learning mechanism combining genetic programming with quality-diversity optimization. By constructing a behaviorally diverse archive of elite heuristics, this method provides the online system with a "warm start" and diverse initial perspectives, significantly enhancing real-time responsiveness. In the online phase, a heterogeneous ant colony system is deployed where different ant sub populations are coupled with distinct elite rules from the offline archive. This architecture enables parallel exploration of the solution space from multiple semantic perspectives, granting the system intrinsic robustness and high adaptability. Extensive experiments on real-world mine demonstrate that this method significantly outperforms state-of-the-art techniques in ore flow stability control and exhibits superior robustness and scalability across varying road network topologies, equipment configurations, and dynamic scenarios.

CCS Concepts

• **Computing methodologies** → **Genetic programming**; **Search methodologies**; *Planning and scheduling*; • **Applied computing** → *Industry and manufacturing*.

Keywords

Ore-flow blending, genetic programming, quality-diversity optimization, ant colony optimization

ACM Reference Format:
Hao Tang, Changhe Li*, Guoyu Chen, Shoufei Han, Michalis Mavrovouniotis, and Miqing Li. 2026. A Heterogeneous Ant Colony Framework with Semantic Manifold Learning for Dynamic Ore-Flow Blending. In *Proceedings of The Genetic and Evolutionary Computation Conference 2026 (GECCO '26).* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3795095.3805181

1 Introduction

Ore-flow blending in open pit mining is a stringent dynamic scheduling problem aimed at stabilizing processing plant feed grades. The goal is to orchestrate truck scheduling [10, 18] amidst continuous travel time fluctuations and discrete disruptions such as equipment failures. The fundamental challenge lies in balancing global optimality with rapid responsiveness [8, 17]. Minimizing grade deviations requires deep prospective analysis, whereas stochastic disturbances demand immediate reactions, presenting a conflict that traditional optimization methods struggle to reconcile.

To address dynamic scheduling challenges, various model driven solutions utilize explicit mathematical models. Multi-stage stochastic programming [5] optimizes expected values via scenario trees , while robust optimization [12] conservatively targets worst case

* Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License.
GECCO '26, San Jose, Costa Rica
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2487-9/2026/07
https://doi.org/10.1145/3795095.3805181

scenarios to guarantee feasibility. However, these approaches heavily suffer from the curse of dimensionality. As random variables and decision horizons grow, their computational complexity increases exponentially. Rule-based methods dominate industrial scheduling due to their low computational cost and simple implementation. Priority strategies [22] allocate tasks using predefined heuristics when triggering events occur. However, these heuristics often employ greedy logic that prioritizes immediate local gains over global objectives, making single rules inadequate for complex dynamic scenarios. Meanwhile, rescheduling strategies focus on trigger mechanisms [9] to balance revision costs and stability. Ultimately, rescheduling remains a remedial action after disturbances and lacks true global adaptability.

Swarm intelligence methods exploit population cooperation to tackle complex optimization problems [1]. Ant colony optimization (ACO) is widely applied in routing and scheduling because of its constructive search and positive feedback mechanisms [16]. For dynamic scheduling, ACO incorporates various adaptation strategies. Pheromone evaporation [19] eliminates outdated information after environmental changes to balance past experience with rapid adaptation. Furthermore, immigration schemes [15] and memory pools [20] maintain population diversity by introducing new or historically optimal individuals, while multiple population mechanisms [25] partition subpopulations to balance exploration and exploitation. Despite its strong global search capability, the initial search of ACO in dynamic environments is often blind and leads to slow convergence. Introducing heuristic information accelerates this process, but manually designed heuristics remain too simple to capture the characteristics of complex dynamic environments.

Genetic programming hyper heuristics (GPHH) [6] solve dynamic scheduling by evolving rules rather than direct schedules. To handle diverse scenarios, multitask GP [28] transfers genetic material , and ensemble GP [4] utilizes voting for robustness. Despite fast offline training responses, GPHH faces two severe bottlenecks. First, semantic cloning causes structurally distinct rule trees to exhibit identical behavior. This renders genotype based diversity ineffective and leads to premature convergence. While quality diversity methods like MAP Elites [24] attempt to build diverse libraries, defining an adaptive semantic space remains difficult. Second, GPHH lacks online adaptability. Relying on static offline training paradigms, fixed rules fail and performance collapses during unforeseen distributional shifts or abrupt disturbances.

In summary, existing methodologies for real-time scheduling in dynamic and uncertain environments are trapped in a dilemma of speed versus quality and offline learning versus online adaptation. ACO offers online search ability but is constrained by cold start and computational efficiency; GPHH provides rapid response but is limited by semantic redundancy and offline staticism. To break this impasse, this paper proposes a heterogeneous ant colony dynamic scheduling framework (GP-ACO), tailored to the specific dynamic and uncertain scheduling problem of ore-flow blending. The central motivation of this study is to exploit abundant offline computational resources to thoroughly mine the semantic manifold structure of the rule space and to construct a high quality and behaviorally diverse elite rule library; these insights are then efficiently transferred to the online stage to empower the swarm intelligence algorithm with warm start capability and parallel multi-strategy exploration.

The main contributions of this paper are as follows.

- We abandon the traditional genotype centric evolutionary perspective and propose rule evolution in a behavioral semantic space. By defining a semantic distance based on state snapshots and embedding rules into a low dimensional semantic manifold, we employ a CVT-MAP-Elites [23] algorithm to build a high quality and behaviorally exclusive elite rule manifold.
- we construct a heterogeneous adaptive ant colony online scheduling mechanism: the offline evolved heterogeneous rule library serves as an online strategy repository for ACO, and we design heterogeneous ants for parallel exploration so that the algorithm can search simultaneously along multiple differentiated semantic directions and avoid local optima caused by single heuristics.
- We validate the effectiveness of the proposed method in realistic ore flow blending scenarios, demonstrating its superior performance in controlling grade fluctuation and in coping with equipment failures and changes in the road network.

2 Ore-Flow Blending Problem

This section fundamentally reframes ore-flow blending as a real-time scheduling decision problem under dynamic and uncertain environments. In this framework, the truck scheduling system acts as a dynamic actuator. The arrival order and time intervals of trucks carrying high-grade or low-grade ore collectively determine the instantaneous fluctuations of the ore flow grade.

2.1 Optimization Objective

To quantify and control the temporal stability of the ore flow, we employ a sliding time window mechanism rather than a shift-based average. The truck flow entering the crushing station is grouped into sliding windows based on arrival times. The optimization objective is to minimize the root mean squared error (RMSE) of the ore flow grade across all time windows. This ensures the grade within each window remains close to the target grade.

Unlike static models, the formulation here treats the ore grade as a cumulative outcome of a stochastic process over time. The truck flow entering the crushing station is grouped into windows based on actual arrival times, which are subject to stochastic travel delays. Let $\mathcal{T}$ denote the continuous operation horizon. The optimization objective is to minimize the cumulative grade deviation across all time windows, formulated as an error-accumulation process over time:

$$\min f(x) = \sqrt{\frac{1}{|O|}\sum_{o\in O}\frac{1}{|Q_o|}\sum_{q\in Q_o}\mathbb{E}_{\xi(\cdot)}\left[\left(\hat{u}^g_{o,q}\big(\xi(t)\big) - G^b\right)^2\right]} \quad (1)$$

where O represents the set of crushing stations and Q_o denotes the set of time windows at station o. The term G^b is the target grade. Crucially, $\hat{u}^g_{o,q}(\xi(t))$ represents the realized average ore grade within the q-th window, which is modeled as a function of the time variant stochastic vector $\xi(t)$. Specifically, $\xi(t) = (\tau_k(t), h_k(t))$ encapsulates the primary spatiotemporal uncertainties of the mining system. Here, $\tau_k(t)$ denotes the stochastic travel time of truck k which fluctuates due to weather and road conditions, and $h_k(t)$ indicates the discrete health state of truck k at time t, capturing

sudden equipment failures. By minimizing the expectation $\mathbb{E}_{\xi(\cdot)}$, the objective seeks a robust schedule that suppresses grade oscillations despite these continuous and discrete disturbances.

This model is also subject to a series of complex mine operational constraints including truck routes, equipment capacity, queueing rules, and shovel excavation sequences (detailed in Supplementary Material 1).

2.2 Challenges of Dynamic and Uncertain Environments

The ore-flow blending problem in open-pit mining is characterized by a deep coupling of system dynamics and stochastic uncertainties. Unlike static resource allocation [3], the mining system acts as a time-varying actuator where the truck scheduling system must contend with high-dimensional spatiotemporal constraints.

The core difficulty in this formulation lies in minimizing grade deviations within a highly stochastic environment. The mining system is subject to high-dimensional spatiotemporal uncertainties. These include operational variability where non-stationary variables such as truck travel times and shovel loading durations fluctuate continuously due to weather conditions. Furthermore, the system must contend with sudden discrete disturbance events like equipment failures and road blockages which instantaneously alter the road network topology. These uncertainties create a complex dependency where a scheduling decision made at the current moment has a cumulative and prospective impact on the grade stability of multiple future time windows. A myopic or greedy strategy that ignores these dynamic factors will inevitably lead to severe grade oscillations when unforeseen disturbances occur.

3 Algorithm Framework

3.1 Motivation and Framework Overview

To address the inherent conflict between global stability and real-time agility, we propose the GP-ACO (Genetic Programming-based Heterogeneous Ant Colony Optimization) framework. The design of this framework is explicitly driven by the limitations of existing single-method approaches: the "cold start" and slow convergence of traditional ACO in dynamic environments, and the "semantic redundancy" and lack of online adaptability of standard GPHH. The framework bridges these gaps through a two-phase architecture:

Phase I: Overcoming Semantic Cloning via Manifold Learning (Offline). Standard genetic programming (GP) often suffer from premature convergence where the population collapses into structurally similar but behaviorally identical rules (semantic cloning). To address this, we abandon the genotype-centric perspective. We employ a Quality-Diversity (QD) algorithm, specifically CVT-MAP-Elites, to evolve a "Semantic Manifold." By mapping scheduling rules into a low-dimensional metric space based on their actual decision-making behaviors, we actively pressure the algorithm to discover a diverse archive of elite heuristics. This ensures that the resulting rule library covers distinct behavioral niches, providing a comprehensive toolkit for unknown future scenarios.

Phase II: Addressing Blind Search via Heterogeneous Agents (Online). In dynamic scheduling, traditional ACO typically starts with a homogeneous colony that searches blindly, leading to slow response times. To overcome this, we design a Heterogeneous Ant Colony System. Instead of starting from scratch, the colony is initialized with the diverse elite rules learned in Phase I. We partition the colony into sub-populations, where each group is bound to a distinct behavioral "instinct" (elite rule). This heterogeneity allows the system to conduct a parallel exploration of the solution space from multiple semantic perspectives simultaneously. By integrating these specific heuristic signals with pheromone feedback, the system achieves a "warm start" and intrinsic robustness, capable of rapidly switching strategies when dynamic disturbances occur.

3.2 Offline Semantic Manifold Learning

Standard GP tends to converge prematurely to a single optimum. To overcome this, we utilize a semantic geometric approach to map the behavior of rules into a queryable metric space.

3.2.1 Genetic Programming Hyper-heuristics. We employ a tree-based GPHH [27] approach to explore the heuristic space. GPHH evolves a scoring function $S_R(s, a)$ within a greedy template. This function automatically generates a heuristic rule R that guides the low-level constructive heuristics in building solutions. Unlike traditional static heuristics, these evolved hyper-heuristics demonstrate strong transferability and can be directly applied to unseen instances. Each individual in the GPHH population is encoded as a syntax tree. The internal nodes are drawn from a function set $\mathcal{F}$, while the leaf nodes are drawn from a terminal set $\mathcal{T}$.

Function Set ($\mathcal{F}$). We define a comprehensive set of arithmetic and logical operators to ensure sufficient expressive power:

$$\mathcal{F} = \{+, -, \times, /, \min, \max, \sin, \text{if}, \text{sigmoid}, \tanh\} \tag{2}$$

Terminal Set ($\mathcal{T}$). The terminals are designed to capture the current environmental state and action features. The set is defined as $\mathcal{T} = \{t^c_*, \ldots, h^a_m\}$, containing the following specific features:

- Truck & System State: Capacity of the current truck (t^c_*); target ore grade of the beneficiation plant (G^b); average ore grade across all loading points (G^l).
- Loading Point Features: Average ore grade at loading point i (G^l_i); current remaining average ore grade across all loading points (g^r) and at specific point i (g^r_i); current remaining ore proportion across all points (p^l) and at point i (p^l_i); average ore grade of currently excavated ore (g^l).
- Transportation Features: Ore grade of the current truck heading to point m ($t^g_{*,m}$); average ore grade of other trucks heading to point m ($t^g_{-,m}$).
- Unloading & Objective Features: Average ore grade of the most recent sliding window across all unloading points (u^g) and at specific point j (u^g_j); RMSE of ore flow grades across all unloading points (u^e) and at point j (u^e_j); expected grade error in the recent window after dispatching the current truck to point m (h^a_m).

For a given problem instance D_m with an initial state s_0, a complete solution trajectory $\tau_R(D_m) = (a^m_0, \ldots, a^m_{T_m})$ is constructed using the heuristic rule R decoded from the GP tree. At each decision step t, the action a_t is selected to maximize the evolved scoring

function:

$$a_t = \arg\max_{a \in \mathcal{A}(s_t)} S_R(s_t, a) \tag{3}$$

where $\mathcal{A}(s_t)$ represents the set of available actions.

Since the scales of objective values vary across the training set $\mathcal{D} = \{D_m\}_{m=1}^{M}$, we employ the relative optimality gap to normalize the fitness evaluation. We maintain a record of the historical best objective value, denoted as f_m^*, found so far for each instance D_m during the training process. This reference value f_m^* is dynamically updated whenever a superior solution is discovered. Consequently, the fitness of an individual is evaluated as the average relative gap across the training instances:

$$F_R = \frac{1}{M} \sum_{m=1}^{M} \frac{f(\tau_R(D_m)) - f_m^*}{|f_m^*| + \epsilon} \tag{4}$$

For numerical stability, a small value $\epsilon > 0$ (e.g., 10^{-8}) is added to the denominator.

3.2.2 Semantic Distance Metric. Defining a distance metric based on tree syntax is ineffective for capturing functional differences. Instead, we define Semantic Distance based on the actual decisions made by rules in identical scenarios. We construct a reference set $\Omega = \{S_1, S_2, \ldots, S_E\}$ containing E distinct state snapshots sampled from historical operational data. For any two rules R_i and R_j, the semantic distance D_{sem} is defined as the probability that they select different actions for the same state:

$$D_{sem}(R_i, R_j) = \frac{1}{E} \sum_{k=1}^{E} \mathbb{I}(\pi(R_i, S_k) \neq \pi(R_j, S_k)) \tag{5}$$

Here, $\pi(R, S_k)$ denotes the action index (e.g., the target destination for the currently triggered truck) selected by rule R in state S_k, and $\mathbb{I}(\cdot)$ is the indicator function, it returns 1 if the condition is true, and 0 otherwise.

3.2.3 Adaptive Anchor Selection Strategy. To embed these rules into a continuous coordinate system (the Semantic Manifold), we require a set of reference points, or "Anchors." We propose a Max-Min Dispersion strategy to select K anchors $\mathcal{A} = \{A_1, \ldots, A_K\}$ that maximize the volume of the spanned semantic space.

(1) Initialization: We generate a large initial random pool P_{large} and filter it to retain the top γ% based on fitness, forming the candidate pool Φ_{valid}.
(2) Core Anchor (A_1): The first anchor is the global best individual in Φ_{valid}, fixing the origin in the high-quality region.
(3) Fringe Anchors (A_k): Subsequent anchors are selected to be maximally distant from all previously selected anchors.

$$A_k = \arg\max_{R \in \Phi_{valid}} \left(\min_{j<k} D_{sem}(R, A_j) \right) \tag{6}$$

Any rule R is then mapped to a coordinate vector $\mathbf{v}_R \in \mathbb{R}^K$:

$$\mathbf{v}_R = [D_{sem}(R, A_1), \ldots, D_{sem}(R, A_K)]^\top \in [0, 1]^K \tag{7}$$

Description: This embedding map projects abstract rule behaviors into a K-dimensional Euclidean space. The k-th component of the vector represents the degree of behavioral difference of rule R relative to anchor point A_k. This relative coordinate representation allows us to utilize Euclidean geometric algorithms (such as K-Means) to process non-Euclidean semantic space structures.

3.2.4 CVT-MAP-Elites Archive Evolution. We use Centroidal Voronoi Tessellation (CVT) to partition the high-dimensional semantic space into N distinct regions (cells). The MAP-Elites algorithm maintains an archive $\mathcal{M}$ where each cell stores the best-performing rule found so far for that specific behavioral niche.

During evolution, a new offspring R_{child} is generated via GP crossover and mutation. Its coordinate $\mathbf{v}_{child}$ is calculated, and it is assigned to the nearest Voronoi cell c. The archive is updated if the cell is empty or if the child outperforms the current elite:

$$\mathcal{M}[c] \leftarrow R_{child}, \quad \text{if } \mathcal{M}[c] = \emptyset \text{ or } F_{R_{child}} < F_{\mathcal{M}[c]} \tag{8}$$

3.3 Online Heterogeneous Ant Colony

In the online phase, the system must respond to truck requests in real-time. We extract the top L distinct elite rules from the archive $\mathcal{M}$ to form the online strategy library $\mathcal{H} = \{H_1, \ldots, H_L\}$.

3.3.1 Heterogeneous Strategy Binding. We modify the standard ACO by introducing agent heterogeneity. The ant colony of size N_{ant} is partitioned into L groups. Each ant k is strictly bound to a specific heuristic rule $H_{(k \mod L)}$ from the library. This turns the colony into a multi-view decision system where different ants evaluate paths based on different "instincts".

3.3.2 Dual-Normalization Mechanism. A critical challenge is integrating the symbolic output of GP rules (which are unbounded and can be negative) with the pheromone trails of ACO (which are strictly positive). We propose a Dual-Normalization Mechanism to align these two signals mathematically.

Softmax Transformation for Heuristic. We convert the raw score S_{ij} assigned by rule H to a candidate node j into a probability distribution η_{ij} using a Softmax function with a temperature parameter T.

$$\eta_{ij} = \frac{\exp((S_{ij} - S_{max})/T)}{\sum_{l \in \mathcal{N}_i} \exp((S_{il} - S_{max})/T)} \tag{9}$$

This ensures $\eta_{ij} \in (0, 1)$ and preserves the relative ranking of choices.

Local Max Scaling for Pheromones. To match the scale of the heuristic info, we scale the pheromone concentration τ_{ij} relative to the maximum pheromone value in the current neighborhood $\mathcal{N}_i$.

$$\tau'_{ij} = \frac{\tau_{ij}}{\max_{l \in \mathcal{N}_i}(\tau_{il})} \tag{10}$$

This ensures that the most established path always has a coefficient of 1.0 and perfectly preserves the ratio characteristics of historical data, making the weights α and β physically meaningful.

3.3.3 Probabilistic Transition and Update. The probability P_{ij}^k that ant k moves from node i to node j is calculated by fusing the normalized signals:

$$P_{ij}^k = \frac{(\tau'_{ij})^\alpha \cdot (\eta_{ij})^\beta}{\sum_{l \in \mathcal{N}_i} (\tau'_{il})^\alpha \cdot (\eta_{il})^\beta} \tag{11}$$

Here, α and β control the relative weights of historical experience (pheromones) and current heuristic information (GP rules), respectively. Since τ' and η have been scaled to the same magnitude, parameters α and β can accurately reflect the importance of the

two information sources, avoiding parameter failure issues caused by vastly different numerical ranges.

Subsequently, the pheromone trails are updated using the Max-Min Ant System (MMAS) [21] mechanism. To prevent stagnation and premature convergence, global evaporation is applied to all edges, while pheromone reinforcement is exclusively restricted to the iteration-best solution:

$$\tau_{ij} \leftarrow (1-\rho) \cdot \tau_{ij} + \Delta\tau_{ij}^{best} \tag{12}$$

where $\rho \in (0,1)$ denotes the evaporation rate. The term $\Delta\tau_{ij}^{best}$ represents the pheromone increment, which is proportional to the quality of the best solution found in the current iteration if edge (i,j) belongs to it, and zero otherwise.

Finally, to ensure continuous exploration capability, the pheromone values are clamped within $[\tau_{min}, \tau_{max}]$ to prevent stagnation. This bounding strategy guarantees that the probability of selecting any specific path never drops to zero, allowing the system to rapidly adapt its strategy when dynamic events render the current best path suboptimal.

4 Experimental Environment

To systematically evaluate the performance of the proposed GP-ACO framework, we established a high-fidelity simulation platform calibrated with real-world mining data. This section details the dataset generation, parameter configurations, and the evaluation metrics used for comparative analysis.

4.1 Simulation Platform and Datasets

To guarantee that the experimental results reflect real-world operational complexities, the simulation environment is driven by empirical data collected from open-pit mining operations. Detailed information about the mine datasets is provided in Supplementary Material 2. The stochastic nature of the mining process is modeled through the following mechanisms:

Stochastic Loading and Haulage. Rather than using static averages, the shovel loading duration is modeled as a random variable following a normal distribution ($\mu = 6.80, \sigma = 0.97$), derived from statistical analysis of on-site operation logs. Similarly, truck travel times are not constant; they are generated via a deep learning-based predictor that accounts for road segment attributes (gradient, distance, curvature), weather, and vehicle physics, ensuring accurate reflection of dynamic hauling conditions.

Topological and Scale Diversity. The datasets encompass three distinct mining road network topologies (C1, C2, and C3) characterized by increasing complexity in terms of intersection density and route multiplicity. For each topology, we defined four equipment scale configurations (D1 through D4) to test scalability. These configurations range from small fleets (e.g., 3 shovels and 6 trucks) to large-scale operations (e.g., 9 shovels and 27 trucks), adhering to optimal truck-to-shovel ratios derived from fleet utility analysis.

4.2 Dynamic and Uncertain Scenarios

To rigorously assess the robustness of the scheduling strategies, the experiments incorporate five distinct scenarios characterized by varying degrees of dynamics and uncertainty (E1-E5):

- E1 (Baseline): A standard operational environment subject only to inherent stochastic fluctuations in loading and travel times without external disruption events.
- E2 (Truck Failure): Simulates sudden equipment breakdowns where specific trucks become unavailable for random durations, forcing the scheduler to manage reduced haulage capacity.
- E3 (Shovel Failure): Simulates loading equipment malfunctions, which creates bottlenecks at source nodes and necessitates flow redirection.
- E4 (Weather Fluctuation): Introduces global perturbations where sudden weather changes (e.g., heavy rain or fog) degrade travel speeds across the entire road network simultaneously.
- E5 (Road Blockage): Simulates localized disruptions where specific road segments become impassable, altering the graph topology and requiring the discovery of alternative paths.

Combining the three mine topologies, four equipment scales, and five dynamic scenarios, the test suite comprises a diverse set of problem instances to evaluate algorithmic generalization.

4.3 Baseline Methods

We compare the proposed GP-ACO against a suite of representative algorithms, spanning evolutionary computation, swarm intelligence, and deep reinforcement learning, widely used in dynamic scheduling and routing. The following methods are selected as benchmarks:

- EA [7]: A standard Evolutionary Algorithm that evolves scheduling solutions directly.
- ACO [2]: The canonical Ant Colony Optimization algorithm, serving as a baseline for constructive search methods.
- MMAS-LS [14]: Max-Min Ant System augmented with Local Search, known for preventing premature convergence.
- MACO [20]: Memory-based Ant Colony Optimization designed to recall high-quality solutions in changing environments.
- ACFGA [26]: An adaptive genetic algorithm utilizing fractional sub-populations to maintain diversity.
- DPCGA [13]: A dynamic clustering-based genetic algorithm that leverages density peaks to guide the search process.
- DRL [11]: An deep learning model trained via reinforcement learning, representing state-of-the-art neural combinatorial optimization for learning routing heuristics.

4.4 Parameter Settings

Baseline Parameters. The hyperparameters for the baseline methods were configured according to their original literature and empirically tuned for the mining datasets. For population-based methods (EA, DPCGA, ACFGA), the population size is set to 100 with a maximum of 100,000 evaluations. For ACO-based variants, the pheromone importance factor α is set to 1.0, the heuristic factor β to 5.0, and the evaporation rate ρ to 0.3.

GP-ACO Parameters. The proposed framework requires specific configurations for both the offline semantic learning phase and the online heterogeneous ant colony phase. Table 1 details these

settings. In the offline phase, the GPHH utilizes a population of 100 individuals evolved over 1,000 generations to populate the MAP-Elites archive. The semantic manifold is constructed using $K = 5$ anchor points to ensure sufficient behavioral resolution. In the online phase, the heterogeneous colony consists of 50 ants, with the top 5 distinct elite rules ($L = 5$) extracted from the archive to form the strategy library.

Table 1: Parameter Settings for GP-ACO

Parameter	Value	Description
Offline (GPHH)		
N_{gp}	100	Number of trees in GP population
G_{max}	1,000	Maximum evolutionary generations
p_{c}	0.9	Probability of subtree crossover
p_{m}	0.1	Probability of subtree mutation
D_{tree}	10-15	Min/Max depth of heuristic trees
Offline (MAP-Elites)		
N_{cells}	200	Number of CVT Voronoi cells
K	5	Number of semantic anchor points
P_{large}	10,000	Samples for initial centroid generation
E	2,000	Number of distinct state snapshots
γ	50%	Fitness gating threshold
Online (ACO)		
N_{ant}	40	Total number of ants
L	5	Number of elite rules used online
α	1.0	Pheromone weight
β	4.0	Heuristic weight
ρ	0.3	Pheromones evaporation rate

4.5 Evaluation Metrics

To provide a comprehensive assessment of solution quality and the internal behavior of the algorithm, we employ the following two metrics:

Average Optimality Gap (Gap). Consistent with the problem objective, we evaluate the ore flow stability using the Root Mean Squared Error (RMSE) of the grade. To standardize the comparison across heterogeneous instances, we report the Average Optimality Gap. This metric quantifies the percentage deviation of an algorithm's solution R_i from the best-known solution R_i^{best} found across all comparison methods for instance i:

$$Gap = \frac{1}{N}\sum_{i=1}^{N}\frac{R_i - R_i^{best}}{R_i^{best}} \times 100\% \tag{13}$$

A lower Gap value indicates superior performance in controlling grade fluctuations.

Semantic Diversity (Div). To verify whether the proposed Semantic Manifold Learning effectively encourages behavioral diversity, we introduce a semantic diversity metric. For the set of elite rules $\mathcal{M}$ evolved in the offline phase, we calculate the average pairwise semantic distance. This metric reflects the degree of behavioral non-redundancy among the strategies in the offline phase:

$$Div(\mathcal{M}) = \frac{2}{M(M-1)}\sum_{i=1}^{M}\sum_{j=i+1}^{M} D_{sem}(R_i, R_j) \tag{14}$$

where $D_{sem}(R_i, R_j)$ represents the semantic distance between rules R_i and R_j. A higher Div value implies that the colony possesses a broader repertoire of distinct decision-making strategies, theoretically enhancing robustness against dynamic disturbances.

5 Results And Discussion

5.1 Analysis of Semantic Manifold Learning

To validate the quality-diversity mechanism, we compared the top-50 rules evolved by our semantic manifold learning against a standard fitness-based genetic programming (Standard GP) approach. Fig. 1 illustrates their distribution in the quality-diversity space. Standard GP individuals cluster tightly in low-diversity regions, indicating premature convergence to structurally similar local optima (semantic cloning). While their average quality is acceptable, this behavioral redundancy limits their adaptability to diverse dynamic events. In contrast, GP-ACO maps rules into a broader semantic space, successfully discovering and preserving distinct, high-performing behaviors. This confirms our offline QD approach effectively constructs a comprehensive and behaviorally diverse rule library.

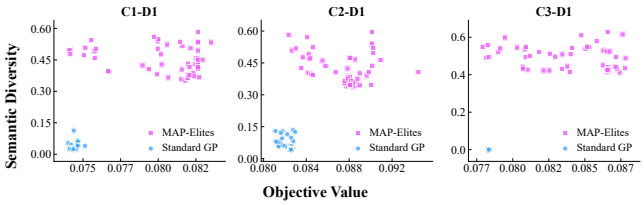


Figure 1: Distribution of the top 50 elite rules in the quality-diversity space across three mining scenarios

5.2 Impact of Heterogeneous Strategy Library

To evaluate the heterogeneous ant colony mechanism, we varied the strategy library size L (number of elite rules deployed online) from 0 to 5, where $L = 0$ implies strict reliance on manual heuristics. Fig. 2 shows that introducing evolved rules ($L = 1$) substantially reduces the average RMSE, proving that learned rules capture complex dependencies better than manual heuristics. Furthermore, increasing L from 1 to 5 consistently improves performance, confirming that a heterogeneous colony outperforms a homogeneous one. By equipping ant sub-populations with distinct decision-making instincts, GP-ACO conducts parallel exploration from multiple semantic perspectives, effectively bypassing local optima. Marginal gains stabilize around $L = 4$ to 5, indicating that a compact, diverse rule set sufficiently covers the necessary behavioral niches.

5.3 Real Time Response Analysis

To explicitly evaluate the real time response capabilities of the algorithms, we analyzed their dynamic convergence behavior across 200

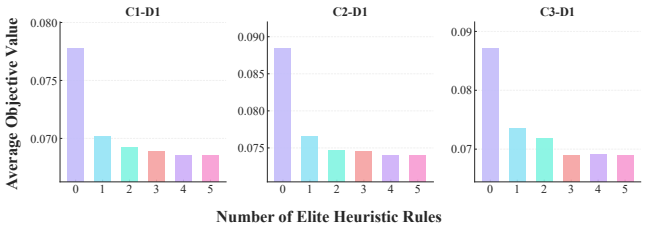


Figure 2: Impact of the heterogeneous strategy library size across three mining scenarios

test instances of the C3-D1 scenario. Fig. 3 tracks the average objective value against the number of evaluations. The periodic spikes in the convergence curves represent sudden dynamic disruptions that invalidate the currently executing schedule.

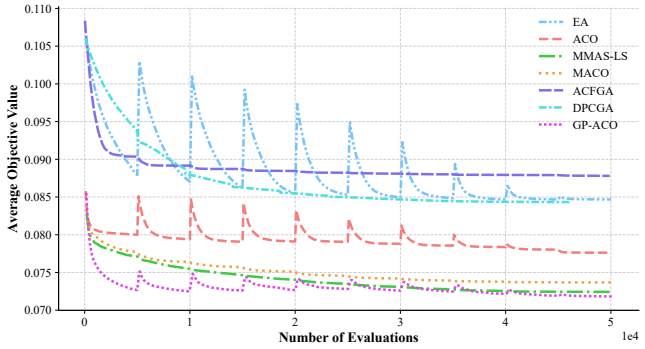


Figure 3: Convergence curves of different algorithms on the C3-D1 test scenario.

As illustrated, traditional evolutionary algorithms (EA, ACFGA, DPCGA) suffer from severe performance degradation when environmental changes occur, requiring extensive blind evaluations to recover. Standard swarm intelligence methods (ACO, MMAS-LS, MACO) adapt slightly better due to memory and local search mechanisms but still exhibit significant lag. In stark contrast, GP-ACO demonstrates exceptional real time agility. It achieves a significantly lower initial objective value and recovers almost instantaneously after each dynamic disruption. By directly exploiting the offline learned heterogeneous strategy library, the ant colony bypasses random exploration and rapidly converges to robust schedules, perfectly fulfilling the rapid response requirement of dynamic ore flow blending.

5.4 Parameter Sensitivity Analysis

To systematically evaluate the robustness of the GP-ACO framework to its key hyperparameters, we conducted an extensive sensitivity analysis using 200 test instances from the C1-D1 scenario. The results are presented in Fig. 4.

The heuristic weight β dictates the influence of the semantic rules. The performance trend stabilizes when β reaches 4 and beyond, making 4 the optimal parameter setting. Lower values cause the colony to underutilize valuable evolved heuristics. Regarding the ant population size N_{ant}, the algorithm performs best at a size of 40. Overall, this parameter has a relatively minor impact on performance compared to others.

The temperature T in the softmax transformation regulates the selection pressure of the heuristic rules. The optimal performance is observed at T equal to 1.0. Values near zero enforce extreme greediness, while higher temperatures dilute the heuristic signals into uniform random noise. Finally, the pheromone evaporation rate ρ controls the memory retention of the swarm. A rate around 0.3 provides the ideal balance, effectively preventing both search stagnation caused by zero evaporation and the loss of high quality historical paths associated with extreme evaporation.

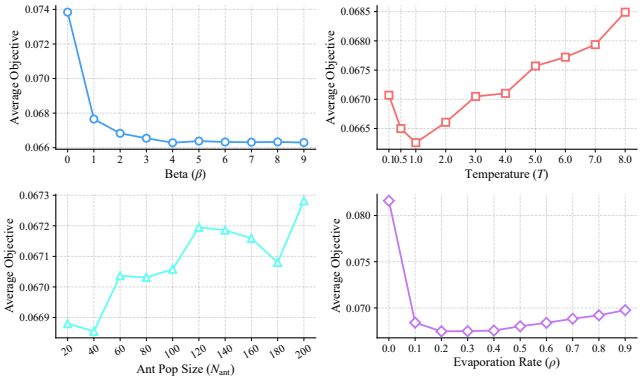


Figure 4: Sensitivity analysis of key hyperparameters evaluated on the C1-D1 test scenario.

5.5 Comprehensive Comparative Evaluation

We benchmarked GP-ACO against seven state-of-the-art baselines across the 60 test scenarios. As detailed in Table 2, GP-ACO achieves an overwhelming 89.3% win rate (based on pairwise Wilcoxon signed-rank tests) and maintains the lowest relative optimality gap globally. This superiority is evident across three dimensions:

- **Scalability (D1–D4):** While standard evolutionary methods (e.g., EA, ACFGA) degrade rapidly in large-scale search spaces, GP-ACO maintains stable performance by leveraging transferable heuristic rules.
- **Robustness (E1–E5):** GP-ACO exhibits exceptional resilience against dynamic disturbances. Our heterogeneous library enables the swarm to rapidly pivot to robust alternative behaviors during equipment failures or road blockages.
- **Generalization (C1–C3):** The framework consistently adapts to varying road network topologies without manual rule tuning.

Collectively, these results confirm that GP-ACO effectively balances offline knowledge acquisition with online agility, establishing a new state-of-the-art for dynamic ore-flow blending.

6 Conclusions

This paper proposes a heterogeneous ant colony framework with semantic manifold learning (GP-ACO) to resolve the conflict between global optimality and real-time responsiveness in dynamic

Table 2: Detailed Performance Comparison on 60 Test Scenarios. The Obj column displays the mean objective value alongside the Wins/Ties/Losses ($W/T/L$) computed via pairwise Wilcoxon signed-rank tests ($\alpha = 0.05$) against the other 7 algorithms based on the 40 instances. The Gap/% column presents the Mean ± Std of the relative gap. For each scenario, the algorithm achieving the maximum Wins has its $W/T/L$ highlighted in bold.

Case	Scale	Event	EA		ACO		MMAS-LS		MACO		ACFGA		DPCGA		DRL		GP-ACO	
			Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%	Obj ($W/T/L$)	Gap/%
C1	D1	E1	0.080 (0/2/5)	13.7 ± 8.9	0.075 (4/2/1)	5.8 ± 5.2	0.074 (4/2/1)	4.7 ± 4.5	0.076 (3/0/4)	7.6 ± 4.9	0.082 (0/2/5)	17.0 ± 12.4	0.080 (0/2/5)	14.2 ± 9.1	0.074 (4/2/1)	4.6 ± 4.4	0.073 **(7/0/0)**	1.7 ± 2.6
		E2	0.070 (1/2/4)	14.9 ± 10.2	0.065 (4/2/1)	7.2 ± 7.1	0.065 (4/3/0)	6.3 ± 5.7	0.067 (2/2/3)	10.6 ± 6.7	0.075 (0/0/7)	22.7 ± 16.5	0.071 (1/1/5)	17.6 ± 12.2	0.067 (3/3/1)	6.0 ± 6.1	0.063 **(6/1/0)**	2.4 ± 3.1
		E3	0.077 (0/3/4)	18.2 ± 15.0	0.070 **(5/2/0)**	5.3 ± 5.6	0.069 **(5/2/0)**	3.6 ± 4.4	0.071 (4/1/2)	7.2 ± 5.1	0.080 (0/3/4)	24.3 ± 19.8	0.077 (0/3/4)	19.4 ± 17.7	0.080 (0/3/4)	24.2 ± 23.7	0.070 (4/3/0)	5.1 ± 7.4
		E4	0.072 (0/3/4)	16.9 ± 12.2	0.065 **(5/2/0)**	4.5 ± 3.7	0.065 **(5/2/0)**	5.0 ± 4.4	0.067 (3/1/3)	8.0 ± 5.2	0.073 (0/2/5)	22.7 ± 19.9	0.071 (1/2/4)	17.4 ± 14.0	0.071 (0/4/3)	10.2 ± 13.2	0.064 **(5/2/0)**	1.3 ± 2.0
		E5	0.082 (1/1/5)	14.8 ± 12.5	0.074 (4/3/0)	3.9 ± 3.5	0.074 (4/3/0)	3.6 ± 4.0	0.076 (3/1/3)	6.7 ± 4.5	0.083 (0/0/7)	20.7 ± 18.4	0.082 (1/1/5)	18.8 ± 17.0	0.076 (3/3/1)	4.4 ± 4.1	0.073 **(5/2/0)**	2.5 ± 3.3
	D2	E1	0.077 (0/2/5)	62.3 ± 43.7	0.058 (4/2/1)	10.7 ± 7.5	0.058 (4/2/1)	12.1 ± 8.2	0.060 (3/0/4)	16.3 ± 11.6	0.075 (0/2/5)	49.9 ± 29.4	0.076 (0/2/5)	49.2 ± 28.1	0.057 (4/2/1)	5.6 ± 9.0	0.054 **(7/0/0)**	0.6 ± 1.8
		E2	0.079 (0/2/5)	39.1 ± 23.2	0.064 (4/2/1)	9.2 ± 4.8	0.064 (4/2/1)	9.3 ± 5.5	0.067 (3/0/4)	16.0 ± 9.4	0.080 (0/2/5)	34.7 ± 23.5	0.078 (0/2/5)	36.4 ± 23.6	0.063 (4/3/0)	4.3 ± 5.8	0.061 **(6/1/0)**	4.2 ± 6.5
		E3	0.076 (0/2/5)	42.8 ± 25.5	0.059 (4/2/1)	6.9 ± 4.6	0.060 (4/2/1)	6.8 ± 5.8	0.061 (3/0/4)	12.3 ± 6.9	0.076 (0/2/5)	44.2 ± 27.4	0.076 (0/2/5)	40.8 ± 22.2	0.059 (4/2/1)	6.9 ± 8.4	0.057 **(7/0/0)**	2.4 ± 3.2
		E4	0.079 (0/2/5)	50.5 ± 32.5	0.060 (4/2/1)	8.1 ± 7.6	0.061 (4/2/1)	9.5 ± 7.0	0.063 (3/0/4)	15.3 ± 7.2	0.079 (0/1/6)	52.5 ± 32.5	0.078 (1/1/5)	47.1 ± 28.7	0.060 (4/3/0)	6.1 ± 9.7	0.057 **(6/1/0)**	2.5 ± 2.8
		E5	0.081 (0/1/6)	38.2 ± 24.4	0.065 (4/2/1)	8.2 ± 5.9	0.065 (4/1/2)	8.1 ± 5.2	0.069 (3/0/4)	14.3 ± 6.9	0.082 (0/1/6)	35.1 ± 20.6	0.079 (2/0/5)	32.5 ± 21.6	0.063 (5/1/1)	4.9 ± 6.7	0.062 **(7/0/0)**	1.2 ± 1.5
	D3	E1	0.078 (0/1/6)	63.6 ± 35.8	0.057 (4/2/1)	10.8 ± 7.7	0.058 (4/2/1)	11.4 ± 9.8	0.061 (3/0/4)	20.9 ± 11.4	0.078 (0/1/6)	65.1 ± 36.4	0.076 (2/0/5)	61.9 ± 36.4	0.053 **(7/0/0)**	1.9 ± 3.6	0.057 (4/2/1)	12.0 ± 14.6
		E2	0.083 (0/0/7)	62.9 ± 30.3	0.059 (4/1/2)	11.2 ± 6.9	0.060 (4/1/2)	13.7 ± 7.6	0.064 (3/0/4)	21.7 ± 9.8	0.081 (1/1/5)	58.9 ± 29.7	0.081 (1/1/5)	56.6 ± 28.1	0.055 **(7/0/0)**	1.5 ± 2.4	0.057 (6/0/1)	6.5 ± 7.4
		E3	0.086 (0/1/6)	66.2 ± 42.4	0.061 (4/2/1)	11.5 ± 7.6	0.061 (4/2/1)	11.1 ± 8.3	0.065 (3/0/4)	20.2 ± 10.1	0.086 (0/2/5)	71.4 ± 48.3	0.085 (1/1/5)	62.9 ± 41.7	0.057 **(7/0/0)**	2.9 ± 4.9	0.060 (4/2/1)	7.7 ± 8.7
		E4	0.078 (0/1/6)	69.7 ± 35.0	0.054 (4/2/1)	11.8 ± 10.8	0.055 (4/2/1)	12.5 ± 10.7	0.058 (3/0/4)	19.4 ± 10.7	0.078 (0/1/6)	69.3 ± 34.3	0.076 (2/0/5)	62.4 ± 30.3	0.051 **(6/1/0)**	6.0 ± 7.3	0.054 (4/3/0)	9.8 ± 10.0
		E5	0.083 (0/1/6)	61.3 ± 35.3	0.061 (4/2/1)	12.0 ± 7.7	0.061 (4/2/1)	10.5 ± 8.7	0.064 (3/0/4)	19.1 ± 9.1	0.083 (0/1/6)	60.7 ± 34.7	0.081 (2/0/5)	57.9 ± 34.7	0.057 **(6/1/0)**	3.9 ± 5.6	0.059 (4/3/0)	9.1 ± 11.5
	D4	E1	0.080 (0/1/6)	92.8 ± 45.7	0.056 (4/1/2)	25.2 ± 14.4	0.057 (4/1/2)	26.8 ± 13.3	0.059 (3/0/4)	33.9 ± 14.2	0.080 (0/1/6)	91.1 ± 43.4	0.077 (2/0/5)	77.4 ± 34.7	0.049 (6/0/1)	9.6 ± 10.2	0.045 **(7/0/0)**	0.3 ± 0.7
		E2	0.080 (1/0/6)	92.2 ± 43.1	0.053 (4/1/2)	24.1 ± 12.4	0.056 (4/1/2)	26.3 ± 12.8	0.058 (3/0/4)	38.0 ± 15.7	0.081 (0/0/7)	100.3 ± 52.5	0.078 (2/0/5)	93.5 ± 50.2	0.048 (6/0/1)	13.8 ± 11.2	0.043 **(7/0/0)**	0.0 ± 0.0
		E3	0.080 (1/0/6)	86.8 ± 39.1	0.056 (4/1/2)	23.2 ± 13.7	0.057 (4/1/2)	23.0 ± 13.1	0.061 (3/0/4)	35.5 ± 19.1	0.082 (0/0/7)	90.1 ± 41.8	0.079 (2/0/5)	81.4 ± 36.2	0.050 (6/0/1)	12.2 ± 11.4	0.046 **(7/0/0)**	0.1 ± 0.3
		E4	0.073 (1/1/5)	88.3 ± 46.1	0.050 (4/1/2)	23.8 ± 15.5	0.051 (4/1/2)	23.9 ± 17.8	0.054 (3/0/4)	34.9 ± 21.7	0.075 (0/0/7)	102.3 ± 55.1	0.072 (1/1/5)	84.7 ± 46.8	0.044 (6/0/1)	9.2 ± 9.6	0.042 **(7/0/0)**	1.6 ± 3.1
		E5	0.084 (0/1/6)	90.2 ± 43.7	0.058 (4/1/2)	23.8 ± 11.7	0.059 (4/1/2)	25.2 ± 13.7	0.063 (3/0/4)	35.3 ± 15.8	0.085 (0/1/6)	93.4 ± 45.3	0.082 (2/0/5)	85.3 ± 42.7	0.052 (6/0/1)	12.1 ± 10.8	0.048 **(7/0/0)**	0.0 ± 0.0
C2	D1	E1	0.090 (0/2/5)	16.9 ± 12.5	0.084 (4/0/3)	8.9 ± 5.3	0.083 (5/1/1)	6.6 ± 5.1	0.085 (3/0/4)	11.3 ± 7.0	0.091 (0/2/5)	16.5 ± 13.7	0.090 (0/2/5)	17.7 ± 13.2	0.081 (5/1/1)	6.6 ± 5.6	0.079 **(7/0/0)**	0.7 ± 1.5
		E2	0.089 (2/1/4)	12.7 ± 9.9	0.084 (4/2/1)	6.3 ± 5.5	0.084 (4/2/1)	6.4 ± 4.5	0.087 (2/1/4)	8.9 ± 5.6	0.094 (0/1/6)	18.7 ± 12.0	0.092 (0/1/6)	13.1 ± 9.4	0.083 (4/2/1)	5.9 ± 5.4	0.079 **(7/0/0)**	0.1 ± 0.2
		E3	0.088 (0/2/5)	18.6 ± 14.6	0.082 (4/2/1)	8.4 ± 6.9	0.081 (4/2/1)	6.7 ± 6.6	0.084 (3/1/3)	10.3 ± 5.9	0.088 (0/2/5)	17.4 ± 13.3	0.088 (0/2/5)	19.6 ± 15.7	0.083 (3/3/1)	6.1 ± 5.1	0.078 **(7/0/0)**	2.5 ± 3.9
		E4	0.087 (0/2/5)	13.5 ± 10.7	0.079 (4/2/1)	5.4 ± 3.9	0.079 (4/2/1)	5.9 ± 5.7	0.081 (2/2/3)	10.0 ± 7.2	0.088 (0/2/5)	16.9 ± 13.3	0.087 (0/3/4)	14.0 ± 11.9	0.080 (3/3/1)	6.4 ± 5.8	0.076 **(7/0/0)**	1.7 ± 3.2
		E5	0.078 (1/1/5)	16.2 ± 11.0	0.072 (3/2/2)	7.6 ± 4.9	0.071 (5/1/1)	6.0 ± 6.1	0.073 (3/2/2)	8.6 ± 6.4	0.083 (0/0/7)	23.3 ± 17.3	0.078 (1/1/5)	16.2 ± 13.2	0.072 (3/3/1)	8.1 ± 7.3	0.068 **(7/0/0)**	2.4 ± 3.4
	D2	E1	0.097 (0/1/6)	39.5 ± 26.8	0.078 (4/1/2)	12.7 ± 10.3	0.079 (4/1/2)	14.4 ± 9.9	0.082 (3/0/4)	17.5 ± 10.2	0.097 (0/2/5)	45.0 ± 31.4	0.095 (1/1/5)	35.3 ± 22.5	0.073 **(6/1/0)**	2.8 ± 3.4	0.074 **(6/1/0)**	4.3 ± 6.2
		E2	0.087 (0/2/5)	40.7 ± 30.5	0.070 (4/2/1)	10.0 ± 8.5	0.070 (4/2/1)	9.2 ± 6.7	0.074 (3/0/4)	14.1 ± 6.9	0.088 (0/2/5)	38.5 ± 26.1	0.087 (0/2/5)	37.3 ± 25.8	0.067 **(6/1/0)**	4.9 ± 6.0	0.068 (4/3/0)	6.4 ± 7.8
		E3	0.088 (0/2/5)	36.5 ± 20.4	0.072 (4/2/1)	13.3 ± 10.7	0.071 (4/3/0)	12.5 ± 11.4	0.075 (3/0/4)	19.2 ± 11.8	0.088 (0/2/5)	39.7 ± 24.4	0.087 (0/2/5)	41.4 ± 27.1	0.069 **(5/2/0)**	4.9 ± 7.5	0.071 (4/3/0)	9.1 ± 11.0
		E4	0.092 (1/1/5)	38.5 ± 24.3	0.074 (4/2/1)	11.0 ± 7.8	0.074 (4/2/1)	9.7 ± 8.7	0.077 (3/0/4)	12.0 ± 6.6	0.094 (0/0/7)	43.5 ± 33.2	0.092 (1/1/5)	44.2 ± 31.2	0.069 **(6/1/0)**	3.6 ± 4.7	0.073 (4/3/0)	6.3 ± 8.4
		E5	0.088 (1/0/6)	41.5 ± 21.0	0.070 (4/1/2)	10.5 ± 8.2	0.068 (5/1/1)	8.6 ± 6.6	0.073 (3/0/4)	15.8 ± 8.4	0.090 (0/0/7)	46.3 ± 25.8	0.087 (2/0/5)	40.6 ± 21.5	0.065 **(6/1/0)**	3.7 ± 4.9	0.069 (4/3/0)	11.0 ± 12.7
	D3	E1	0.093 (0/1/6)	63.9 ± 34.9	0.067 (4/1/2)	15.9 ± 10.6	0.068 (4/1/2)	14.2 ± 8.9	0.071 (3/0/4)	24.1 ± 11.5	0.094 (0/1/6)	63.7 ± 32.8	0.091 (2/0/5)	63.9 ± 37.7	0.060 **(6/1/0)**	3.0 ± 3.5	0.061 **(6/1/0)**	2.3 ± 3.8
		E2	0.089 (0/2/5)	63.5 ± 33.2	0.065 (4/1/2)	21.5 ± 11.9	0.066 (4/1/2)	22.0 ± 13.9	0.069 (3/0/4)	28.2 ± 13.7	0.089 (0/2/5)	70.8 ± 42.3	0.088 (0/2/5)	63.4 ± 33.8	0.056 **(6/1/0)**	4.0 ± 4.8	0.057 **(6/1/0)**	5.8 ± 7.9
		E3	0.096 (0/1/6)	66.5 ± 43.9	0.069 (4/1/2)	15.0 ± 8.5	0.070 (4/2/1)	14.0 ± 9.3	0.073 (3/0/4)	21.3 ± 11.7	0.095 (0/2/5)	64.9 ± 44.9	0.094 (1/1/5)	62.7 ± 42.0	0.064 **(6/1/0)**	7.0 ± 8.7	0.066 (5/2/0)	8.1 ± 11.7
		E4	0.095 (0/2/5)	51.0 ± 26.8	0.076 (4/1/2)	17.9 ± 9.1	0.076 (4/1/2)	18.3 ± 9.6	0.080 (3/0/4)	23.7 ± 10.5	0.095 (0/1/6)	53.5 ± 28.8	0.094 (1/1/5)	49.6 ± 26.9	0.067 **(6/1/0)**	4.2 ± 5.0	0.066 **(6/1/0)**	2.4 ± 4.2
		E5	0.093 (0/2/5)	59.8 ± 29.7	0.070 (4/1/2)	20.3 ± 12.5	0.071 (4/1/2)	20.6 ± 12.6	0.075 (3/0/4)	28.9 ± 11.9	0.092 (0/2/5)	64.2 ± 35.8	0.091 (0/2/5)	58.0 ± 28.2	0.061 **(6/1/0)**	3.5 ± 5.1	0.061 **(6/1/0)**	2.9 ± 4.3
	D4	E1	0.094 (0/1/6)	81.9 ± 38.0	0.069 (5/0/2)	29.2 ± 12.7	0.071 (4/0/3)	32.3 ± 13.3	0.072 (3/0/4)	37.3 ± 16.5	0.094 (0/1/6)	85.5 ± 43.3	0.092 (2/0/5)	84.4 ± 47.1	0.060 (6/0/1)	11.4 ± 9.9	0.053 **(7/0/0)**	0.0 ± 0.0
		E2	0.096 (0/1/6)	103.9 ± 56.7	0.067 (4/1/2)	35.1 ± 15.2	0.067 (4/1/2)	31.3 ± 13.1	0.070 (3/0/4)	44.1 ± 20.1	0.095 (0/2/5)	108.5 ± 60.1	0.094 (1/1/5)	98.5 ± 53.9	0.056 (6/0/1)	12.4 ± 10.7	0.050 **(7/0/0)**	0.0 ± 0.0
		E3	0.097 (0/1/6)	79.7 ± 44.5	0.072 (4/1/2)	27.5 ± 13.9	0.072 (4/1/2)	25.7 ± 12.0	0.075 (3/0/4)	35.1 ± 15.2	0.097 (0/1/6)	76.2 ± 37.7	0.095 (2/0/5)	75.8 ± 42.9	0.062 (6/0/1)	7.3 ± 7.3	0.057 **(7/0/0)**	0.1 ± 0.3
		E4	0.095 (0/1/6)	78.8 ± 42.7	0.070 (4/1/2)	26.4 ± 12.4	0.070 (4/1/2)	26.8 ± 16.0	0.074 (3/0/4)	37.5 ± 16.7	0.096 (0/1/6)	84.5 ± 47.8	0.094 (2/0/5)	80.2 ± 48.4	0.061 (6/0/1)	10.3 ± 8.9	0.055 **(7/0/0)**	0.0 ± 0.0
		E5	0.095 (0/1/6)	79.4 ± 33.5	0.065 (4/1/2)	24.7 ± 11.8	0.066 (4/1/2)	30.1 ± 16.4	0.068 (3/0/4)	33.0 ± 11.2	0.093 (0/2/5)	83.2 ± 43.0	0.092 (1/1/5)	76.9 ± 38.6	0.058 (6/0/1)	12.5 ± 9.8	0.051 **(7/0/0)**	0.0 ± 0.0
C3	D1	E1	0.083 (1/1/5)	29.1 ± 21.6	0.074 (4/2/1)	9.5 ± 8.4	0.075 (4/2/1)	10.5 ± 8.9	0.077 (3/1/3)	14.8 ± 9.6	0.087 (0/0/7)	34.2 ± 24.2	0.084 (1/1/5)	24.2 ± 16.6	0.075 (3/3/1)	10.6 ± 10.6	0.069 **(7/0/0)**	0.1 ± 0.3
		E2	0.088 (1/1/5)	24.5 ± 20.2	0.078 (4/1/2)	10.4 ± 6.8	0.077 (5/1/1)	7.7 ± 5.7	0.080 (3/1/3)	13.5 ± 8.9	0.090 (0/0/7)	34.4 ± 29.3	0.087 (1/1/5)	25.9 ± 20.6	0.078 (3/3/1)	9.2 ± 5.6	0.072 **(7/0/0)**	0.3 ± 0.7
		E3	0.078 (0/3/4)	22.6 ± 18.4	0.068 (5/0/2)	9.3 ± 6.9	0.067 **(6/1/0)**	6.5 ± 6.0	0.070 (4/0/3)	13.4 ± 9.7	0.079 (2/5/5)	22.7 ± 17.6	0.077 (1/2/4)	29.9 ± 28.4	0.078 (0/3/4)	18.3 ± 18.3	0.067 **(6/1/0)**	2.0 ± 3.7
		E4	0.088 (1/1/5)	21.9 ± 14.3	0.080 (4/1/2)	11.3 ± 7.9	0.079 (5/1/1)	9.7 ± 7.8	0.082 (3/0/4)	13.0 ± 7.7	0.090 (0/0/7)	24.9 ± 14.5	0.087 (1/1/5)	20.9 ± 12.3	0.079 (4/2/1)	8.7 ± 5.8	0.073 **(7/0/0)**	0.0 ± 0.0
		E5	0.085 (1/1/5)	27.0 ± 19.2	0.075 (4/1/2)	12.1 ± 9.2	0.074 (5/1/1)	10.1 ± 8.2	0.076 (3/1/3)	13.7 ± 8.0	0.089 (0/1/6)	31.0 ± 20.6	0.087 (0/2/5)	25.4 ± 13.7	0.075 (3/3/1)	8.8 ± 6.7	0.068 **(7/0/0)**	0.0 ± 0.2
	D2	E1	0.087 (0/2/5)	62.9 ± 35.2	0.066 (4/1/2)	19.4 ± 9.7	0.066 (4/1/2)	21.4 ± 13.8	0.069 (3/0/4)	27.1 ± 12.9	0.088 (0/2/5)	68.3 ± 39.9	0.087 (0/2/5)	67.0 ± 42.0	0.062 (6/0/1)	11.9 ± 9.1	0.055 **(7/0/0)**	0.0 ± 0.0
		E2	0.083 (1/1/5)	105.2 ± 83.9	0.062 (4/2/1)	24.7 ± 16.4	0.061 (4/2/1)	24.1 ± 15.1	0.064 (3/0/4)	30.2 ± 17.1	0.088 (0/0/7)	119.7 ± 92.8	0.082 (1/1/5)	101.0 ± 76.5	0.058 (4/2/1)	19.9 ± 20.6	0.051 **(7/0/0)**	0.0 ± 0.0
		E3	0.085 (1/0/6)	83.8 ± 65.8	0.062 (4/2/1)	19.8 ± 15.4	0.062 (4/2/1)	17.2 ± 12.2	0.065 (3/0/4)	26.3 ± 13.8	0.087 (0/0/7)	83.8 ± 61.3	0.083 (2/0/5)	84.8 ± 71.1	0.060 (4/2/1)	19.1 ± 20.6	0.054 **(7/0/0)**	0.0 ± 0.0
		E4	0.089 (1/1/5)	73.1 ± 56.2	0.066 (5/0/2)	18.7 ± 10.6	0.067 (4/0/3)	17.2 ± 8.1	0.070 (3/0/4)	24.3 ± 11.3	0.092 (0/0/7)	87.5 ± 71.0	0.090 (1/1/5)	76.1 ± 54.9	0.062 (6/0/1)	12.1 ± 10.4	0.056 **(7/0/0)**	0.0 ± 0.0
		E5	0.101 (0/1/6)	47.0 ± 28.9	0.085 (4/1/2)	21.5 ± 8.6	0.086 (4/1/2)	21.7 ± 8.8	0.088 (3/0/4)	25.5 ± 11.8	0.103 (0/1/6)	48.8 ± 30.9	0.100 (2/0/5)	47.9 ± 31.4	0.077 (6/0/1)	9.0 ± 8.0	0.070 **(7/0/0)**	0.0 ± 0.0
	D3	E1	0.090 (1/0/6)	134.8 ± 85.6	0.061 (4/1/2)	35.8 ± 23.7	0.060 (4/1/2)	34.9 ± 24.0	0.063 (3/0/4)	44.6 ± 26.7	0.092 (0/0/7)	142.4 ± 92.2	0.088 (2/0/5)	128.7 ± 81.9	0.050 (6/0/1)	10.3 ± 11.2	0.046 **(7/0/0)**	0.0 ± 0.0
		E2	0.097 (1/0/6)	71.2 ± 41.0	0.074 (4/1/2)	23.0 ± 11.4	0.075 (4/1/2)	25.4 ± 14.1	0.078 (3/0/4)	33.1 ± 14.9	0.099 (0/0/7)	74.4 ± 45.7	0.095 (2/0/5)	67.5 ± 42.0	0.064 (6/0/1)	8.6 ± 8.1	0.061 **(7/0/0)**	0.0 ± 0.0
		E3	0.089 (1/0/6)	85.8 ± 44.4	0.063 (4/1/2)	24.1 ± 17.1	0.062 (4/1/2)	21.8 ± 14.2	0.066 (3/0/4)	30.6 ± 12.7	0.091 (0/0/7)	90.0 ± 46.7	0.088 (2/0/5)	83.6 ± 43.3	0.054 (6/0/1)	10.8 ± 11.8	0.052 **(7/0/0)**	2.4 ± 4.1
		E4	0.089 (0/1/6)	104.6 ± 48.4	0.059 (5/0/2)	24.6 ± 13.8	0.061 (4/0/3)	28.6 ± 16.4	0.064 (3/0/4)	36.5 ± 15.9	0.090 (0/1/6)	105.8 ± 50.1	0.087 (2/0/5)	98.7 ± 45.4	0.053 (6/0/1)	12.7 ± 10.1	0.047 **(7/0/0)**	0.0 ± 0.0
		E5	0.087 (1/1/5)	103.6 ± 58.7	0.062 (4/1/2)	34.4 ± 22.6	0.061 (4/1/2)	34.8 ± 23.3	0.064 (3/0/4)	45.3 ± 24.0	0.089 (0/0/7)	112.7 ± 65.6	0.086 (1/1/5)	101.9 ± 58.0	0.052 (6/0/1)	13.8 ± 11.6	0.047 **(7/0/0)**	0.0 ± 0.0
	D4	E1	0.092 (1/0/6)	133.1 ± 79.8	0.056 (5/0/2)	29.5 ± 17.0	0.058 (4/0/3)	30.9 ± 14.0	0.060 (3/0/4)	39.3 ± 17.3	0.094 (0/0/7)	137.3 ± 81.2	0.090 (2/0/5)	125.2 ± 77.7	0.047 (6/0/1)	14.6 ± 16.8	0.043 **(7/0/0)**	0.0 ± 0.0
		E2	0.094 (1/0/6)	119.9 ± 72.4	0.063 (4/1/2)	31.5 ± 18.4	0.062 (4/1/2)	29.8 ± 19.2	0.066 (3/0/4)	38.8 ± 19.1	0.096 (0/0/7)	125.2 ± 78.0	0.093 (2/0/5)	115.5 ± 68.0	0.052 (6/0/1)	10.7 ± 12.9	0.049 **(7/0/0)**	1.4 ± 2.8
		E3	0.088 (0/1/6)	126.9 ± 70.5	0.053 (4/1/2)	28.4 ± 19.6	0.053 (4/1/2)	28.3 ± 15.9	0.057 (3/0/4)	35.7 ± 17.0	0.088 (0/1/6)	122.0 ± 67.2	0.085 (2/0/5)	119.1 ± 69.9	0.045 **(6/1/0)**	5.4 ± 6.8	0.044 **(6/1/0)**	2.2 ± 4.3
		E4	0.093 (0/1/6)	128.6 ± 67.2	0.058 (5/0/2)	35.5 ± 20.1	0.060 (4/0/3)	37.6 ± 17.3	0.063 (3/0/4)	45.7 ± 22.0	0.094 (0/1/6)	131.6 ± 70.6	0.091 (2/0/5)	123.1 ± 66.8	0.047 **(6/1/0)**	8.5 ± 12.3	0.045 **(6/1/0)**	3.8 ± 5.6
		E5	0.092 (0/1/6)	110.0 ± 55.0	0.059 (4/1/2)	28.3 ± 17.6	0.060 (4/1/2)	28.0 ± 17.2	0.062 (3/0/4)	35.9 ± 15.7	0.093 (0/1/6)	110.4 ± 51.8	0.089 (2/0/5)	102.6 ± 48.9	0.051 (6/0/1)	16.4 ± 20.3	0.047 **(7/0/0)**	2.2 ± 3.6
Overall ($W/T/L$)			**23/71/326**		**247/78/95**		**250/81/89**		**179/14/227**		**1/61/358**		**69/54/297**		**300/71/49**		**375/42/3**	
Win Rate			**5.5%**		**58.8%**		**59.5%**		**42.6%**		**0.2%**		**16.4%**		**71.4%**		**89.3%**	

ore flow blending. By integrating offline quality-diversity optimization with an online heterogeneous ant colony system, the framework effectively transforms diverse high quality heuristic rules into agile decision making strategies. This architecture successfully mitigates the semantic cloning issue of evolutionary algorithms and enables robust parallel exploration. Experimental results confirm that GP-ACO significantly outperforms state of the art baselines, demonstrating exceptional scalability and resilience against complex dynamic disturbances.

Future work will focus on two key directions: extending the framework to multi-objective optimization to simultaneously address grade stability and energy efficiency, and developing a joint shovel truck scheduling mechanism to achieve finer grained control over the entire production process.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62476006, 62503012, and 62306279, in part by the Hubei Provincial Natural Science Foundation of China under Grant 2023AFA049, and in part by the Fundamental Research Funds of the AUST under Grant 2024JBZD0007.

References

[1] Abdel Monaem FM AbdAllah, Daryl L Essam, and Ruhul A Sarker. 2017. On solving periodic re-optimization dynamic vehicle routing problems. *Applied Soft Computing* 55 (2017), 1–12.

[2] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. 2000. Inspiration for optimization from social insect behaviour. *Nature* 406, 6791 (2000), 39–42.

[3] Lu Chen, Qinghua Gu, Rui Wang, Zhidong Feng, and Chao Zhang. 2022. Comprehensive utilization of mineral resources: optimal blending of polymetallic ore using an improved NSGA-III algorithm. *Sustainability* 14, 17 (2022), 10766.

[4] Xinan Chen, Ruibin Bai, Rong Qu, Jing Dong, and Yaochu Jin. 2025. Deep Reinforcement Learning Assisted Genetic Programming Ensemble Hyper-Heuristics for Dynamic Scheduling of Container Port Trucks. *IEEE Transactions on Evolutionary Computation* 29, 4 (2025), 1371–1385.

[5] Daniela Dering and Christopher LE Swartz. 2022. A stochastic optimization framework for integrated scheduling and control under demand uncertainty. *Computers & Chemical Engineering* 165 (2022), 107931.

[6] Jianguo Duan, Fanfan Liu, Qinglei Zhang, Jiyun Qin, and Ying Zhou. 2024. Genetic programming hyper-heuristic-based solution for dynamic energy-efficient scheduling of hybrid flow shop scheduling with machine breakdowns and random job arrivals. *Expert Systems with Applications* 254 (2024), 124375.

[7] Agoston E Eiben and Jim Smith. 2015. From evolutionary computation to the evolution of things. *Nature* 521, 7553 (2015), 476–482.

[8] Mageed Ghaleb, Hossein Zolfagharinia, and Sharareh Taghipour. 2020. Real-time production scheduling in the Industry-4.0 context: Addressing uncertainties in job arrivals and machine breakdowns. *Computers & Operations Research* 123 (2020), 105031.

[9] Alper Hamzadayi and Gokalp Yildiz. 2016. Event driven strategy based complete rescheduling approaches for dynamic m identical parallel machines scheduling problem with a common server. *Computers & Industrial Engineering* 91 (2016), 66–84.

[10] Da Huo, Yuksel Asli Sari, and Qian Zhang. 2024. Smart dispatching for low-carbon mining fleet: A deep reinforcement learning approach. *Journal of Cleaner Production* 435 (2024), 140459.

[11] Wouter Kool, Herke van Hoof, and Max Welling. 2019. Attention, Learn to Solve Routing Problems!. In *International Conference on Learning Representations*. https://openreview.net/forum?id=ByxBFsRqYm

[12] Arie MCA Koster, Jenny Segschneider, and Nicole Ventsch. 2024. Γ-robust optimization of project scheduling problems. *Computers & Operations Research* 161 (2024), 106453.

[13] Fei Li, Qiang Yue, Yuanchao Liu, Haibin Ouyang, and Fangqing Gu. 2024. A fast density peak clustering based particle swarm optimizer for dynamic optimization. *Expert Systems with Applications* 236 (2024), 121254.

[14] Michalis Mavrovouniotis, Felipe M. Müller, and Shengxiang Yang. 2017. Ant Colony Optimization With Local Search for Dynamic Traveling Salesman Problems. *IEEE Transactions on Cybernetics* 47, 7 (2017), 1743–1756. doi:10.1109/TCYB.2016.2556742

[15] Michalis Mavrovouniotis and Shengxiang Yang. 2015. Ant algorithms with immigrants schemes for the dynamic vehicle routing problem. *Information Sciences* 294 (2015), 456–477.

[16] Michalis Mavrovouniotis, Shengxiang Yang, Mien Van, Changhe Li, and Marios Polycarpou. 2020. Ant colony optimization algorithms for dynamic optimization: A case study of the dynamic travelling salesperson problem [research frontier]. *IEEE Computational Intelligence Magazine* 15, 1 (2020), 52–63.

[17] Su Nguyen, Yi Mei, and Mengjie Zhang. 2017. Genetic programming for production scheduling: a survey with a unified framework. *Complex & Intelligent Systems* 3, 1 (2017), 41–66.

[18] Roberto Noriega, Yashar Pourrahimian, and Hooman Askari-Nasab. 2025. Deep Reinforcement Learning based real-time open-pit mining truck dispatching system. *Computers & Operations Research* 173 (2025), 106815.

[19] Marcella Samà, Paola Pellegrini, Andrea D'Ariano, Joaquin Rodriguez, and Dario Pacciarelli. 2016. Ant colony optimization for the real-time train routing selection problem. *Transportation Research Part B: Methodological* 85 (2016), 89–108.

[20] Lin Shi, Zhi-Hui Zhan, Di Liang, and Jun Zhang. 2022. Memory-based ant colony system approach for multi-source data associated dynamic electric vehicle dispatch optimization. *IEEE Transactions on Intelligent Transportation Systems* 23, 10 (2022), 17491–17505.

[21] Thomas Stützle and Holger H Hoos. 2000. MAX–MIN ant system. *Future generation computer systems* 16, 8 (2000), 889–914.

[22] Marko Đurasević and Domagoj Jakobović. 2018. A survey of dispatching rules for the dynamic unrelated machines environment. *Expert systems with applications* 113 (2018), 555–569.

[23] Vassilis Vassiliades, Konstantinos Chatzilygeroudis, and Jean-Baptiste Mouret. 2017. Using centroidal voronoi tessellations to scale up the multidimensional archive of phenotypic elites algorithm. *IEEE Transactions on Evolutionary Computation* 22, 4 (2017), 623–630.

[24] Meng Xu, Frank Neumann, Aneta Neumann, and Yew Soon Ong. 2025. Quality Diversity Genetic Programming for Learning Scheduling Heuristics. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1090–1098.

[25] Danial Yazdani, Ran Cheng, Cheng He, and Juergen Branke. 2020. Adaptive control of subpopulations in evolutionary dynamic optimization. *IEEE Transactions on Cybernetics* 52, 7 (2020), 6476–6489.

[26] Danial Yazdani, Ran Cheng, Cheng He, and Jürgen Branke. 2022. Adaptive Control of Subpopulations in Evolutionary Dynamic Optimization. *IEEE Transactions on Cybernetics* 52, 7 (2022), 6476–6489. doi:10.1109/TCYB.2020.3036100

[27] Fangfang Zhang, Yi Mei, Su Nguyen, Kay Chen Tan, and Mengjie Zhang. 2021. Multitask genetic programming-based generative hyperheuristics: A case study in dynamic scheduling. *IEEE Transactions on Cybernetics* 52, 10 (2021), 10515–10528.

[28] Fangfang Zhang, Yi Mei, Su Nguyen, Mengjie Zhang, and Kay Chen Tan. 2021. Surrogate-assisted evolutionary multitask genetic programming for dynamic flexible job shop scheduling. *IEEE Transactions on Evolutionary Computation* 25, 4 (2021), 651–665.