Computational Engineering and Technology Innovations

www.ceti.reapress.com

Comput. Eng. Technol. Innov. Vol. 3, No. 2 (2026) 142–160.

Paper Type: Original Article

Nature-Inspired Optimization Across the Mine Value Chain: A Critical Review

Khris Benjamin Swain[1,*], Elif Ceviz[2], Maduabuchi Okafor[3], Rina V. Petrova[4], Hiroshi Tanabe[5]

[1] Department of Computational Chemistry, School of Engineering & Applied Science, Heriot-Watt University, Edinburgh, United Kingdom; k.swain@hw-chem.ac.uk.

[2] Department of Chemical Engineering, Bogazici University, Istanbul, Turkiye.

[3] Department of Computational Science, University of Lagos, Akoka, Nigeria.

[4] Department of Computational Chemistry and Biochemistry, Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland.

[5] Department of Computational Science, RIKEN Center for Computational Science, Kobe, Japan.

Citation:

Received: 10 November 2025 *Revised: 06 January 2026* *Accepted: 03*	Swain, K. B., Ceviz, E., Okafor, M., Petrova, R. V., Tanabe, H. (2026). Nature-inspired optimization across the mine value chain: A critical review. *Computational Engineering and Technology Innovations*, *3*(2), 142-160.

Abstract

Nature-inspired optimization offers flexible search mechanisms for mining decisions that combine discrete schedules, nonlinear process responses and competing resource demands. Its practical value, however, depends on the formulation being solved and the evidence supporting the reported improvement. This critical narrative review examines these methods across fragmentation, extraction and blending, transportation, mineral processing, water management, and waste and closure planning. It connects foundational search mechanisms with selected mining studies and distinguishes predictive accuracy, improvement within a mathematical model, and operational validation. The synthesis identifies representation, constraint handling, cross-stage coupling, uncertainty and computational accounting as the principal dimensions of a defensible comparison. A compact mathematical framework makes material balance, recovery, resource limits and residue accumulation explicit. An exactly enumerated, dimensionless example then exposes the consequences of these couplings without claiming empirical algorithm superiority. Of 189 configurations, 14 satisfy all constraints and eight are nondominated in profit, energy and water. A policy that minimizes the fragmentation component at a fixed throughput violates the downstream energy limit, while an alternative fragmentation mode restores feasibility and improves the profit score. Capacity and penalty analyses further show that a higher unconstrained objective can reflect constraint violation and that the penalty weight needed to exclude infeasible optima depends on the instance and normalization. The review concludes with a comparison protocol, elementary arithmetic checks for rank-based results, and a research agenda centered on transparent hybridization, independent process validation and transferable benchmarks. Nature-inspired search is most persuasive when its computational mechanisms are matched to mine-chain structure and its benefits remain traceable to feasible, reproducible decisions.

Keywords: Mine value chain, Nature-inspired optimization, Metaheuristics, Stochastic mine planning, Multiobjective optimization, Constraint handling, Reproducible benchmarking.

1 | Introduction

A mining operation creates value through a sequence of connected transformations. Rock is fragmented and extracted, materials are blended and transported, valuable constituents are recovered, and residual streams

Corresponding Author: k.swain@hw-chem.ac.uk

https://doi.org/10.48314/ceti.v3i2.89

Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

require continuing management. A decision that improves one stage can alter the costs, capacities or risks of another. Lower expenditure on fragmentation, for example, can change the material delivered to the crusher; a processing policy that favors recovery can increase energy or water demand; and a production schedule can transfer liabilities to waste placement and closure. The operations-research literature established the breadth of these linked decisions [1], while simultaneous stochastic planning explicitly connected multiple components and uncertain material supply within mineral value chains [2], [3]. The relevant optimization problem is therefore determined by the interactions included in the model, rather than by the name of the search procedure.

Mining also presents a difficult computational setting. Extraction precedences create combinatorial structure, destinations and processing routes introduce discrete choices, and recovery and blending responses can be nonlinear. Geological uncertainty affects both quantity and quality, whereas equipment availability and operating conditions vary over shorter horizons. Engineering descriptions of the mining system provide the physical context for these decisions [4]. Metaheuristic methods supply flexible ways to explore difficult feasible sets [5], and mathematical programming supplies formulations, relaxations and bounds that can reveal what the search has actually achieved [6]. These traditions are complementary. A fast search is valuable when it addresses an important formulation and produces decisions whose feasibility can be independently established.

Interest in nature-inspired methods has expanded beyond scheduling to prediction-assisted design, process optimization and environmental applications. Tan et al. [7] reviewed algorithms alongside swarm robotics and wider biomimicry technologies. Lin et al. [8] focused on metaheuristic approaches to long-term open-pit production scheduling. These scopes are useful but answer different questions: A broad technology landscape does not establish comparative search performance, and a scheduling review does not by itself resolve the interfaces between fragmentation, processing and environmental constraints. The present review uses these interfaces as its organizing principle. It asks which decision is optimized, how an upstream action changes a downstream response, and what evidence is needed before a reported improvement can be transferred to another setting.

Three contributions follow from this perspective: 1) The review compares method families through their operators, representations and constraint-handling requirements, and relates those features to decisions across the mine chain, 2) it presents a common modeling vocabulary and an exact finite example that makes local-versus-coupled decisions, resource bottlenecks and penalty behavior inspectable, and 3) it specifies how future comparisons can separate feasibility, solution quality, computational effort and uncertainty, including arithmetic checks on reported average ranks and Friedman statistics. The numerical example is an explanatory reference instance with declared assumptions; it is not a field experiment or an empirical ranking of metaheuristics. The objective is to make the connection between algorithmic claims and engineering evidence sufficiently explicit to support better research design.

2|Review Scope and Development of the Field

This is a focused critical narrative review. Sources were selected to connect foundational optimization mechanisms, established mine-planning formulations, recent stage-specific applications and comparison methodology. Targeted searches combined terms for nature-inspired optimization, metaheuristics, mine production scheduling, blasting, flotation, electric transportation, mine water and reclamation. Reference tracing connected the mining-complex literature to its methodological foundations. Bibliographic and scope checks used publisher records, official journal sites, author-institution repositories and the official tailings-standard repository, with a final verification date of 22 October 2026. A source ledger accompanies the manuscript. The review does not claim an exhaustive database census, a bibliometric estimate of the field, or a systematic-review screening flow.

The selection favored sources that disclosed a recognizable decision problem or clarified a methodological issue central to evaluation. Books supplied foundational context; primary articles supplied specific formulations and applications; and existing reviews supplied research positioning. Where a publisher restricted full-text access, the synthesis was limited to the verified abstract, indexed methods description or author-deposited record. No inaccessible study supplied a pooled effect size or a numerical superiority claim. This boundary matters because a paper can establish a promising application without providing enough comparable information for a cross-paper ranking. Differences in input data, constraints, tuning, hardware and stopping criteria prevent an unqualified aggregation of published objective values.

The historical sequence in *Table 1* reflects changes in research emphasis, rather than a claim that one generation displaced another. Kennedy and Eberhart [9] introduced particle swarm optimization as a population-based search guided by individual and social information. Dorigo et al. [10] developed the ant system around probabilistic construction and shared memory. Storn and Price [11] introduced differential evolution for continuous nonlinear search, while Kalyanmoy [12] developed NSGA-II to combine nondominated selection with diversity preservation. These methods offered reusable computational mechanisms. Their relevance to mining depended on how those mechanisms were connected to schedules, routes, destinations or process controls.

Table 1. Selected developments linking search mechanisms to mine-chain questions.

Development	Representative Sources	Change in the Research Question
1995–2002: Reusable search mechanisms	PSO [9], ant system [10], DE [11], NSGA-II [12]	How do information sharing, construction, mutation and nondominated selection organize search?
2003–2012: Engineering structure and bounds	Metaheuristic handbook [5]; mine-planning review [1]; structured scheduling [13]	Which decisions, precedences and capacity limits must a useful search preserve?
2015–2017: Coupled stochastic planning	Montiel and Dimitrakopoulos [14]; Goodfellow and Dimitrakopoulos [2]	How should processing routes, blending and uncertain supply be optimized together?
2022–2024: Adaptive and surrogate-assisted search	Integrated overview [3]; learn-to-perturb [15]; flotation and blasting studies [16–18]	Can learned responses and adaptive neighborhoods reduce search effort without weakening validation?
2025–2026: Broader integration and environmental decisions	Biomimicry review [7]; OPPS review [8]; progressive reclamation [19]	How can search support connected production, resource use and life-cycle liabilities?

The terminology of nature inspiration can obscure this connection. Lones [20] examined the underlying metaheuristic components used by nature-inspired algorithms, and Sörensen [21] criticized the proliferation of metaphor-led methods that did not establish a substantive computational contribution. Their arguments suggest a practical reading rule: The scientifically useful description identifies the state representation, candidate-generation operator, information retained between iterations, selection rule and constraint treatment. A new biological analogy is insufficient to explain why a mine schedule became better. Conversely, an established search mechanism can make an important contribution when it supports a more faithful formulation or an effective representation of engineering constraints.

The distinction is visible in *Table 2* Genetic search can manipulate destination assignments or priority vectors, but the decoder determines whether those vectors become valid extraction schedules. A particle swarm can optimize continuous setpoints, but categorical route choices require an explicit mapping. Ant-colony construction is naturally compatible with some graph decisions, although graph compatibility alone does not enforce charging, availability or capacity constraints. Differential evolution can exploit continuous structure,

but mixed categorical variables and incompatible scales require additional design. Pareto selection exposes trade-offs only after objectives and feasibility have been defined consistently. Constraint handling also has its own methodological history: Deb [22] examined feasibility-oriented selection as an alternative to conventional penalty tuning. These observations concern the match between a method and a problem; they are not a universal ordering of method families.

Table 2. Comparison of computational mechanisms and the mining-specific questions they raise.

Method Family	Computational Mechanism and Suitable Representation	Mining-Specific Test of Usefulness
Genetic search	Selection and variation on schedules, destination codes or bounded controls; a decoder can enforce precedences.	Separate the effects of encoding, repair and local improvement; compare with a simple constructive schedule [5].
Particle swarm	Personal and shared information guide positions in a numerical search space.	For discrete choices, disclose the position-to-decision mapping and account for rejected or repaired moves [9].
Ant colony	Probabilistic construction uses pheromone memory and problem information on a graph.	Demonstrate feasible route or dispatch construction under changing availability and energy limits [10], [23].
Differential evolution	Difference-based variation acts on numerical candidate vectors.	Respect units, scaling and categorical variables; retain a continuous or appropriately decoded decision interpretation [11].
Pareto evolutionary search	Nondominated selection and diversity preservation approximate trade-offs.	Use fixed objective scaling and reference points; retain only independently verified feasible solutions [12].
Hybrid and adaptive search	Local search, mathematical programming, surrogates or learned operator selection augment a search engine.	Use ablations and count training, tuning, exact-solver and simulator work in the total cost [20], [15].

A second development concerned the scope of the formulation. Montiel and Dimitrakopoulos [14] modeled multiple processing and transportation alternatives together under uncertainty. Goodfellow and Dimitrakopoulos [2] developed simultaneous stochastic optimization across mining-complex components. The later overview by Dimitrakopoulos and Lamghari [3] explained the associated integration of uncertain supply, processing transformations and market-facing decisions. The central advance in this line was not simply a different search metaphor: It was the ability to preserve interactions that sequential planning could suppress. Attributing every resulting gain to the optimizer would confound the value of integration with the value of the search mechanism.

Structured optimization remained a necessary comparator. Chicoisne et al. [13] exploited the structure of open-pit scheduling through relaxation, decomposition, rounding and local improvement. Their work illustrated why a useful baseline can include a mathematical bound and a tailored heuristic rather than only other population methods. More recently, Yaakoubi and Dimitrakopoulos [15] used learning to adapt the selection of perturbation heuristics within a simulated-annealing framework. This moved adaptation toward the choice of neighborhoods. It also created an evaluation obligation: Learning effort, transfer to new instances and the contribution of individual components must be distinguished from the performance of the complete hybrid.

3|Decisions and Couplings Across the Mine Value Chain

Fig. 1 summarizes the main structural argument. Material properties and decisions propagate downstream, while water, energy and storage limits constrain several stages at once. Information and resource feedback mean that the chain cannot always be represented as a sequence of independent optima. Nevertheless, simultaneous optimization is not automatically preferable in every application. Its additional coupling must be supported by a response model, consistent data and an operational reason to coordinate the stages. Where interfaces are weak or poorly characterized, a simpler model with explicit sensitivity analysis can be more informative than an elaborate model whose coefficients are unvalidated.

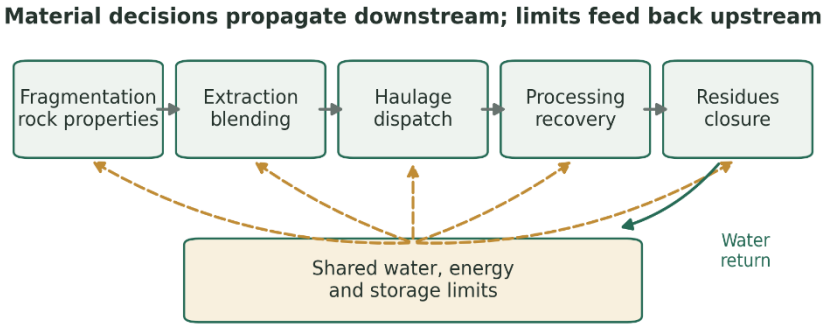


Fig. 1. Conceptual mine-chain dependencies.

Solid horizontal arrows indicate downstream material decisions. dashed arrows indicate shared limits. The water-return arrow represents feedback from residue management to the common resource system. This is a modeling schematic, not a process-flow design for a particular mine.Fragmentation provides an early example of these issues. A prediction model can relate design variables and rock characteristics to a fragmentation indicator. A search procedure can then optimize the predicted response jointly with cost. Guo et al. [16] combined genetic-algorithm tuning of least-squares support-vector prediction with multiobjective particle-swarm optimization of blasting parameters. Such a workflow links prediction and decision support, but it contains several distinct tests. The predictor must generalize to relevant rock conditions; the optimizer must respect its stated domain and constraints; and the selected design must be evaluated against the physical response. A low prediction error averaged over historical observations does not ensure accuracy at the extreme operating point chosen by an optimizer.

The downstream interface is equally consequential. An optimization that values a particular fragmentation indicator should explain how that indicator affects loading, crushing or processing. A size target may be only an intermediate proxy for the economic decision. If no validated response function connects fragmentation to downstream performance, the analysis should report sensitivity to alternative connections instead of presenting a unique chain-wide optimum. This requirement also helps separate model development from engineering prescription. Search outputs become actionable only after their domain, constraints and response assumptions have been established for the operating setting.

Extraction and blending introduce a different representation problem. A feasible schedule must respect access and precedence relationships, temporal capacities and destinations. Blending requires conservation of material and its attributes; averaging grades without mass weights can make a nominally attractive feed physically incorrect. Stockpiles add memory because their inventories and compositions depend on previous decisions. Montiel and Dimitrakopoulos [14] included processing, stockpiling and transport alternatives in a common stochastic formulation. This provides a stronger basis for discussing integration than a list of algorithms assigned to isolated stages. It also shows why geology and destination decisions should be held constant when the performance contribution of a solver is being isolated.

Transportation adds time, congestion and resource availability to the material-flow picture. A route can be short and still be unsuitable when a vehicle cannot complete it within its battery range or when the destination queue becomes excessive. Zhang et al. [23] addressed underground electric transportation with an improved ant-colony method incorporating learning and local search. Their application illustrates the need to encode transport-specific constraints within candidate construction or subsequent validation. The research question is broader than whether one algorithm found a shorter path: It includes whether the resulting schedule remains feasible under the modeled energy demand, vehicle availability and system interactions. Transfer to another mine requires a new assessment of these conditions.

Mineral processing often combines continuous controls with nonlinear and partly observed responses. Amankwaa-Kyeremeh et al. [17] used Gaussian-process predictive objectives to examine the recovery–grade trade-off with controlled-NSGA-II and paretosearch. Hasidi et al. [18] combined neural-network-based monitoring with genetic-algorithm optimization of flotation settings. Both approaches made predictive modeling part of the decision pipeline. Their evidence must be interpreted at the appropriate level: A surrogate supports a modeled trade-off, and industrial monitoring observations support the monitoring model, whereas simulated setpoint optimization does not automatically establish an implemented operational gain.

This distinction becomes sharper when an optimizer repeatedly interrogates a surrogate. The search can concentrate on regions where model error is favorable to the objective. Consequently, uncertainty should be evaluated at optimizer-selected points and under relevant shifts in feed properties, not only across a random test split. Process constraints also need an independent check. Predicted recovery, grade and throughput should jointly describe a possible material transformation. When concentrate quality is an objective, its relationship to recovered metal and total concentrate mass must be made explicit. Otherwise, improvements in one reported indicator can conceal an inconsistent mass balance or an unreported loss in another.

Water management couples quantity, quality and treatment effort. Liu et al. [24] formulated cooperative mine-water scheduling with water-quality and water-quantity constraints and a hybrid genetic/simulated-annealing approach. The relevant interface is the allocation of suitable water to individual users, rather than only the reduction of total freshwater volume. Recirculation can change process-water characteristics and energy requirements. A water-saving result therefore needs a clear boundary: Withdrawal, consumption, treatment, storage and return are different quantities. A model that tracks only total reuse should not be interpreted as having established water-quality suitability for every downstream operation.

Waste, tailings and closure decisions extend the planning horizon further. Freire Guimaraes and Dimitrakopoulos [19] integrated waste management and progressive reclamation through encapsulation within stochastic mining-complex optimization. This line of research treated waste properties and placement as decisions with consequences for future liabilities. Tailings-facility safety has a wider scope than a production-scheduling constraint. The Global Industry Standard on Tailings Management addressed responsibility, information disclosure and facility management across the life cycle [25]. Meeting a storage cap in an optimization model does not establish the stability of a facility or satisfaction of its broader obligations.

Tables 3 and *4* connect these stages to their evidence boundaries. They deliberately compare problem structure and validation needs rather than pooling reported percentage improvements. The selected studies used different objectives and engineering settings, and their complete computational costs were not harmonized into a common experiment here. A useful cross-study conclusion is therefore conditional: A mechanism is promising for a particular representation, a model preserves a particular interaction, or a study validates a particular stage of the decision pipeline. A claim of general superiority would require a different study design with common instances, budgets and feasibility rules.

Table 3. Decisions, cross-stage effects and the constraints that govern interpretation.

Stage and Decision	Downstream Coupling	Constraint or Evidence that Changes Interpretation
Fragmentation: Design mode and target size distribution	Diggability, crusher capacity, grinding energy and feed variability	A predicted size reduction does not establish a chain-wide economic gain. Validate downstream response [16].
Extraction and blending: Timing and destinations	Feed quantity, grade, mineralogy and stockpile inventories	Preserve precedences, selective extraction and mass-weighted quality; test geological scenarios [1], [14].
Transport: Dispatch, route and charging	Plant starvation, queues, cycle time and energy demand	Vehicle availability, road restrictions and battery range must survive stochastic evaluation [23].
Processing: Intensity, routing and setpoints	Recovery, concentrate quality, water and residue generation	Check surrogate validity outside historical operating points and distinguish simulation from deployment [17], [18].
Water: Allocation, treatment and return	Freshwater demand, recirculation quality and energy	Close quantity and quality balances across users; a reuse objective alone is insufficient [24].
Waste and closure: Placement and timing	Storage demand, haulage, geochemistry and future liabilities	Separate planning constraints from facility safety, governance and post-closure obligations [19], [25].

Table 4. Selected studies and the limits of transferring their evidence across problems.

Study	Verified Model or Contribution	Boundary for Cross-Study Comparison
Chicoisne et al. [13]	LP relaxation decomposition, rounding and local improvement for open-pit scheduling.	A structured baseline; guarantees and bounds depend on the stated formulation.
Montiel and Dimitrakopoulos [14]	Stochastic MINLP and metaheuristic solver joined processing and transportation alternatives.	Benefits combine integration, uncertainty treatment and search; these effects require separate comparisons.
Goodfellow and Dimitrakopoulos [2]	Simultaneous stochastic optimization linked extraction, blending and downstream transformations.	Case-study evidence supports the modeled setting; it does not rank all metaheuristic families.
Yaakoubi and Dimitrakopoulos [15]	Learned heuristic selection adapted neighborhoods within simulated annealing.	Include learning cost and assess transfer to held-out mining complexes.
Guo et al. [16]	GA-tuned least-squares support-vector prediction supported multiobjective blast optimization.	Prediction, optimizer behavior and field validation answer different questions.
Amankwaa-Kyeremeh et al. [17]	Gaussian-process objectives supported controlled-NSGA-II and paretosearch for flotation.	A surrogate Pareto front needs validation against the underlying process.
Hasidi et al. [18]	Neural-network monitoring and GA-based flotation optimization were combined.	Industrial monitoring data and simulated optimization outcomes have distinct evidential roles.
Zhang et al. [23]	Improved ACO with learning and local search addressed electric underground transport.	Assess individual components, energy constraints and instance dependence before transfer.

4 | A Coupled Formulation and an Exact Reference Example

4.1 | Modeling the Interfaces

A compact formulation helps identify what a comparison must keep fixed. Let v denote the components of a material network, t a planning period and ω a scenario. Let $I_{vt\omega}$ be bulk inventory, $s_{vt\omega}$ supplied material, $d_{vt\omega}$ terminal withdrawals, and $f_{uvt\omega}$ a nonnegative flow from component u to component v. The sets A_v^- and A_v^+ contain the incoming and outgoing neighbors of v. A bulk balance takes the form

$$I_{vt\omega} = I_{v,t-1,\omega} + s_{vt\omega} + \sum_{u \in A_v^-} f_{uvt\omega} - \sum_{u \in A_v^+} f_{vut\omega} - d_{vt\omega} \geq 0. \tag{1}$$

Processing nodes must account for every output stream, including residues, within this balance. Attribute balances then track valuable constituents and other relevant properties. For a blend with positive total feed $Q = \sum_i q_i$, feed grade G and recovered valuable mass M can be written as

$$G = \frac{\sum_i g_i q_i}{Q}, M = \eta(G, z, \omega) G Q, Q > 0. \tag{2}$$

Here g_i is the valuable-material fraction of source i, q_i its feed mass, and η the recovery fraction under process controls z. Zero-feed periods require a separate zero-production definition. Recovered valuable mass is not generally the total mass of concentrate; a concentrate-grade model is needed to connect the two. These bookkeeping distinctions prevent a search from exploiting an invalid transformation.

For a decision policy x, a standard economic objective is expected discounted net value,

$$\mathrm{NPV}(x) = \sum_{\omega \in \Omega} \pi_\omega \sum_{t=1}^{H} \frac{p_{t\omega} M_{t\omega}(x) - C_{t\omega}(x) - L_{t\omega}(x)}{(1+\delta)^t}. \tag{3}$$

The probabilities satisfy $\pi_\omega \geq 0$ and $\sum_\omega \pi_\omega = 1$. The horizon is H, the discount rate is δ, p is the unit value of recovered material, C includes the modeled operating and capital charges, and L represents separately identified rehabilitation or liability costs. Common decisions must obey nonanticipativity across scenarios sharing the same information. A policy that chooses extraction after seeing future geology solves a different, more informed problem. The stochastic integration literature supplied a basis for treating these interfaces explicitly [2], [3]; the equations here state common modeling requirements rather than a new mine-planning theorem.

Resource constraints should remain visible even when an environmental cost appears in the objective. For energy use E, freshwater demand W, residue inventory J, generated residue T and an authorized removal or retreatment flow D, representative restrictions are

$$\begin{aligned} &E_{t\omega}(x) \leq \bar{E}_t, W_{t\omega}(x) \leq \bar{W}_t. \\ &0 \leq J_{t\omega}(x) = J_{t-1,\omega}(x) + T_{t\omega}(x) - D_{t\omega}(x) \leq \bar{J}_t. \end{aligned} \tag{4}$$

Removed material must have an accounted destination. Mass, volume and storage geometry require a validated conversion when facility capacity is expressed volumetrically. The feasible set $\mathcal{X}_f$ also includes precedences, equipment limits, quality restrictions and the appropriate engineering constraints. If total expected energy and water use are denoted by E_Σ and W_Σ, an illustrative multiobjective problem is

$$\min_{x \in \mathcal{X}_f} f(x) = (-\mathrm{NPV}(x), E_\Sigma(x), W_\Sigma(x)). \tag{5}$$

A feasible decision dominates another when it is no worse in every component and strictly better in at least one. A Pareto set therefore describes a trade-off conditional on the chosen objectives, feasible set and uncertainty model. It does not remove the need to select an operating policy or validate the response

functions. A weighted scalar objective can be useful for a stated preference, but its optimum should not be described as the complete trade-off frontier, particularly when decisions are discrete.

4.2 | Exact Finite Instance

The following example isolates cross-stage coupling in a form that can be checked completely. All quantities, coefficients and limits are assumed dimensionless values. The instance is an explanatory decision model, not measured mine data, a calibrated process simulator or an engineering design recommendation. Throughput q, fragmentation mode b, processing intensity p and recirculation mode r have domains

$$q \in \{60,70,80,90,100,110,120\}, b, p, r \in \{0,1,2\}. \quad (6)$$

The assumed recovery fraction R, energy score E and freshwater score W are

$$(R = 0.76 + 0.035b + 0.025p. \quad (7)$$

$$E = q(1.30 - 0.18b + 0.08p + 0.04r). \quad (8)$$

$$W = q(0.90 - 0.20r + 0.03p). \quad (9)$$

Higher fragmentation mode increases its direct cost but reduces the downstream energy coefficient and increases recovery. Processing intensity increases recovery and resource requirements; recirculation reduces freshwater demand while consuming additional energy and cost. These relationships are declared assumptions, not fitted causal effects. An idealized pure recovered product P, residue mass T and total cost score C are defined by

$$P = 0.12qR, T = q - P. \quad (10)$$

$$C = q(1.25 + 0.12b + 0.08p + 0.10r) + 0.45. \quad (11)$$

The factor 0.12 is the assumed valuable-material fraction. Thus P represents recovered valuable mass and T includes gangue and unrecovered valuable material. The pure-product assumption makes the bulk balance transparent; a real concentrate would require a separate concentrate-grade relation. The profit score F and hard limits are

$$F = 50P - C. \quad (12)$$

$$E \leq 125, W \leq 65, T \leq 95, P \geq 8.4. \quad (13)$$

$$q \leq 110 + 8b - 6p. \quad (14)$$

The final inequality is an assumed throughput-capacity response. It represents another coupling: More intense processing reduces capacity, while improved fragmentation increases it. The residue limit is a planning capacity in this idealized mass model and has no geotechnical interpretation. *Table 5* states the complete enumeration protocol.

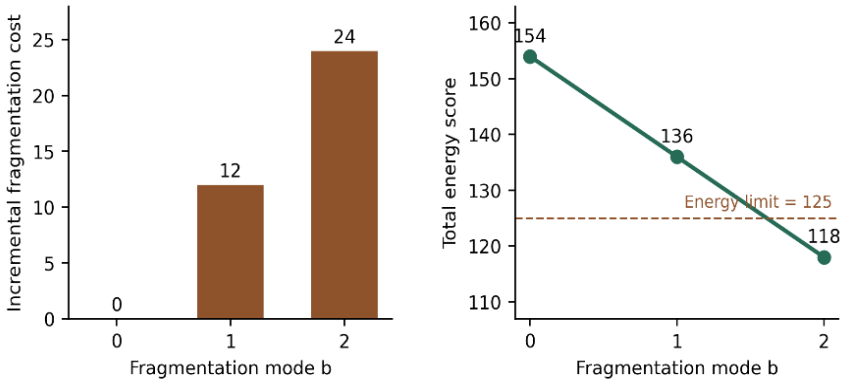
Table 5. Reference-instance specification and computational boundary.

Element	Declared Specification
Decision domains	Seven throughput levels from 60 to 120 in increments of 10; three ordered modes each for fragmentation, processing and recirculation (see Eq. (6)).
Response assumptions	Eqs. (7)–(9) define recovery, energy, water, product, residue, cost and profit. Every coefficient is assumed and dimensionless.
Hard limits	Five restrictions on energy, water, residue, minimum recovered product and throughput, stated in Eq. (9).
Reference calculation	Enumerate all 189 configurations with rational arithmetic. Apply all five limits before selecting profit or Pareto solutions.
Objectives	Maximize profit; for trade-off analysis minimize energy and water simultaneously with profit maximization.
Released records	All configurations, 14 feasible points, 8 Pareto points, one-at-a-time capacity sweeps, penalty sweep and executable evaluator.

There are $7 \times 3^3 = 189$ possible configurations. Exact rational arithmetic is used for response calculations, constraint comparisons, dominance tests and the penalty threshold below; decimal output is provided for readability. Fourteen configurations satisfy all five limits. Eight are nondominated when profit is maximized and energy and water are minimized. Enumeration supplies the exact finite reference set, so these counts and objective values do not depend on random seeds or a search algorithm. The companion records contain every configuration, including infeasible candidates, rather than only the selected solutions.

At fixed $q = 90$, $p = 2$ and $r = 2$, the lowest fragmentation mode $b = 0$ minimizes the direct fragmentation-cost component. It yields $E = 138.6$, which exceeds the energy limit of 125, and $F = 230.13$. Setting $b = 2$ at the same throughput and controls yields $E = 106.2$ and $F = 260.91$ while satisfying every limit. The direct fragmentation charge rises by 21.6, but energy falls by 32.4 and recovered product rises by 0.756. At the stated coefficients, the resulting profit gain is 30.78. This is a paired comparison of two policies, not a claim that the second policy is globally optimal at that throughput.

Fig. 2 makes the same interface visible at $q = 100$, $p = 2$, $r = 2$. A higher direct fragmentation charge accompanies lower energy demand; only mode $b = 2$ meets all constraints at that throughput and those controls. The example does not prove that increasing fragmentation effort always reduces real operating costs. It demonstrates why a stage-specific cost objective can select an infeasible or inferior chain-level decision when its downstream effects are omitted.



a. **b.**

Fig. 2. Declared fragmentation–energy coupling at fixed throughput 100, processing mode 2 and recirculation mode 2; a. upstream cost, b. downstream response.

Incremental fragmentation cost is measured relative to mode 0. The responses are assumptions of the dimensionless reference instance; mode 0 also violates the throughput-capacity constraint at these fixed settings. The maximum feasible profit is 289.90 at $(q, b, p, r) = (100,2,2,2)$, with energy 118 and water 56. The minimum-water feasible configuration uses throughput 80 and attains water 44.8 with profit 231.92. The minimum-energy configuration uses $(90,2,0,1)$ and attains energy 88.2 with profit 265.41. *Table 6* and *Fig. 3* show why a single best-profit number does not describe the entire decision problem. Several feasible configurations remain attractive under different preferences, and the minimum-energy choice does not coincide with the minimum-water choice. No objective weights were fitted to make one policy appear universally preferable.

Table 6. Selected exact configurations.

Configuration	(q, b, p, r)	Energy	Water	Profit	Feasibility
Stage-local	(90, 0, 2, 2)	138.60	50.40	230.13	energy
Coupled, same q	(90, 2, 2, 2)	106.20	50.40	260.91	Feasible
Minimum water	(80, 2, 2, 2)	94.40	44.80	231.92	Feasible
Minimum energy	(90, 2, 0, 1)	88.20	63.00	265.41	Feasible
Maximum feasible profit	(100, 2, 2, 2)	118.00	56.00	289.90	Feasible
Unconstrained maximum	(120, 2, 2, 0)	132.00	115.20	376.20	energy, water, tailings, throughput

Values are dimensionless model scores; infeasibility lists identify violated constraints.

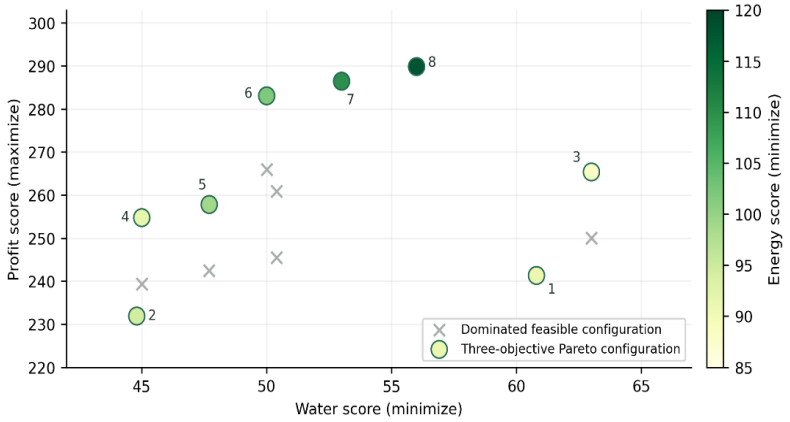


Fig. 3. All 14 feasible configurations projected onto water and profit.

Color shows energy, and numbered circles identify the eight nondominated configurations in the three-objective space. Numbers follow the row order of pareto_configurations.csv. Crosses denote dominated feasible configurations; a two-dimensional projection alone does not determine three-objective dominance. Capacity sensitivity changes the decision set itself. *Fig. 4* varies one limit at a time, holding the remaining limits at their nominal values and enumerating all configurations again. The best feasible profit is nondecreasing as any upper resource limit is relaxed, because the feasible sets are nested. The step pattern follows from the finite decision domains. Plateaus indicate that another restriction or the available domain prevents additional value; they do not imply a continuous engineering shadow price. The water limit must reach 44.8 before any configuration can satisfy all other nominal requirements. These observations are exact for this instance, while their numerical magnitudes have no external calibration.

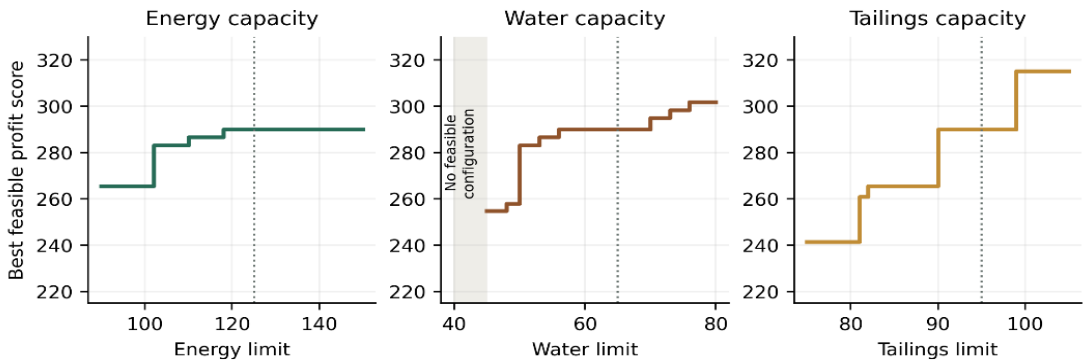


Fig. 4. Sensitivity of exact best-feasible profit to integer-valued energy, water, and residue limits.

Dotted lines mark nominal values, while the shaded water region indicates infeasibility.

4.3 | Why a Penalty Score Does Not Establish Feasibility

Constraint handling can reverse a reported comparison. Let $g_\ell(x) \le 0$ denote each of the five limits after bringing all terms to the left, and let $s_\ell > 0$ be declared scaling constants. A normalized violation and penalized profit are

$$V(x) = \sum_{\ell=1}^{5} \frac{\max\{0, g_\ell(x)\}}{s_\ell}, S_\lambda(x) = F(x) - \lambda V(x). \quad (15)$$

For this calculation, the scales are 125, 65, 95, 8.4 and 110 for energy, water, residue, minimum product and throughput respectively. The throughput scale is a fixed reference value, not the decision-dependent capacity. A finite penalty allows an objective improvement to compensate for a violation. Deb [22] developed a feasibility-oriented alternative to tuning conventional penalty parameters. Whatever approach is adopted, the final decision should be checked against the original constraints rather than accepted because its penalized score is high.

The reference instance permits an exact statement about the penalty. Let F^* be the maximum feasible profit. Among infeasible configurations that have higher unpenalized profit, define

$$\lambda^* = \max_{\substack{x \notin \chi_f \\ F(x) > F^*}} \frac{F(x) - F^*}{V(x)} = \frac{1199375}{1857} \approx 645.867. \quad (16)$$

For every $\lambda > \lambda^*$, every infeasible candidate scores strictly below a best feasible candidate. This follows directly by rearranging $F(x) - \lambda V(x) < F^*$; infeasible candidates with $F(x) \le F^*$ also satisfy the strict inequality because $V(x) > 0$ and $\lambda > 0$. At equality, a tie can remain. At weight 500, the maximizer is still infeasible: It has throughput 110, residue 98.714 and profit 315.15. At weight 1000, the penalized maximizer is the feasible profit optimum. This threshold is specific to the finite instance, profit units and violation scales. It is neither a universal penalty recommendation nor a convergence guarantee for a heuristic solver.

5 | What a Defensible Algorithm Comparison Must Establish

5.1 | Feasibility, Effort and Decision Quality

A comparison should identify its target before selecting algorithms. Beiranvand et al. [26] reviewed best practices for optimization comparisons, and Bartz-Beielstein et al. [27] organized benchmarking around problem specification, metrics, experimental design and reproducibility. For mining, the target can be a scheduling formulation, a simulation-based operating policy, a surrogate optimization pipeline or an adaptive controller. Each requires a different instance description and validation boundary. Comparing their headline percentage gains without reconstructing those boundaries creates an apparent ranking from incompatible experiments.

Table 7 specifies a practical comparison design. A constructive policy and local improvement provide a transparent starting point; a suitable mathematical-programming approach can supply feasible incumbents or bounds. Population methods should receive comparable tuning opportunities and evaluation budgets. Iterations alone are not a common currency: A population iteration may require many objective evaluations, and a repair routine or a scenario-based objective can conceal substantial work. Report objective and simulator calls, scenario replications, wall time, hardware and parallelism. Offline model fitting, hyperparameter search and learning costs should be disclosed separately from online decision latency. Their relevance depends on how often a trained system will be reused.

Table 7. Controls for interpretable mining-optimization comparisons.

Comparison Element	Required Control	Interpretive Purpose
Instances and scenarios	Separate development and evaluation instances; disclose geology, topology, capacities and scenario construction.	Tests transfer beyond an advantageous single configuration.
Baselines	Include a constructive policy, local search and a suitable mathematical-programming or bound-producing method.	Distinguishes useful search from weak comparator selection.
Computation	Match objective/simulator calls; report wall time, hardware, parallelism and all tuning or training work.	Prevents an iteration count from concealing unequal effort.
Constraint handling	Disclose residuals, tolerances, repairs and rejected candidates; validate final physical balances independently.	Separates nominal objective quality from deployable feasibility.
Randomness and uncertainty	Nest seeds within instances; vary process scenarios separately; provide interval estimates and paired instance effects.	Avoids treating repeated runs as independent mines.
Reporting and release	Publish run-level outcomes, failures, time-to-target traces, settings, code version and data provenance.	Allows recomputation of summaries, ranks and statistical tests.

Feasibility is the first reported outcome. For N completed evaluation runs with returned candidates x_a, a feasibility rate can be defined as

$$FR = \frac{1}{N}\sum_{a=1}^{N} 1\{x_a \in x_f\}. \quad (17)$$

The denominator must include failed and timed-out runs according to a predeclared reporting rule. Residuals, tolerances and missing outputs should remain visible; silently discarding failures biases the apparent quality of the method. When a repair changes a candidate, both the repair rule and the cost of producing the repaired solution belong to the method. Independent feasibility checks should use the original physical and operational constraints. Feasibility with respect to a surrogate or an approximate decoder can differ from feasibility in the full evaluation model.

For feasible scalar outcomes on instance i, a normalized reference gap can be written as

$$G_i = \frac{F_i^{ref} - F_i}{s_i}, s_i > 0. \quad (18)$$

The scale s_i must be fixed before inspecting the evaluated methods and expressed in compatible objective units. It should not be silently replaced by a vanishing reference objective. An exact optimum, a bound and a best-known feasible value have different meanings and must be labeled accordingly. If the reference is a feasible incumbent, a negative gap may indicate that the method improved it. For multiobjective comparisons, report the feasible approximation set and justify objective normalization and any reference point used by an indicator. An improved aggregate indicator can coexist with a loss in a decision-relevant region of the frontier.

Dolan and Moré [28] proposed performance profiles to summarize comparative performance across problems. In a mining study, an analogous summary can be useful when the target, failure handling and denominator are clearly specified. Time-to-target and quality-versus-budget curves can reveal practical differences hidden by a final mean objective. They should be accompanied by instance-level results, because an aggregate curve can conceal failure on a difficult geology, bottleneck or operating regime. Where an exact

or tightly bounded small instance is available, it can validate code and feasibility logic before larger experiments are interpreted.

Random seeds address search variability; they do not create new mines. Runs on one instance share its geology, topology, constraints and response functions. Consequently, seeds should be nested within instances, while scenario replications used to estimate an uncertain objective should be tracked separately. A study claiming transfer should evaluate held-out instances or operating regimes, with tuning isolated from that evaluation. Learning-assisted methods additionally need a clear separation between training data and decision-quality assessment. The evaluation hierarchy in *Fig. 5* makes these dependencies explicit and preserves the path from a run-level record to a general conclusion.

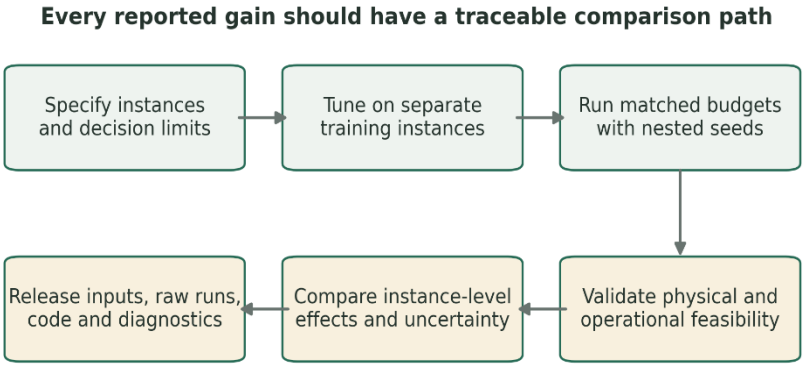


Fig. 5. A traceable comparison workflow.

Repeated seeds are nested within declared instances, and feasibility is assessed before objective comparisons are generalized. Releasing the underlying records allows independent checks of summary tables and uncertainty estimates.

5.2 | Arithmetic Checks Before Statistical Interpretation

Rank-based methods can support comparisons across heterogeneous instances, but their inputs must be internally possible. Demšar [29] discussed Friedman and related procedures for complete multi-method comparisons across data sets. The same arithmetic can check an optimization table before any inferential conclusion is drawn. Let r_{ij} be the rank of algorithm j on instance i, with k algorithms and n complete instances. If each row contains the ranks 1 through k, or tied positions receive their corresponding midranks, then

$$\sum_{j=1}^{k} r_{ij} = \frac{k(k+1)}{2}, \sum_{j=1}^{k} \bar{r}_j = \frac{k(k+1)}{2}. \tag{19}$$

The second identity follows by averaging the first over instances. It is an arithmetic invariant, not a hypothesis test. If displayed average ranks violate it beyond what documented rounding permits, the reported table cannot arise from the stated complete ranking procedure. Missing methods, unequal instance subsets or another ranking convention require an explicitly different analysis. They cannot be repaired by applying a complete-block test to incompatible inputs.

For complete rankings without ties, the uncorrected Friedman statistic is

$$Q_F = \frac{12n}{k(k+1)} \sum_{j=1}^{k} \left(\bar{r}_j - \frac{k+1}{2} \right)^2 , 0 \le Q_F \le n(k-1). \quad (20)$$

The upper bound has a short algebraic explanation. Every centered untied rank row is a permutation of the same vector and has squared norm $k(k^2-1)/12$. The squared norm of their average is no greater than the average squared norm. Substituting this inequality into the displayed expression yields $n(k-1)$. For example, a table with 12 complete instances and eight algorithms must have average ranks summing to 36, and its untied statistic cannot exceed 84. These values illustrate the check; they are not results from an algorithm experiment. Ties require the appropriate correction, and the corrected procedure must be reported separately.

Passing these checks establishes arithmetic consistency only. It does not establish independence, adequate power, a practically important effect or an appropriate sampling unit. García and Herrera [30] examined adjusted inference for multiple pairwise comparisons. A mining study should predeclare the comparison family, identify its control or all-pairs question, report adjusted probabilities when required, and retain effect sizes and uncertainty alongside ranks. A collection of many seed-level observations from a few instances should not be presented as equally many independent operating environments. The most useful statistical statement remains tied to a defined decision problem, feasible outcomes and a meaningful magnitude of improvement.

6 | Research Agenda and Implications

The strongest opportunity is to connect better representations with better evidence. A first priority is to distinguish the benefit of integration from the benefit of the solver. A controlled study can hold geological inputs, response functions and feasible decisions fixed while comparing sequential and simultaneous policies. A second comparison can then hold the integrated formulation fixed while changing the search procedure. This separation would identify whether a gain came from preserving a cross-stage interaction, finding a better solution to the same problem, or both. The exact reference example illustrates the first issue, but its small domain and assumed responses cannot establish performance on industrial-scale planning problems.

A second priority is transparent hybridization. A learned neighborhood selector, local improvement routine, constraint-aware decoder or exact subproblem solver can each improve a method. Their combined performance, however, does not reveal which component was necessary. Ablations should remove one component at a time while keeping the total evaluation effort comparable. Training and inference should also be separated. A method that pays a large offline cost can be attractive for frequent reuse, while a simpler procedure may suit a one-off planning problem. The decision context determines how these costs should be interpreted; a lower iteration count alone cannot resolve the trade-off.

A third priority concerns surrogate reliability. Optimizer-selected points deserve a stronger validation test than a general predictive summary. Held-out time periods, geological regimes and operating ranges can expose the conditions under which a response model ceases to be useful. Uncertainty diagnostics should accompany independent evaluations of proposed settings, and the permitted search domain should follow the evidence supporting the surrogate. The relevant outcome is the quality and feasibility of the resulting decision under model error. A reduction in prediction error is valuable evidence, but it answers a narrower question.

A fourth priority is environmental integration with explicit scope. Water quantity, water quality, waste properties, placement, energy and progressive reclamation affect different parts of a mine's operation and life cycle. A study can begin with a transparent proxy when data are limited, provided the proxy is identified and its sensitivity is examined. Stronger conclusions require validated response models and appropriate engineering assessment. A low tailings-mass objective or a satisfied storage bound should not stand in for facility safety. Similarly, closure liabilities should not disappear because the optimization horizon ends before the associated expenditure occurs.

Table 8 translates these priorities into testable study designs. Shared benchmarks should span bottleneck regimes, uncertainty structures and different decision representations, while retaining interpretable physical interfaces. Small exactly checkable instances can expose implementation errors; larger public or suitably anonymized instances can evaluate scale and transfer. Common formats for material properties, schedules, constraints and run records would make comparisons more cumulative. Such a benchmark program would support a more informative claim than a new metaphor accompanied by a small set of favorable objective values.

Table 8. Research gaps and the evidence needed to address them.

Research Gap	Concrete Study Design	Evidence Needed for a Stronger Claim
Fragmented stage objectives	Hold geological inputs and feasible decisions fixed; compare sequential and simultaneous planning.	Chain-wide value, resource balances and bottleneck shifts, with uncertainty intervals.
Hybrid complexity without attribution	Ablate the learning, local-search and repair components under a common total budget.	A measurable contribution from each component and a credible simpler baseline.
Surrogate exploitation	Use time- or geology-held-out data, uncertainty diagnostics and independent evaluation of optimizer-selected points.	Decision quality under model error, beyond average predictive accuracy.
Environmental proxies without life-cycle accounting	Jointly model water quality, residue properties, placement and progressive reclamation.	Validated environmental response and explicit separation of planning from facility safety.
Limited transfer and reproducibility	Release interoperable instances spanning commodities, scales, bottlenecks and uncertainty regimes.	Repeatable effects across independent instances, failures included [26], [27].

The no-free-lunch perspective reinforces the need to specify the problem class. Wolpert and Macready [31] related algorithmic performance to assumptions about the problems being solved. This does not imply that all methods perform equally on mining problems. Mining has structure: Precedences, material conservation, spatial relationships, nonlinear transformations and operating limits can all be exploited. The implication is that a performance claim should state which structure the method uses and where the evidence applies. Domain knowledge is a source of computational advantage when its contribution is made explicit and tested.

The review itself has limits. Its source selection is purposive, and restricted access prevented a uniform full-text audit of all studies. It provides no pooled estimate of economic improvement, publication-frequency trend or universal ranking. The exact example omits time, stochastic geology, queues, water chemistry, concentrate impurities and detailed residue behavior. Its purpose is to make coupling, feasibility and trade-offs fully inspectable. These limits point toward a practical progression: Verify the formulation and implementation on transparent instances, validate responses and uncertainty on relevant data, and only then assess transfer or operational impact.

7 | Conclusion

Nature-inspired optimization can contribute to mine-chain decisions when its representation and operators match the structure of the engineering problem. The literature examined here showed a progression from reusable search mechanisms toward coupled stochastic formulations, surrogate-supported decisions and adaptive hybrids. Its evidence remains interpretable only when predictive performance, modeled optimization gains and operational validation are distinguished.

The exact finite example demonstrated three consequences of that distinction. A low upstream cost could select a policy that violated a downstream limit; feasible profit, energy and water objectives produced different preferred configurations; and a high penalized score could conceal infeasibility until an instance-specific penalty threshold was exceeded. These are formulation and reporting issues that precede any claim of algorithm superiority. Future progress depends on preserving material and resource interfaces, comparing credible baselines under disclosed computational effort, and releasing records sufficient to reproduce both feasibility checks and statistical summaries. Under those conditions, nature-inspired search becomes a defensible component of mining decision support rather than a substitute for engineering evidence.

Data and code availability

The accompanying research package contains the complete dimensionless reference-instance specification, enumeration code, all configuration-level outputs, sensitivity and penalty records, figure sources, the verified-source ledger, and matching LaTeX and BibTeX material. The calculations use no private or measured mine dataset. No metaheuristic.

Acknowledgments

The authors have no acknowledgments to declare.

Author Contributions

All authors contributed to the conceptualization, literature review, analysis, and writing of the manuscript. All authors reviewed and approved the final version.

Funding

This research received no external funding.

Data Availability

No new empirical data were generated or analyzed in this review. All information is available from the cited sources.

Conflicts of Interest

The authors declare no conflict of interest.

Reference

[1] Newman, A. M., Rubio, E., Caro, R., Weintraub, A., & Eurek, K. (2010). A review of operations research in mine planning. *Interfaces*, *40*(3), 222–245. https://doi.org/10.1287/inte.1090.0492

[2] Goodfellow, R., & Dimitrakopoulos, R. (2017). Simultaneous stochastic optimization of mining complexes and mineral value chains. *Mathematical Geosciences*, *49*(3), 341–360. https://doi.org/10.1007/s11004-017-9680-3

[3] Dimitrakopoulos, R., & Lamghari, A. (2022). Simultaneous stochastic optimization of mining complexes-mineral value chains: An overview of concepts, examples and comparisons. *International Journal of Mining, Reclamation and Environment*, *36*(6), 443–460. https://doi.org/10.1080/17480930.2022.2065730

[4] Darling, P. (2011). *SME mining engineering handbook*. Society for Mining, Metallurgy & Exploration. https://books.google.com/books?id=5uq-kdfHLWUC&printsec=frontcover

[5] Glover, F. W., & Kochenberger, G. A. (2003). *Handbook of metaheuristics*. Springer Science & Business Media. https://doi.org/10.1007/b101874

[6] Ye, D. G. L., & Yinyu, A. (2016). *Linear and nonlinear programming*. Springer. https://doi.org/10.1007/978-3-319-18842-3

[7] Tan, J., Melkoumian, N., Harvey, D., & Akmeliawati, R. (2025). Nature inspired solutions for sustainable mining: applications of NIAs, swarm robotics, and other biomimicry-based technologies. *Biomimetics*, *10*(3), 181. https://doi.org/10.3390/biomimetics10030181

[8] Lin, J., Asad, M. W. A., Topal, E., & Chang, P. (2026). Metaheuristic approaches for open pit mine production scheduling: A comprehensive review. *International Journal of Mining, Reclamation and Environment*, 1–33. https://doi.org/10.1080/17480930.2026.2674140

[9] Kennedy, J., & Eberhart, R. (1995). Particle SWARM optimization. *Proceedings of ICNN'95 international conference on neural networks* (pp. 1942–1948). IEEE. https://doi.org/10.1109/ICNN.1995.488968

[10] Dorigo, M., Maniezzo, V., & Colorni, A. (1996). Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, *26*(1), 29–41. https://doi.org/10.1109/3477.484436

[11] Storn, R., & Price, K. (1997). Differential evolution-a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, *11*(4), 341–359. https://doi.org/10.1023/A:1008202821328

[12] Kalyanmoy, D. (2002). A fast and elitist multi-objective genetic algorithm: NSGA-II. *IEEE Trans on Evolutionary Computation*, *6*(2), 182–197. https://doi.org/10.1109/4235.996017

[13] Chicoisne, R., Espinoza, D., Goycoolea, M., Moreno, E., & Rubio, E. (2012). A new algorithm for the open-pit mine production scheduling problem. *Operations Research*, *60*(3), 517–528. https://doi.org/10.1287/opre.1120.1050

[14] Montiel, L., & Dimitrakopoulos, R. (2015). Optimizing mining complexes with multiple processing and transportation alternatives: An uncertainty-based approach. *European Journal of Operational Research*, *247*(1), 166–178. https://doi.org/10.1016/j.ejor.2015.05.002

[15] Yaakoubi, Y., & Dimitrakopoulos, R. (2023). Learning to schedule heuristics for the simultaneous stochastic optimization of mining complexes. *Computers & Operations Research*, *159*, 106349. https://doi.org/10.1016/j.cor.2023.106349

[16] Guo, J., Zhao, Z., Zhao, P., & Chen, J. (2024). Prediction and optimization of open-pit mine blasting based on intelligent algorithms. *Applied Sciences*, *14*(13), 5609. https://doi.org/10.3390/app14135609

[17] Amankwaa-Kyeremeh, B., McCamley, C., Ehrig, K., & Asamoah, R. K. (2024). Multiobjective optimisation of flotation variables using controlled-NSGA-II and paretosearch. *Resources*, *13*(11), 157. https://doi.org/10.3390/resources13110157

[18] Hasidi, O., El Alaoui-Chrifi, M. A., Chahid, R., Qazdar, A., Qassimi, S., Zaizi, F. Z., …, & Bendaouia, A. (2024). Data driven system for intelligent monitoring and optimization of froth flotation circuits using artificial neural networks and genetic algorithms. *Journal of Process Control*, *137*, 103198. https://doi.org/10.1016/j.jprocont.2024.103198

[19] Guimaraes, V. F., & Dimitrakopoulos, R. (2026). Simultaneous stochastic optimisation of mining complexes: Integrating waste management and progressive reclamation with encapsulation. *International Journal of Mining, Reclamation and Environment*, 1–14. https://doi.org/10.1080/17480930.2026.2654147

[20] Lones, M. A. (2014). Metaheuristics in nature inspired algorithms. *Proceedings of the companion publication of the 2014 annual conference on genetic and evolutionary computation* (pp. 1419–1422). Association for Computing Machinery. https://doi.org/10.1145/2598394.2609841

[21] Sörensen, K. (2015). Metaheuristics the metaphor exposed. *International Transactions in Operational Research*, *22*(1), 3–18. https://doi.org/10.1111/itor.12001

[22] Deb, K. (2000). An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering*, *186*(2–4), 311–338. https://doi.org/10.1016/S0045-7825(99)00389-8

[23] Zhang, Y., Guo, Y., Huang, Y., & Ge, S. (2025). A low-carbon scheduling method based on improved ant colony algorithm for underground electric transportation vehicles. *Complex & Intelligent Systems*, *11*(2), 166. https://doi.org/10.1007/s40747-024-01775-8

[24] Liu, Y., Zhang, Z., Zhu, D., Bo, L., Yang, S., Yue, Y., & Wang, Y. (2024). Mine water cooperative optimal scheduling based on improved genetic algorithm. *Heliyon*, *10*(6). https://doi.org/10.1016/j.heliyon.2024.e27289

[25] United Nations Environment Programme, International Council on Mining and Metals, Principles for Responsible Investment. (2020). *Global industry standard on tailings management*. https://wedocs.unep.org/handle/20.500.11822/36139

[26] Beiranvand, V., Hare, W., & Lucet, Y. (2017). Best practices for comparing optimization algorithms. *Optimization and Engineering*, *18*(4), 815–848. https://doi.org/10.1007/s11081-017-9366-1

[27] Bartz Beielstein, T., Doerr, C., Berg, D. van den, Bossek, J., Chandrasekaran, S., Eftimov, T., …, & Weise, T. (2020). *Benchmarking in optimization: Best practice and open issues.* https://doi.org/10.48550/arXiv.2007.03488

[28] Dolan, E. D., & Moré, J. J. (2002). Benchmarking optimization software with performance profiles. *Mathematical Programming*, *91*(2), 201–213. https://doi.org/10.1007/s101070100263

[29] Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, *7*(01), 1–30. https://www.jmlr.org/papers/v7/demsar06a.html

[30] Garcia, S., & Herrera, F. (2008). An extension on" statistical comparisons of classifiers over multiple data sets" for all pairwise comparisons. *Journal of Machine Learning Research*, *9*(12), 1–18. https://www.jmlr.org/papers/volume9/garcia08a/garcia08a.pdf

[31] Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, *1*(1), 67–82. https://doi.org/10.1109/4235.585893