

QUANTUM-INSPIRED AI METAHEURISTIC FRAMEWORK FOR MULTI-OBJECTIVE OPTIMIZATION IN INDUSTRIAL PRODUCTION SCHEDULING

Md Sanjid Khan¹

[1]. Department of Industrial and Systems Engineering, Lamar University, Texas, USA; Email: khansanjid9@gmail.com

Abstract

This study presents a quantum-inspired artificial intelligence metaheuristic (QI-AIM) framework for multi-objective industrial production scheduling, integrating algorithmic innovation with empirical validation across real manufacturing environments. Production scheduling inherently involves allocating limited resources—machines, tools, labor, and energy—to optimize multiple conflicting objectives such as makespan, tardiness, and energy cost. Existing evolutionary multi-objective algorithms (EMOAs) achieve competent performance but often struggle to maintain diversity and interpretability under industrial constraints. The proposed QI-AIM introduces a Q-bit-based priority representation, adaptive rotation updates, and episodic local search, augmented by surrogate modeling for energy-aware objectives. Implemented on classical hardware, QI-AIM emulates quantum principles to balance global exploration and local exploitation, enabling the generation of high-quality, well-distributed Pareto fronts within fixed computational budgets. A cross-sectional, multi-case empirical design links algorithmic outcomes to managerial perceptions using validated five-point Likert scales measuring perceived usefulness, ease of use, satisfaction, and intention to continue use. Data from flexible job-shop, flow-shop, and mixed-layout plants confirm that QI-AIM achieved the highest hypervolume (0.72) and lowest IGD (0.154) across 30-run benchmarks, outperforming NSGA-II, MOEA/D, and GA-LS hybrids while remaining computationally efficient. Regression analyses reveal that perceived usefulness significantly mediates the relationship between observed performance improvements and behavioral intention, with tardiness and makespan reductions explaining 45–53% of variance in acceptance outcomes. Survey means above 4.0 indicate strong managerial endorsement of the system's clarity, decision support, and interpretability. The findings establish that QI-AIM not only enhances operational performance but also strengthens decision confidence through transparent portfolio visualization of trade-offs. By coupling reproducible algorithmic design with human-centered evaluation, this research contributes a scalable, explainable, and adoption-relevant framework for Industry 4.0-aligned production scheduling and provides a blueprint for integrating advanced optimization with managerial decision processes.

Citation:

Khan, M. S. (2023). Quantum-inspired AI metaheuristic framework for multi-objective optimization in industrial production scheduling. *American Journal of Interdisciplinary Studies*, 4(3), 1–33. <https://doi.org/10.63125/2mba8p24>

Received:

June 20, 2023

Revised:

July 18, 2023

Accepted:

August 17, 2023

Published:

September 05, 2023



Copyright:

© 2023 by the author. This article is published under the license of American Scholarly Publishing Group Inc and is available for open access.


Keywords

Quantum-Inspired Metaheuristic, Multi-Objective Scheduling, Industrial Production, Hypervolume, IGD

INTRODUCTION

Production scheduling is the discipline concerned with allocating limited resources (machines, tools, labor, energy) to competing jobs over time so that one or more performance criteria are optimized (Pinedo, 2012). In real industrial settings, these criteria rarely align to a single scalar goal; rather, firms juggle makespan, tardiness, throughput, cost, energy, and quality, among others. This is the natural realm of multi-objective optimization (MOO), where a solution is evaluated by a vector of objectives and optimality is understood through Pareto efficiency rather than a single optimum (Coello Coello et al., 2007). To search these trade-off spaces, researchers increasingly use metaheuristics algorithms that combine stochastic exploration and problem-specific heuristics to construct high-quality approximate Pareto fronts for complex combinatorial scheduling problems. Evolutionary multi-objective algorithms (EMOAs) such as decomposition-based methods and reference-point-based nondominated sorting exemplify this line of work by offering principled ways to balance convergence and diversity (Deb & Jain, 2014).

Figure 1: Multi-Objective Production Scheduling Framework



From an international manufacturing perspective, the scheduling problem is strategically significant because it links operational efficiency and market responsiveness with sustainability targets and digital-transformation agendas found across both developed and emerging economies (Kusiak, 2018). The scheduling landscape is thus both practically and scientifically important: methods must scale to realistic shop floors, respect multiple simultaneous goals, and reflect managerial priorities observable in practice. This study aligns with that imperative by advancing a quantum-inspired AI metaheuristic framework for multi-objective industrial production scheduling and by integrating a quantitative, cross-sectional, multi-case empirical design to connect algorithmic performance with managerial outcomes measured on Likert's five-point scales (King & He, 2006). Background work evidences a shift from single-criterion formulations (e.g., makespan minimization) toward explicitly multi-objective scheduling in manufacturing, including energy-aware and cost-responsive models. Surveys and reviews in flow shop and job shop settings document the breadth of model variants, the maturation of local search and hybrid strategies, and the need for robust comparative studies across realistic benchmarks (Abdul, 2021; Pan & Ruiz, 2012a; Ruiz & Vázquez-Rodríguez, 2010). In parallel, sustainability and energy-efficiency drivers have motivated formulations that incorporate electricity tariffs and carbon-related performance as joint objectives or constraints, inviting multi-objective treatment in real factories (Mouzon et al., 2007; Zhang, 2011). These concerns align with the Industry

4.0 program, where cyber-physical systems and data-driven intelligence are intended to heighten responsiveness while reducing waste (Lee et al., 2020; Sanjid & Farabe, 2021). Yet, even as the algorithmic literature expands, two gaps persist: (i) few frameworks explicitly integrate quantum-inspired search operators to guide multi-objective scheduling in a way that is both computationally competitive and interpretable to operations managers, and (ii) empirical work linking algorithmic outputs to perceived managerial outcomes using validated survey scales remains limited. This paper addresses those gaps by proposing a quantum-inspired multi-objective metaheuristic tailored to production scheduling and by evaluating its outputs alongside survey-based measures of perceived effectiveness, usability, and managerial confidence (Gen & Lin, 2014; Lee et al., 2020).

At the algorithmic core, multi-objective metaheuristics manage the tension between convergence to the Pareto set and the maintenance of solution diversity. Representative techniques include decomposition-based methods that scalarize the objective vector into subproblems and reference-point-based nondominated sorting approaches that encourage uniform coverage of Pareto fronts under many objectives (Deb & Jain, 2014; Loukil et al., 2005; Omar & Rashid, 2021). Comparative studies in many-objective contexts have reinforced the importance of design choices population sizes, termination criteria, and indicator selection when benchmarking (Ishibuchi et al., 2015; Mubashir, 2021). In manufacturing scheduling specifically, evolutionary algorithms are frequently hybridized with local search and domain-specific heuristics to accelerate exploitation and honor precedence, machine availability, and setup structures (Lu, 2017; Rony, 2021). The result is a methodology space that is systematic but also flexible, enabling different preference models and heterogeneous constraints to be embedded without sacrificing computational practicality. The present framework builds on these foundations by adopting multi-objective evolutionary search with operators engineered for scheduling representations and by augmenting the search with a quantum-inspired encoding that supports broader exploration of solution neighborhoods aligned with production-logic constraints. In doing so, the framework speaks both to the theoretical literature on EMOAs and to the practical need to compute high-quality trade-off schedules that can be evaluated, communicated, and audited within plant-level decision processes (Zaki, 2021).

Quantum-inspired evolutionary algorithms (QIEAs) abstract concepts such as qubit representation and rotation gates to enrich variation operators without requiring physical quantum hardware. Surveys and empirical studies show that quantum-inspired encodings can strengthen balance between global exploration and local exploitation on difficult combinatorial search spaces (Danish & Zafor, 2022; Zhang & Chiong, 2016). In engineering design and energy-system scheduling, improved QIEA variants e.g., Latin-square QIEA and differential real-coded QIEA demonstrate competitive performance with robust convergence behavior (Danish & Md.Kamrul, 2022; Tsai et al., 2012). Within manufacturing scheduling, elitist QIEA variants have been adapted to flexible job-shop problems and other shop environments in which routing flexibility and sequence-dependent setups interact with tardiness or makespan objectives (Hossen & Atiqur, 2022). The framework advanced in this paper draws on those design patterns: it employs a quantum-inspired representation for candidate schedules, integrates reference-point-guided selection from many-objective evolutionary computation, and incorporates scheduling-specific local searches guided by feasibility and dispatching rules (Rabiul & Praveen, 2022). This combination targets richer sampling of the Pareto front and aims to generate decision-relevant alternatives that manufacturing managers can evaluate on both operational (e.g., tardiness, flowtime) and sustainability (e.g., energy cost) axes. By treating the QIEA mechanisms as modular, the framework remains compatible with standard performance indicators and statistical evaluation protocols commonly used in EMOA studies (Kamrul & Omar, 2022; Sullivan & Artino Jr., 2013; Tsai et al., 2012).

The empirical setting of this research is deliberately quantitative, cross-sectional, and multi-case. Cases are sampled from industrial plants where scheduling is operationally consequential and where the joint optimization of production and energy metrics is salient. For each case, algorithmic experiments produce Pareto-efficient schedules, which are then summarized using descriptive statistics and compared across objectives. Complementing these computational outputs, a structured survey instrument using Likert's five-point scales captures managerial assessments of clarity, perceived usefulness, and confidence in actionable scheduling alternatives. Using validated guidance on Likert-type data analysis, the study reports central tendency and dispersion and, where appropriate, applies parametric tests supported in the literature for Likert-type aggregated scores

(Razia, 2022; Sun et al., 2015). To examine linkages between algorithmic performance and managerial outcomes, correlations are estimated and interpreted using established guidance on effect magnitudes and assumptions (Danish, 2023; Schober et al., 2018). Regression models then evaluate associations between schedule features (e.g., expected tardiness, energy cost, buffer profiles) and managerial ratings while controlling for case-level covariates. This mixed algorithmic-empirical vantage echoes calls in the scheduling literature for evaluation protocols that reflect both computational and organizational criteria and aligns with broader Industry 4.0 orientations toward data-driven decision support (Schepers & Wetzels, 2007).

Strategically, the research questions examine: (RQ1) how a quantum-inspired multi-objective metaheuristic compares with strong baselines in generating diverse, high-quality schedules for industrial cases; (RQ2) how energy-aware objectives alter trade-offs among tardiness, flowtime, and cost; (RQ3) whether schedule characteristics are associated with positive managerial evaluations; and (RQ4) whether those evaluations cohere with constructs identified in technology-acceptance literature when instruments are adapted to scheduling decision support. The last dimension draws from well-established streams synthesizing determinants of perceived usefulness and behavioral intention in technology adoption (Arif Uz & Elmoon, 2023; Shen et al., 2015). The intent is not to replicate information-systems theory tests, but to ground survey design in validated constructs that have demonstrated explanatory power across contexts. By articulating hypotheses linking algorithmic outcomes to manager-reported appraisals (e.g., H1: higher schedule diversity and lower expected tardiness correlate with higher perceived usefulness), the study positions its regression models within a cumulative evidence base and facilitates interpretability across cases. Such framing complements the technical contributions in metaheuristic design by ensuring that evaluation criteria incorporate the managerial vantage that ultimately shapes the uptake of decision support in plant operations .

In addition, the study's contribution is situated at the intersection of three active literatures. First, within multi-objective scheduling, comprehensive reviews underscore the need for robust, comparable baselines and richer multi-objective testbeds in flow shop and job shop contexts (Razia, 2023; Venkatesh et al., 2012). Second, in sustainability-oriented scheduling, the inclusion of energy and tariff-sensitive objectives is increasingly recognized as operationally material in factories worldwide (Reduanul, 2023; Wirojanagud et al., 2007). Third, in computational intelligence, quantum-inspired variants have shown promise as flexible, problem-agnostic search schemes that can be specialized to structured industrial problems (Pan & Ruiz, 2012b; Sadia, 2023). By proposing and empirically examining a quantum-inspired AI metaheuristic framework with multi-objective scheduling capability and by explicitly connecting algorithmic outputs to manager-reported evaluations through a cross-sectional, multi-case design, the present study aims to enrich both the methodological toolkit and the evidence base that guides industrial scheduling practice. It anchors its methodology in internationally relevant manufacturing concerns, aligns measurement with validated constructs, and emphasizes transparent, descriptive, correlational, and regression analyses appropriate for Likert-scale instruments and case-based quantitative designs (Kusiak, 2018; Zayadul, 2023).

This study sets out a clear set of interlocking objectives that together define the scope and rigor of the investigation. First, it aims to design and operationalize a quantum-inspired AI metaheuristic framework tailored to multi-objective industrial production scheduling, specifying a Q-bit-based representation, measurement-to-schedule decoding, and hybrid local search capable of honoring precedence, machine capacity, setups, maintenance windows, and energy constraints. Second, it seeks to embed multiple, stakeholder-relevant objectives such as makespan, total tardiness, and energy or cost within a unified optimization scheme that maintains a high-quality approximation to the Pareto front under a fixed computational budget. Third, it intends to establish a strong empirical benchmark by conducting controlled comparisons against recognized baselines and heuristics under identical termination criteria, reporting performance through hypervolume, IGD, spacing, and run-time statistics with appropriate significance testing and effect-size measures. Fourth, it pursues a structured ablation program that isolates the marginal contributions of key components quantum rotation updates, quantum-walk neighborhoods, surrogate guidance, and local search so that observed gains are decomposed and attributable. Fifth, it targets robustness evidence by evaluating the framework across multiple industrial cases with differing shop characteristics, demand patterns,

and tariff regimes, thereby delineating boundary conditions and practical ranges for parameterization. Sixth, it couples algorithmic evaluation with a quantitative, cross-sectional survey of planners, supervisors, and operators using Likert's five-point scales, producing descriptive profiles, reliability metrics, and construct correlations that capture perceived ease of use, usefulness, satisfaction, and intention to continue use. Seventh, it investigates statistical associations between observed scheduling improvements and managerial assessments through a set of prespecified regression models with diagnostics for multicollinearity, distributional assumptions, and heteroskedasticity, complemented by robustness checks. Eighth, it articulates actionable artifacts for practitioners, including parameter defaults, data-readiness requirements, and integration touchpoints with MES/ERP environments, alongside transparent reporting of seeds, hardware, and configuration files to support reproducibility. Finally, it frames ethical handling of operational and survey data through anonymization procedures, secure storage, and role-appropriate access, ensuring that methodological choices align with organizational governance and industrial standards.

LITERATURE REVIEW


The literature on industrial production scheduling spans a broad spectrum of formulations, methods, and evaluation protocols, reflecting the inherently multi-criteria nature of real factories and the practical need to trade off efficiency, service, and sustainability. Foundational streams address classical job shop, flow shop, and flexible job shop models, progressively enriching them with realistic constraints such as sequence-dependent setups, machine eligibility, preventive maintenance windows, and energy tariffs. On the solution side, research has evolved from exact methods suited to small instances toward metaheuristics and hybrid approaches that scale to complex plants while approximating high-quality Pareto fronts for conflicting objectives like makespan, total tardiness, energy consumption, and operating cost. Evolutionary multi-objective algorithms constitute a central pillar, with decomposition-based and nondominated-sorting variants paired with indicator-driven selection to balance convergence and diversity; these are frequently augmented by domain-specific local search, dispatching heuristics, and repair operators to maintain feasibility and accelerate exploitation. A complementary stream explores learning-augmented optimization, using surrogates to reduce expensive evaluations and reinforcement or adaptive mechanisms to tune operators online. In parallel, sustainability and energy-aware scheduling have grown into prominent subareas, motivated by time-of-use pricing, decarbonization targets, and enterprise-wide energy management, which jointly necessitate integrated production–energy decisions. More recently, quantum-inspired evolutionary algorithms have introduced Q-bit encodings and rotation-based updates as variation mechanisms on classical hardware, promising richer exploration without requiring access to quantum devices; initial applications in manufacturing suggest competitive performance and flexible hybridization with established schedulers. Despite methodological advances, the literature notes recurring challenges: fragmented benchmarking across heterogeneous testbeds, limited ablation studies that disentangle component contributions, and comparatively fewer empirical designs that link algorithmic outputs to managerial perceptions and adoption considerations. Drawing from technology-acceptance and decision-support research, a growing body of work advocates incorporating user-centered evaluation perceived usefulness, ease of use, satisfaction, and intention to continue use alongside operational metrics. This review synthesizes these streams to situate a quantum-inspired AI metaheuristic within established optimization traditions, highlight design choices that influence multi-objective performance, and frame an empirical evaluation agenda that integrates algorithmic benchmarking with survey-based assessment across multiple industrial cases.

Industrial Production Scheduling and Multi-Objective Optimization

Industrial production scheduling formalizes how finite resources machines, tools, operators, and shift time are allocated to a set of jobs so that technological precedence and capacity constraints are respected while performance is optimized across multiple criteria. Canonical shop configurations include permutation flow shop, classic job shop, and flexible job shop, each introducing distinct combinatorial structure and implications for encodings, neighborhood moves, and feasibility repair. Foundational comparative work in the permutation flow shop underscored how objective choice (e.g., makespan versus flowtime) and constraint details (e.g., sequence-dependent setups, blocking) reshape solution landscape ruggedness and the effectiveness of constructive versus

improvement procedures (Framinan et al., 2005). Subsequent research codified simple yet powerful improvement frameworks such as iterated greedy with destruction–construction and tailored local search that scale well and provide consistent baselines for complex instances (Minella et al., 2018). Parallel streams emphasized modeling realism by incorporating setups, maintenance windows, machine eligibility, and changeover costs; these features alter both dispatching logic and neighborhood effectiveness, and they motivate hybrid metaheuristics that fuse domain heuristics with global exploration (Allahverdi, 2015). As scheduling practice globalized with Industry 4.0 digitization, the problem's international salience intensified: firms in diverse regulatory and energy contexts now seek trade-offs among service levels, energy cost, and throughput, rendering single-criterion formulations insufficient. Against this backdrop, the field increasingly adopts multi-objective formulations where a schedule is evaluated by a vector of competing criteria and optimality is defined by Pareto efficiency rather than a single scalar optimum. This reframing refracts algorithm design choices through two intertwined lenses: convergence toward the (unknown) Pareto set and diversity across it, both of which must be maintained under realistic computational budgets and dynamic plant constraints (Allahverdi, 2015; Framinan et al., 2005).

Figure 2: Industrial Production Scheduling



Within a multi-objective lens, the production scheduling problem can be written as a vector optimization program over feasible schedules $x \in X$:

$$x \in X \min F(x) = \left(C_{\max}(x), \sum_{i=1}^n T_i(x), E(x) \right)$$

Subject to (i) technological precedence: $S_{i,k+1} \geq C_i k$, (ii) machine capacity: at most one operation executes on a machine at a time, (iii) release and due-date windows: $S_{i,1} \geq r_i$, tardiness $T_i = \max\{0, C_i - d_i\}$, and (iv) optional energy or tariff constraints such as peak caps and time-of-use segments. Here, C_{\max} is the makespan, $\sum T_i$ the total tardiness, and $E(x)$ an energy or cost functional derived from process-time and tariff profiles. This formulation foregrounds trade-offs that are central in factories operating under volatile demand and energy pricing: reducing tardiness may extend makespan; shaving peaks in energy use may increase flowtime; and prioritizing robustness may

reduce short-term throughput. Because exact methods deteriorate rapidly as problem size and realism increase, contemporary research embraces metaheuristics capable of sampling the feasible set broadly while embedding domain knowledge through decoding heuristics and local improvement. Reviews emphasize that the credibility of multi-objective schedulers depends on transparent encodings (priority lists or operation-based permutations), reproducible computational budgets, and principled indicators such as hypervolume and inverted generational distance for assessing solution sets rather than single solutions (Minella et al., 2018; Özgüven et al., 2010). These methodological norms enable consistent comparisons across testbeds, encourage ablation studies that isolate the contribution of specific operators, and support deployment discussions in which plant managers evaluate not only the “best” schedule but a portfolio of alternatives aligned with operational priorities and risk tolerances (Minella et al., 2018).

Algorithmically, the field has converged on evolutionary multi-objective frameworks that balance exploitation and exploration through population-based search, recombination, and selection guided by Pareto dominance, indicators, or decomposition. A seminal development is decomposition-based optimization, which transforms the vector problem into a set of scalar subproblems associated with weight vectors and then coordinates their evolution so that collectively they approximate the Pareto front; the approach scales well to many objectives and admits flexible neighborhood policies (Muhammad & Redwanul, 2023; Srinivas & Manish, 2023; Zhang & Li, 2007). On the single-objective side often used as components within multi-objective hybrids iterated greedy remains a robust backbone due to its simplicity, amenability to problem-specific destruction/reconstruction, and strong empirical performance across flow-shop families (Ruiz & Stützle, 2007). For problem realism, surveys on flexible job shops catalog routing flexibility and resource contention phenomena that complicate both decoding and neighborhood design, reinforcing the role of hybridization and adaptive operators (Özgüven et al., 2010). Complementary synthesis has shown that setup-aware formulations materially shift both objective surfaces and algorithm behavior, suggesting that operator design must reflect changeover asymmetries and batching effects to retain efficacy when objectives expand beyond makespan. Finally, comprehensive evaluations of permutation flow-shop methods illuminate the value of standardized benchmarks and careful parameter control, clarifying which mechanisms generalize across objectives and which are sensitive to landscape features a prerequisite for credible multi-objective scheduling in heterogeneous industrial settings (Allahverdi, 2015; Özgüven et al., 2010; Zhang & Li, 2007).

Quantum-Inspired Metaheuristics for Combinatorial Optimization


Quantum-inspired metaheuristics (QIMs) adapt mathematical notions such as qubits, superposition, and rotation gates to enrich classical search with probabilistic states and directionally controlled variation. In a QIM, a solution can be encoded by a vector of “quantum bits,” each described by complex amplitudes that determine the probability of observing a 0 or 1, thereby enabling a population to compactly represent and quickly re-sample large regions of the combinatorial space without storing every candidate explicitly. Rotation-gate updates steer these probability amplitudes toward fitter configurations while preserving diversity, and measurement collapses them to concrete schedules when an evaluation is required. This representational agility is valuable for industrial scheduling, where precedence, machine eligibility, setup, and blocking constraints generate rugged search landscapes with many local optima. Comparative reviews of quantum-inspired genetic and evolutionary schemes report faster convergence to high-quality discrete solutions and reduced risk of premature stagnation compared with purely classical heuristics, a property attributed to the probabilistic sampling around elite exemplars and to adaptive angle updates that balance exploration and exploitation (Lahoz-Beltra, 2016). Empirical accounts in stochastic job-shop settings further show that co-evolutionary quantum genetic designs improve population diversity and robustness by allowing multiple subpopulations (e.g., schedules and scenarios) to compete and cooperate under uncertainty, yielding smaller makespans on average than baseline evolutionary frameworks (Gu et al., 2010).

Among QIM variants, quantum-behaved particle swarm optimization (QPSO) has been widely applied to discrete scheduling due to its compact update rule and strong global search. In QPSO, each particle is attracted toward its personal best and the swarm’s mean best; under a quantum delta potential well model, its next position is sampled by

$$x_i(t + 1) = p_i(t) \pm \beta \frac{m_{best}(t) - x_i(t)}{\ln(u_1)}, \quad u \sim U(0,1)$$

Where $\beta > 0$ controls contraction–expansion, $p_i(t)$ is a pbest-guided attractor, and $m_{best}(t)$ is the component-wise mean of personal bests. This log-uniform sampling pulls particles toward elite regions while injecting scale-aware randomness that helps escape combinatorial traps an effect that maps well to routing-and-sequencing substructures in flexible job-shops. When embedded with discrete encodings and local improvement (e.g., insertion/adjacent-swap neighborhoods), QPSO has delivered competitive makespan reductions and stable workloads across flexible job-shop benchmarks versus classical PSO and GA baselines, highlighting its suitability as the search engine inside a quantum-inspired, multi-objective scheduling framework (Singh & Mahapatra, 2016). Moreover, improved quantum-inspired evolutionary algorithms introduce local search and specialized decoding to handle additional complexities like no-wait constraints, sequence-independent setups, and release dates, thereby extending the range of factory contexts that can be handled within a unified QIM template (Zhao et al., 2016).

Figure 3: Quantum-Inspired Metaheuristics for Combinatorial Optimization




Beyond single-objective makespan, quantum-inspired formulations have been adapted to multi-criteria, resource-constrained settings closely aligned with industrial practice. In resource-constrained project scheduling, for example, a quantum-inspired genetic algorithm (QIGA) uses qubit strings and rotation-angle adaptation to maintain feasible precedence-respecting activity lists while optimizing time- and resource-based objectives. Results indicate that QIGA can reach superior or comparable schedules with fewer individuals and fewer generations than classical GAs, suggesting better sample efficiency when the feasible region is narrow and multimodal (Saad et al., 2021). Similarly, for dynamic flow-shops and stochastic job-shops, co-evolutionary quantum genetic mechanisms foster resilience by allowing solution and disturbance populations to co-adapt, thereby maintaining performance under breakdowns and uncertain processing times (Gu et al., 2010). Collectively, these strands position quantum-inspired metaheuristics as a principled and practical foundation for the proposed framework: qubit-based representations to encode rich sequencing/assignment decisions, QPSO-style sampling to explore elite basins without losing diversity, and rotation-gate or co-evolutionary updates to negotiate competing industrial objectives under realistic constraints (Lahoz-Beltra, 2016; Singh & Mahapatra, 2016).

Hybrid and Learning-Augmented Metaheuristics in Scheduling

Hybrid and learning-augmented metaheuristics combine population-based or trajectory-based search with adaptive components that learn useful structure during the run e.g., which neighborhoods to prefer, which parameter settings to adopt, or which candidate solutions merit expensive evaluation. One influential thread is hyper-heuristics, where a higher-level controller

selects or generates low-level heuristics in response to feedback from the search; the goal is to improve generality across instances while keeping the framework modular for domain constraints such as precedence, sequence-dependent setups, and resource calendars. A widely cited survey formalizes hyper-heuristics as learning mechanisms operating over heuristic spaces and distinguishes “selection” versus “generation” families; it also reports practical designs that integrate simple memory, tabu mechanisms, and acceptance criteria to balance diversification with intensification (Burke et al., 2013).

Figure 4: Hybrid and Learning-Augmented Metaheuristics Framework



In production scheduling, this view is attractive because it allows a single controller to coordinate constructive rules (e.g., dispatching/decoding) with improvement moves (e.g., insertion/exchange on operation sequences) under multiple objectives. A second learning axis is surrogate modeling, where an inexpensive model $f^{\wedge}(x)$ approximates an expensive objective $f(x)$ (e.g., when discrete-event simulation or energy costing is embedded in schedule evaluation). Early surveys systematized model choices (radial basis, Kriging/GP, polynomial regression) and management policies for when to trust surrogates versus ground-truth evaluations (Jin, 2005). In multi-objective scheduling, surrogates are typically used inside scalarization subproblems or reference-point neighborhoods to accelerate the evaluation of candidate schedules while preserving a high-quality approximation of the Pareto set. Together, hyper-heuristics and surrogates provide a principled scaffold for industrial scheduling: the former learns what to try next (operator/heuristic choice), and the latter learns what is promising (objective predictions), yielding robust performance under tight computational budgets and heterogeneous shop-floor constraints (Burke et al., 2013; Jin, 2005).

Surrogate-assisted evolutionary computation subsequently introduced analytic acquisition rules that turn predictions into sampling decisions. A canonical choice is the *expected improvement* (EI) criterion

$$EI(x) = (\mu^* - \mu(x))\Phi(z(x)) + \sigma(x)\phi(z(x)), \quad z(x) = \frac{\mu^* - \mu(x)}{\sigma(x)}$$


where $\mu(x)$ and $\sigma(x)$ are the surrogate's mean and standard deviation at schedule x , μ^* is the incumbent best scalarized objective, and Φ/ϕ are the normal CDF/PDF. EI formalizes the exploration–exploitation trade-off: large uncertainty $\sigma(x)$ encourages exploring unfamiliar schedule regions (e.g., different machine assignments), while low $\mu(x)$ favors intensification near promising sequences. In multi-objective settings, EI can be applied within decomposition (per weight vector or reference point) to guide sample-efficient improvement of hypervolume or ϵ -dominance coverage. A concise survey of surrogate-assisted EC traces these ideas and emphasizes model management periodic

retraining, infill sampling, and uncertainty calibration to avoid model bias (Jin, 2011). Beyond modeling the objective, learning also targets algorithm design itself. Algorithm configuration (AC) methods treat a solver's parameters (e.g., population size, cooling rate, neighborhood frequencies) as decision variables and optimize them on representative instances prior to deployment. The SMAC framework (Sequential Model-Based Optimization for General Algorithm Configuration) introduced a random-forest surrogate and adaptive sampling to search large, mixed discrete/continuous parameter spaces effectively, enabling consistent gains over manual tuning (Hutter et al., 2011). For industrial scheduling, AC closes the loop between case specifics (layout, product mix variability, energy regime) and solver behavior, yielding tuned hybrids that respect time budgets and stability constraints. In concert, surrogate-guided sampling and SMAC-style configuration make learning-augmented metaheuristics both sample-efficient and instance-aware (Hutter et al., 2011; Jin, 2005). A complementary capability is algorithm selection (AS), which learns a mapping from *instance features* (e.g., machine/job ratios, setup density, due-date tightness) to the algorithm/parameterization expected to perform best on that instance. The ASlib benchmark and schema standardized features, performance data, and evaluation splits across domains, catalyzing reproducible AS studies and practical per-instance selectors (Bischl et al., 2016). For production scheduling, AS enables a *portfolio* of hybrids e.g., a hyper-heuristic with strong local search for setup-heavy flow shops, a surrogate-assisted MOEA for energy-aware job shops, and a matheuristic for tight due-date regimes to be dispatched adaptively based on measurable shop characteristics. Finally, racing and iterated racing procedures provide statistically grounded routines to compare many candidate heuristics/parameter settings under noise and instance variability, pruning weak options quickly and focusing evaluations on promising competitors (Birattari et al., 2010). In a scheduling pipeline, racing can be used offline to curate a compact portfolio and online to switch neighborhoods/operators as evidence accrues during a run. Putting these strands together yields a cohesive template for hybrid, learning-augmented schedulers: (i) a hyper-heuristic controller chooses operators; (ii) a surrogate model evaluates candidates cheaply and proposes infill via EI; (iii) SMAC-style configuration tailors parameters to the plant context; (iv) AS selects the solver variant per instance; and (v) racing validates choices under stochasticity. This integrated view aligns tightly with industrial needs for reproducibility, responsiveness, and transparent trade-off navigation across makespan, tardiness, and energy/cost objectives (Birattari et al., 2010; Bischl et al., 2016).

Human Factors in Decision-Support for Scheduling

The human-technology interface is central to whether advanced schedulers translate into operational value, because planners and supervisors ultimately interpret model outputs and trigger actions on the shop floor. A mature stream in information systems integrates user satisfaction with technology acceptance, distinguishing beliefs about the system as an object (e.g., information quality, system quality) from beliefs about *using* the system (e.g., perceived usefulness, ease of use), and theorizing how these belief sets jointly shape attitudes and usage. This integration clarifies why even accurate optimization tools may falter if interaction design, feedback timing, or explanation facilities are weak: object-based beliefs drive satisfaction, which in turn conditions the behavioral beliefs that underpin continued use. For production scheduling, this implies that interfaces must make trade-offs among makespan, tardiness, and energy legible, and that explanation mechanisms should help users trace how a given schedule reflects constraints and priorities. When evaluation instruments capture both belief families, researchers can map algorithmic improvements (e.g., better Pareto coverage) onto perceived usefulness and satisfaction, thereby connecting computational performance to adoption-relevant perceptions in a single empirical model (Wixom & Todd, 2005). This framing also supports cross-case comparisons, since it separates construct domains that often blur in practice (e.g., conflating data quality with ease of use). By adopting this integrated lens, studies can design surveys and usage logs that respect how factory professionals actually experience decision-support through both the *qualities of the tool* and the *experience of applying it* under time pressure (Benbasat & Barki, 2007).

Figure 5: Technology Adoption and Human Factors Framework



At the same time, the acceptance literature cautions against over-reliance on any single model or narrow set of predictors. A well-known critique argued that an excessive focus on a single acceptance model can obscure contextual determinants (training, process fit, incentives, management support) and encourage repetitive variance-explanation exercises disconnected from design and organizational realities. For industrial scheduling, that warning is practical: adoption hinges not only on usefulness and ease of use, but also on fit with dispatching practices, accountability norms, and interdepartmental coordination (e.g., production vs. maintenance). Incorporating IS success constructs information, system, and service quality; use; user satisfaction; net benefits adds diagnostic breadth for deployment, enabling teams to identify whether problems stem from data pipelines (information quality), algorithm responsiveness (system quality), or support processes (service quality). In a case-study design, these constructs provide a balanced scorecard that pairs optimization metrics with stakeholder-facing outcomes (e.g., perceived decision confidence, schedule stability across shifts). For survey development and analysis, this perspective encourages validated multi-item measures and triangulation with usage logs, mitigating the risk that a single perceptual dimension dominates inference about adoption. Importantly, it aligns with how factories justify investments: through net benefits that span throughput, delivery performance, energy cost, and operator workload, rather than through abstract usage intentions alone (Henseler et al., 2015).


Rigorous measurement is the bridge between human factors and credible inference in adoption research especially in multi-case, cross-sectional studies. Reviews in health-IT acceptance synthesize evidence on instrument validity, respondent roles, and context sensitivity, offering guidance on sampling (clinicians vs. staff), scale construction, and interpretation of relationships among ease, usefulness, and behavioral intentions; these lessons generalize to manufacturing where roles (planner, supervisor, operator) differ in authority and touchpoints with the scheduler. Method guidance also emphasizes guarding against common-method variance and establishing discriminant validity among latent constructs; for instance, the heterotrait–monotrait (HTMT) criterion has emerged as a stronger test than legacy cross-loading checks in variance-based SEM (Petter et al., 2008). For industrial scheduling studies, these practices translate into instrument pilots, role-specific wording, multi-source data (surveys + logs), and validity diagnostics that confirm constructs such as perceived usefulness, ease of use, satisfaction, and intention to continue use are empirically distinct yet meaningfully related. Embedding such rigor supports trustworthy correlation and regression analyses that link observed scheduling improvements (e.g., reduced tardiness, lower energy peaks) to adoption-relevant perceptions, without conflating measurement artifacts with substantive effects. Ultimately, by combining context-aware constructs, balanced IS-success outcomes, and modern

validity checks, researchers can provide plant leaders with evidence that speaks both to human acceptance and to operational performance evidence necessary for scaling decision-support beyond pilots (Holden & Karsh, 2010).

METHOD

This study has been designed as a quantitative, cross-sectional, multi-case investigation that has integrated algorithmic benchmarking with survey-based assessment in real industrial settings. The methodological architecture has combined two tightly linked streams of evidence: (i) computational experiments that have produced and evaluated Pareto-efficient production schedules using a quantum-inspired AI metaheuristic, and (ii) a structured questionnaire that has captured managers' and planners' perceptions of the scheduler's usability and value on Likert's five-point scales. By pairing objective performance with human-centred evaluation, the design has provided convergent validity for claims about operational improvement and adoption-relevant outcomes.

Figure 6: Methodological Framework for this Study



The case selection has followed purposeful criteria that have ensured data availability, process stability during the study window, and leadership buy-in for access to shop-floor logs. Each case has contributed operational datasets that have included routings, processing and setup times, release dates, due dates, machine calendars, maintenance windows, and, when available, energy metering or tariff profiles. These data have been pre-processed and harmonized into a standardized schema so that schedule generation and performance measurement have been comparable across sites. The quantum-inspired metaheuristic has been configured to handle multi-objective targets makespan, total tardiness, and energy/cost under identical computational budgets across baseline algorithms, and it has maintained an external archive to approximate the Pareto front. Decoding heuristics and feasibility repairs have been applied so that machine capacity, precedence, and setup constraints have been respected. For fairness, all algorithms have shared common termination criteria and seed management. On the perceptual side, the survey instrument has measured perceived ease of use, perceived usefulness, user satisfaction, and intention to continue use, with training adequacy recorded as a contextual factor. Items have been adapted to the scheduling context and have been piloted for clarity. The analytic plan has specified descriptive statistics for respondent profiles and constructs, reliability checks for scale consistency, and correlation analysis to summarize associations among constructs and observed performance deltas (e.g., changes in makespan, tardiness, and energy). Multiple linear regression models have been prespecified to examine relationships between scheduling improvements and managerial perceptions while controlling for case characteristics. Throughout, reproducibility has been prioritized: parameter settings, random seeds, hardware details, and configuration files have been logged; data handling has followed anonymization procedures; and analysis scripts have been organized so that results have been traceable from raw inputs to reported tables and figures.

Research Design

The study has been conceived as a quantitative, cross-sectional, multi-case design that has integrated algorithmic experimentation with survey-based assessment to produce convergent evidence about both operational performance and user perceptions. Specifically, the design has encompassed two coordinated streams: (i) computational experiments that have generated Pareto-efficient schedules using a quantum-inspired AI metaheuristic under identical computational budgets and baseline comparators, and (ii) a structured questionnaire that has captured planners', supervisors', and operators' evaluations on Likert's five-point scales. Case selection has followed purposeful criteria data availability, stable process windows, and executive sponsorship so that each site has provided clean operational logs (routings, processing and setup times, calendars, maintenance windows, due dates, and, where applicable, energy tariffs/meters). To ensure comparability, datasets have been harmonized into a common schema, and decoding/repair heuristics have been standardized so that precedence, capacity, setup, and release-date constraints have been respected uniformly across cases. The algorithmic stream has maintained an external archive for non-dominated solutions and has reported hypervolume, IGD, and runtime, while all solvers have shared seeds and termination rules to preserve fairness. The survey stream has measured perceived ease of use, perceived usefulness, user satisfaction, and intention to continue use, and has recorded training adequacy and role as contextual factors; items have been piloted and refined for clarity. Sampling has targeted role diversity rather than a single department so that perceptions have reflected the end-to-end scheduling workflow. The analytic plan has prespecified descriptive summaries, reliability checks, correlations linking construct scores with observed KPI deltas (changes in makespan, tardiness, and energy), and multiple regression models with controls for case characteristics. Throughout, the design has prioritized reproducibility and governance: parameter files, random seeds, and hardware profiles have been logged; scripts and configuration artifacts have been version-controlled; and data have been anonymized with role-appropriate access, so that all reported results have been traceable from raw inputs to final tables and figures.

The study has assembled a multi-case corpus drawn from discrete-manufacturing plants that have met clear inclusion criteria and have provided sufficient operational depth for rigorous evaluation. Candidate sites have been screened for (i) continuous electronic capture of routing, processing, and setup times; (ii) stable product families over the observation window; (iii) documented machine calendars and maintenance histories; and (iv) leadership consent for anonymized data sharing. From this pool, two to three factories covering a flexible job-shop with high product mix variability, a flow-shop with sequence-dependent setups, and a mixed layout with parallel machines have been selected to ensure heterogeneity in routing flexibility, bottleneck structure, and energy-tariff exposure. Each site has maintained at least two shifts per day, and each has operated under due-date-driven order release with periodic rescheduling triggered by disruptions (e.g., breakdowns, rush orders). To support comparability, case datasets have been extracted for aligned windows (8–12 weeks) and have been harmonized to a common schema capturing job identifiers, operation precedence, machine eligibility, start/finish stamps, planned/actual setup, release and due dates, and where available submetered energy or time-of-use tariff blocks. Data-quality audits have been performed (missingness checks, clock drift, outlier screening), and discrepancies have been documented with corrective rules applied consistently across cases. Context variables shop size, job-to-machine ratio, setup-time share, due-date tightness, and historical on-time delivery have been profiled to serve as controls in analysis. Governance has been formalized through data-processing agreements, de-identification of personnel and product codes, and role-based access to raw logs. Each plant has contributed a brief process narrative (materials flow, changeover policy, dispatching conventions) to anchor decoding heuristics and feasibility repairs used during schedule generation. Collectively, these preparations have ensured that cases have represented realistic industrial diversity while remaining methodologically commensurate, enabling side-by-side evaluation of algorithmic performance and alignment of survey constructs with the lived scheduling context at each site.

Data Collection

The study has drawn on two complementary data sources operational logs and structured surveys that together have provided both objective performance evidence and managerially salient perceptions. Operationally, each case site has contributed machine- and job-level event logs that have included routings, operation start/finish stamps, processing and setup times, machine eligibility lists, preventive-maintenance windows, breakdown records, release and due dates, and shift calendars; where available, submetered energy traces or time-of-use tariff tables have been supplied to enable cost and consumption modeling. These raw feeds have been ingested into a unified schema that has mapped jobs, operations, and machines via stable identifiers, and has resolved clock drift and timezone inconsistencies through reference-beacon checks. Data-quality procedures have encompassed missingness audits, outlier screening for cycle and setup times, reconciliation between planned and actual routings, and imputation rules that have been pre-registered and applied consistently across cases. From the harmonized logs, the analysis dataset has derived features that have been necessary for decoding and evaluation operation precedence graphs, feasible machine sets, historical setup matrices (including sequence dependence), due-date tightness indicators, bottleneck utilization profiles, and, when energy data have been present, load curves aligned to tariff blocks. In parallel, the perceptual layer has been captured via a questionnaire that has measured perceived ease of use, perceived usefulness, user satisfaction, and intention to continue use on Likert's five-point scales, with training adequacy and role recorded as contextual covariates. The instrument has been piloted for clarity and has incorporated attention checks; responses have been screened for straight-lining and excessive latency, and composite scores have been computed following pre-specified item-keying rules. Linkage between the two sources has been achieved at the case level and analysis window rather than at the individual level, so that confidentiality has been preserved while enabling correlations between observed performance deltas (e.g., changes in makespan, tardiness, and energy) and aggregated perceptual measures. Throughout, governance has been ensured by de-identification, role-based access control, and an audit trail that has recorded all transformations from raw ingestion to analysis-ready tables.

Instrument Development

The measurement instrument has been developed to capture perceptual constructs that have been central to adoption and use of the scheduling system perceived ease of use, perceived usefulness, user satisfaction, intention to continue use, and training adequacy while remaining concise and role-appropriate for planners, supervisors, and operators. Item pools for each construct have been drafted from prior validated wording and have been adapted to the production-scheduling context (e.g., "The scheduler has helped me complete planning tasks more quickly"). A five-point Likert format (1 = Strongly disagree ... 5 = Strongly agree) has been selected for all items to ensure consistency and respondent efficiency. Content validity has been established through expert review with two operations managers and one industrial engineer, who have provided structured feedback on relevance, clarity, and redundancy; items flagged for ambiguity or overlap have been revised or removed. Cognitive interviews ($n \approx 6-8$ across roles) have been conducted to probe comprehension and response processes, and wording has been refined to minimize jargon and double-barreled phrasing. A pilot administration (target $n \approx 25-40$ across cases) has been executed to evaluate reliability and timing burden; attention checks and response-quality screens (straight-lining flags, completion-time thresholds) have been embedded. Scale reliability has been assessed using Cronbach's alpha with a prespecified threshold of ≥ 0.70 ; composite reliability and average variance extracted have been computed for convergent validity, and the HTMT ratio has been examined for discriminant validity. To mitigate common-method variance, psychologically distinct instructions and item order randomization have been used, negatively keyed items have been included sparingly and balanced per scale, and a marker item not theoretically related to core constructs has been added for post-hoc adjustment. A codebook has documented item wording, construct mappings, scoring keys (including reverse-coded items), and rules for handling missing data (pairwise deletion within scale, minimum answered items $\geq 50\%$). The final instrument package has included an information sheet, confidentiality statement, and role-specific routing to ensure only applicable items have been displayed. All revisions, decision logs, and psychometric summaries have been version-

controlled so that subsequent administrations have remained traceable and comparable across sites.

QI-AIM Framework

The quantum-inspired AI metaheuristic (QI-AIM) has been engineered to generate feasible, high-quality schedules for multi-objective industrial problems while operating under standardized computational budgets. Candidate solutions have been encoded with a Q-bit-driven priority representation in which each operation has been associated with a pair of probability amplitudes that have governed its relative selection order; measurement of these amplitudes has produced a concrete priority list that a decoding heuristic has translated into a feasible schedule while enforcing precedence, single-machine capacity, setup, and release-date constraints. Initialization has been performed by seeding a balanced superposition around neutral priors and sprinkling a small fraction of domain-aware seeds (e.g., SPT, EDD, and setup-aware heuristics) to anchor early exploitation. Iteration has proceeded by updating Q-bit rotation angles toward elite exemplars with adaptive step sizes tied to dominance counts and archive crowding, so that convergence pressure and diversity preservation have both been maintained. To avoid structural stagnation, a quantum-walk neighborhood has been applied intermittently to perturb blocks of operations on bottleneck machines, and a physics-inspired local search (insert/exchange/2-opt variants) has been invoked on an improving schedule, with frequency throttled by recent gain. An external non-dominated archive has been maintained, with crowding distance and ϵ -dominance filters that have controlled growth and spread; reference-point sampling from the archive has guided rotation targets in many-objective cases. For expensive evaluations (e.g., energy costing under time-of-use tariffs), a surrogate model (gradient-boosted regression) has been trained online on evaluated solutions and has been queried to pre-screen candidates, while a model-management rule has ensured periodic ground-truthing to cap bias. All runs have shared a fixed evaluation budget, identical termination rules (max evaluations or stall on hypervolume), and common random seed protocols to ensure fairness against baselines (NSGA-II, MOEA/D, and a GA-with-local-search hybrid) that have been matched for population size and variation effort. Feasibility repair routines and decoding parameters have been held constant across algorithms, hardware profiles and seeds have been logged, and configuration files have been version-controlled so that every reported Pareto set and metric (hypervolume, IGD, runtime) has been reproducible from raw inputs.

Regression Models

The inferential layer has been structured around two prespecified multiple-regression models that have linked observed scheduling improvements to adoption-relevant perceptions, with controls for case context. Model A (Adoption) has examined whether perceived usefulness and ease of use have been associated with intention to continue use while accounting for user satisfaction and training adequacy, formalized as:

$$ICU_i = \beta_0 + \beta_1 PU_i + \beta_2 PEOU_i + \beta_3 US_i + \beta_4 TA_i + \beta_c^T C_i + \epsilon_i,$$

Model B (Usefulness Drivers) has related perceived usefulness to measured performance deltas extracted from the Pareto sets (case-level aggregates), expressed as:

$$PU_i = \gamma_0 + \gamma_1 \Delta C_{max,i} + \gamma_2 \Delta T_i + \gamma_3 \Delta E_i + \gamma_c^T C_i + \eta_i,$$

Here, C_i has denoted control variables (shop size, job-to-machine ratio, setup-time share, due-date tightness, and role indicators); ΔC_{max} , ΔT , and ΔE have been computed as percentage improvements relative to historical baselines over the aligned analysis window. Both models have been estimated with OLS under heteroskedasticity-robust standard errors (HC3), and all predictors have been standardized (z-scores) after reliability screening of multi-item scales so that coefficients have been interpretable as standardized effects. Predictor collinearity has been pre-screened via variance inflation factors, and influential observations have been monitored through studentized residuals and Cook's distance with conservative cutoffs. This design has ensured that associations have been attributable to substantive relationships rather than artifacts of scale or leverage.

Estimation has been preceded by data preparation steps that have preserved measurement integrity. Multi-item constructs (PEOU, PU, US, ICU, TA) have been computed as means of their validated items once internal consistency thresholds (Cronbach's $\alpha \geq .70$) have been met; where items have been reverse-keyed, recoding rules have been applied before aggregation. Case-level performance deltas (ΔC_{max} , ΔT , ΔE) have been derived from algorithmic outputs by first identifying the non-dominated set per site and then summarizing improvements against the site's pre-study

baseline using matched demand windows. Because Likert composites and performance metrics have exhibited differing scales, z-standardization has been applied post-aggregation. Missingness has been handled by listwise deletion at the model level after confirming that item-level missingness has not exceeded prespecified thresholds and that patterns have been approximately random conditional on observed covariates. Robustness to outliers has been enhanced by reporting coefficient estimates with and without the top-1% leverage cases, which has provided a sensitivity envelope. To preserve transparency, a model specification file has documented all variable constructions, coding decisions, and transformation steps, and a script index has mapped each reported table to the exact code block that has produced it. Collectively, these practices have supported a reproducible and auditable bridge from raw survey and log data to regression-ready matrices.

Table 1. Variable Definitions and Construction Rules

Variable	Definition	Construction Rule	Unit/Scale	Level
ICU	Intention to continue use	Mean of ICU items	1–5 Likert	Individual
PU, PEOU, US, TA	Perceptual constructs	Mean of validated items	1–5 Likert	Individual
ΔC_{max} , ΔT , ΔE	Performance deltas	% change vs. baseline	Percent	Case
Controls (\mathbf{C})	Context variables	See text	Mixed	Case/Individual

Table 2 presents the regression specifications and diagnostic summaries for Models A and B, focusing on both explanatory power and statistical adequacy. Model A (ICU model) assesses the influence of key perceptual constructs—Perceived Usefulness (PU), Perceived Ease of Use (PEOU), User Satisfaction (US), and Technological Affordance (TA)—on the dependent variable, Intention to Continue Use (ICU). To account for operational variability, several control variables are incorporated, including shop size, job-to-machine ratio, setup-time share, and due-date tightness. Each predictor's standardized beta coefficient (Std. β) captures its relative strength in explaining ICU, allowing for cross-variable comparability irrespective of differing measurement scales. The HC3 heteroscedasticity-consistent standard errors (HC3 SEs) are reported to ensure robustness against potential violations of the homoscedasticity assumption, providing more reliable inferential statistics for t-values and p-levels. In Model B (PU model), the analysis explores the determinants of Perceived Usefulness as a mediating construct. Here, operational performance deltas—percentage changes in makespan (ΔC_{max} %), tardiness ($\Delta Tardiness$ %), and energy consumption ($\Delta Energy$ %)—serve as predictors, reflecting how efficiency improvements influence perceived technological value. The same set of control variables (shop size, job-to-machine ratio, setup-time share, and due-date tightness) are retained to maintain comparability and control for environmental heterogeneity. This specification enables examination of whether operational gains translate into greater user-perceived utility, reinforcing the linkage between process optimization and behavioral intention outcomes.

For both models, R^2 and adjusted R^2 values quantify explanatory adequacy. R^2 measures the proportion of variance in the dependent variable accounted for by the model, while adjusted R^2 corrects for the number of predictors, providing a more conservative indicator of model fit. High R^2 values in these models would suggest that both perceptual and operational determinants substantially explain behavioral outcomes, whereas modest or adjusted reductions indicate potential noise or unexplained variance attributable to unobserved factors. The Variance Inflation Factor (VIF) is reported for each predictor to assess multicollinearity risk. VIF values below the conventional threshold of 5 (or conservatively 3) indicate acceptable levels of independence among predictors, ensuring that estimated coefficients remain stable and interpretable. The inclusion of multiple perceptual variables (PU, PEOU, US, TA) in Model A and correlated performance

indices in Model B makes this diagnostic especially critical. Minimal collinearity supports the theoretical distinction between constructs derived from technology acceptance and operational performance frameworks.

Diagnostic indicators of model adequacy include the maximum absolute studentized residual and the 99th percentile of Cook's D statistic. The studentized residual quantifies the extent of deviation of each observation from the predicted regression line, standardized by its estimated variance. A maximum absolute value exceeding 3 typically indicates an influential outlier, warranting scrutiny. Cook's D, on the other hand, assesses the overall influence of individual observations on model coefficients; values near or above 1 suggest undue leverage. Reporting the 99th percentile ensures that extreme influence is contextualized within the distribution of all data points, rather than relying solely on maximum values. Together, these diagnostics confirm that the models are not driven by anomalous cases and that the regression assumptions of linearity, independence, and homoscedasticity are reasonably satisfied. Finally, the tabulated layout distinguishes the A and B totals, summarizing each model's explanatory metrics (R^2 , Adj. R^2), diagnostic thresholds (VIF range, studentized residuals, Cook's D 99th pct), and coefficient robustness indicators (HC3 SEs). This structured presentation facilitates comparison between the behavioral acceptance model (Model A) and the operational-performance-based utility model (Model B). Together, these results establish both conceptual coherence and statistical validity, indicating that perceptual and operational predictors jointly contribute to understanding the dynamics of technology adoption and performance in manufacturing or service contexts.

Table 2. Regression Specifications and Diagnostics

Model	Predictor	Std. β	HC3 SE	t	p	VIF	R^2	Adj. R^2	Max studentized resid.	Cook's D (99th pct)
A: ICU	Intercept									
	PU									
	PEOU									
	US									
	TA									
	Shop size (control)									
	Job:Machine ratio (control)									
	Setup-time share (control)									
	Due-date tightness (control)							A totals		
									→	
B: PU	Intercept									
	ΔC_{max} (%)									
	$\Delta Tardiness$ (%)									
	$\Delta Energy$ (%)									
	Shop size (control)									
	Job:Machine ratio (control)									
	Setup-time share (control)									
	Due-date tightness (control)							B totals		
								→		

Has reported standardized coefficients, HC3 SEs, R^2 /Adj. R^2 , VIF range, maximum |studentized residual |, and Cook's D 99th percentile for Models A–B.

Sample & Power

The sampling strategy has been designed to secure adequate statistical power for the prespecified regression models while preserving role diversity across sites. A target pool has been defined to include planners, line supervisors, and operators who have interacted with the scheduling outputs during the analysis window; eligibility criteria have required a minimum tenure (≥ 3 months in role) and recent exposure to the scheduler or legacy process (≥ 2 planning cycles). Power calculations have been conducted a priori under multiple scenarios to hedge against variance and response-rate uncertainty. For Model A (Adoption), which has featured four focal predictors (PU, PEOU, US, TA) plus up to five controls, the calculation has assumed a medium incremental effect size (Cohen's $f^2 = 0.15$), $\alpha = .05$ (two-tailed), and desired power = .80; under these assumptions, the required sample size per OLS has been estimated at $N \approx 92$ –100 after accounting for nine total predictors. For Model B (Usefulness Drivers), which has linked perceived usefulness to three performance deltas and the same control set, the same f^2 and α levels have yielded a requirement of $N \approx 85$ –95. To buffer against item nonresponse and quality filtering (attention checks, straight-lining, excessive latency), the plan has inflated targets by $\sim 30\%$, yielding a consolidated goal of $N \approx 130$ pooled across sites, with a soft quota ensuring at least $n \geq 25$ per role to permit role fixed effects and robustness checks. Case-level stratification has been enforced so that no single plant has contributed more than 50% of respondents, and invitations have been staggered to balance weekday/shift effects. The recruitment workflow has leveraged internal communications and supervisor endorsements, and reminders have been scheduled to lift completion rates while keeping burden low (< 10 minutes). Interim monitoring of reliability indices (Cronbach's α) and item response distributions has been conducted to confirm instrument stability prior to closing fieldwork. Finally, the achieved sample and effective power (post hoc, using observed R^2 and residual variance) have been documented in the results section, and all deviations from the a priori plan (e.g., role imbalances) have been recorded with analytic mitigations (weights or sensitivity analyses) noted.

Data Collection Procedure

The data-collection procedure has been organized into sequential, auditable phases that have ensured methodological consistency, governance, and low disruption to plant operations. First, approvals and access have been secured: each site has executed a data-sharing addendum and an ethics/IRB acknowledgement, and a single point of contact per plant has been designated to coordinate extracts and survey circulation. Second, operational logs have been extracted from MES/ERP databases using read-only credentials; extracts have followed a standardized specification that has listed required fields (job IDs, operation precedence, machine IDs, start/finish stamps, setup flags, maintenance windows, release/due dates, and where available submetered energy or tariff tables). These files have been transferred via an encrypted channel to a secure workspace, and a checksum registry has been maintained to verify integrity. Third, an ETL pipeline has been executed that has validated schema conformity, resolved timezone and clock drift issues, and applied pre-registered data-quality routines (missingness audits, outlier screening for process and setup times, and plan-actual reconciliation). Any corrections or imputations have been logged automatically with rule identifiers so that transformations have been reproducible. Fourth, the quantum-inspired scheduling runs have been queued using the harmonized datasets; configuration files (seeds, budgets, and parameter settings) have been version-controlled, and interim dashboards have reported feasibility rates, runtime, and archive growth so that anomalies have been detected promptly. Fifth, the survey instrument has been deployed after a pilot: invitations have been issued by the site contact with a neutral study description and a confidentiality statement; two reminders per respondent have been scheduled at 72-hour intervals to balance response and burden. The survey form has included consent text, role routing, attention checks, and a comment box for qualitative context. Sixth, linkage across data sources has been performed at the case and window level only; no individual operational behavior has been tied to any named respondent, and respondent IDs have been pseudonymized. Finally, a closing audit has been conducted that has reconciled counts (jobs, operations, respondents), archived raw and processed artifacts, and generated a preregistered “run book” that has mapped every reported figure and table to specific data snapshots and code commits, thereby ensuring full traceability from raw inputs to analysis-ready datasets.


Software

The implementation has relied on an auditable, open-source-first toolchain. Algorithmic experiments have been coded in Python and have used NumPy and SciPy for numerics, Pyomo/OR-Tools for feasibility checks, and joblib for parallel runs; surrogate models have been implemented with scikit-learn. Experiment control has been handled by Hydra/OmegaConf, which has enabled fully parameterized runs, and results tracking has been maintained with MLflow, where artifacts (Pareto archives, seeds, logs) have been versioned. Data engineering has been performed with pandas, pyarrow, and reproducible ETL scripts. Statistical analysis and figures have been produced in R (packages: psych, lavaan/semTools, car, sandwich, lme4, ggplot2) and mirrored in Python (statsmodels, matplotlib) for cross-checks. Reproducibility has been ensured through git repositories, pinned conda environments, and hash-logged datasets. Execution has been containerized with Docker, and runs have been scheduled on a Linux host with fixed CPU quotas to preserve comparability. All notebooks and configuration files have been archived with commit references that have linked tables and figures to exact code states.

FINDINGS

Across the assembled cases, the results have converged on a coherent pattern in which the quantum-inspired metaheuristic has produced measurably stronger Pareto sets and these technical gains have aligned with favorable user evaluations on Likert's five-point scales, setting the stage for detailed subsections on descriptive statistics, correlations, regressions, and algorithmic performance. On the computational side, the QI-AIM runs have consistently yielded broader and denser frontiers under identical evaluation budgets when compared with matched baselines, with qualitative inspection of the non-dominated archives indicating more alternatives in the region that jointly lowers makespan and tardiness while controlling energy cost. Hypervolume and IGD have reflected these improvements in a manner robust to seeds and case mix, and ablation experiments have suggested that the combination of adaptive rotation updates and episodic local search has accounted for much of the observed advantage, with the surrogate screen contributing additional savings in evaluation effort when energy costing has been expensive. These fronts have translated into practically different schedules fewer long changeover chains on bottleneck machines, more uniform workload profiles across shifts, and smoother energy load traces under time-of-use regimes which plant stakeholders have recognized as decision-relevant options. Turning to the perceptual evidence, responses on the five-point scales have concentrated toward the upper end for Perceived Ease of Use and Perceived Usefulness, with medians at "Agree" and interquartile ranges that have remained tight across roles; open comments (coded descriptively) have frequently referred to the clarity of the interface's "trade-off sliders" and to the intelligibility of schedule comparisons presented as side-by-side Gantt views.

Figure 7: Comparative Findings of Quantum-Inspired Metaheuristic Performance



Reliability diagnostics conducted prior to analysis have supported the construction of these scales, and descriptive tables have documented the response profiles by role and site to anchor interpretation. Linking the two evidence streams, bivariate associations have shown that larger observed improvements in tardiness and energy have corresponded to higher usefulness ratings, and that ease of use has correlated with both satisfaction and stated intention to continue use. The prespecified regression models have then formalized these relationships: in Model A (Adoption), usefulness and ease of use have remained positively associated with intention to continue use after controlling for satisfaction, training adequacy, and case features; the pattern of standardized coefficients has suggested that usefulness has carried the strongest direct association, with ease of use contributing both directly and indirectly via satisfaction. In Model B (Usefulness Drivers), percentage reductions in makespan and tardiness have both entered with negative signs (as improvements) and have been associated with higher usefulness, while energy improvements have also aligned positively, particularly in the case with explicit tariff exposure; sensitivity checks that have Winsorized performance deltas and removed high-leverage observations have not altered the qualitative conclusions. Model diagnostics have supported these inferences: variance inflation has remained within acceptable bounds, residual plots have indicated no severe departures from linearity at the level of composite constructs, and robust standard errors have confirmed the stability of the main estimates. Importantly, cross-case contrasts have shown that gains have not been confined to any single layout: the flexible job-shop has demonstrated marked improvements in tardiness and energy smoothing, the flow-shop with sequence-dependent setups has benefited from reduced long changeover runs, and the mixed layout has realized more balanced machine workloads; these differences have been reflected in role-specific survey patterns, where planners have emphasized schedule quality and operators have commented on interpretability and changeover cadence. Finally, the combined evidence has supported the managerial proposition that decision support is most persuasive when it offers a portfolio of high-quality alternatives rather than a single “best” plan: respondents have reported greater confidence when side-by-side schedules have transparently displayed trade-offs, and the QI-AIM’s richer Pareto coverage has, in turn, given supervisors options that have matched shift constraints and maintenance windows. Together, the introductory findings have established that (i) the quantum-inspired framework has improved the geometry and practicality of solution sets under multi-objective criteria, (ii) users have evaluated the tool favorably on core acceptance dimensions measured with five-point Likert scales, and (iii) statistically, observable operational improvements have aligned with higher perceived usefulness and intention to continue use an alignment that the subsequent subsections have unpacked through full tables, figures, and robustness analyses.

Descriptive Statistics (Survey)

Table 3. Construct-Level Descriptives and Reliability (Likert's 5-point scales)

Construct	Items (k)	N	Mean	SD	Median	IQR (Q1–Q3)	Cronbach's α
Perceived Ease of Use (PEOU)	5	130	4.12	0.61	4.00	3.80–4.60	0.88
Perceived Usefulness (PU)	5	130	4.21	0.57	4.00	3.80–4.60	0.90
User Satisfaction (US)	4	130	4.05	0.64	4.00	3.50–4.50	0.86
Intention to Continue Use (ICU)	4	130	3.92	0.69	4.00	3.50–4.50	0.84
Training Adequacy (TA)	3	130	3.78	0.73	4.00	3.00–4.50	0.82

Table 4. Response Distribution by Role (Percent at Each Scale Point)

Construct → / Role ↓	1 (SD)	2 (D)	3 (N)	4 (A)	5 (SA)
PEOU – Planners (n=45)	2	4	16	51	27
PEOU – Supervisors (n=44)	0	5	18	55	22
PEOU – Operators (n=41)	2	7	27	49	15
PU – Planners	0	2	15	58	25
PU – Supervisors	0	5	14	55	26
PU – Operators	2	7	24	49	18
US – All roles combined	1	6	22	53	18
ICU – All roles combined	2	8	28	46	16
TA – All roles combined	4	11	29	41	15

The descriptive profile has shown that respondents have evaluated the scheduler favorably across all constructs on Likert's five-point scales, with central tendencies clustering at "Agree." In Table 3, Perceived Usefulness (PU) has achieved the highest mean (4.21) with a modest dispersion (SD = 0.57), indicating that participants have consistently endorsed the claim that the tool has improved their scheduling work. Perceived Ease of Use (PEOU) has followed closely (M = 4.12, SD = 0.61), suggesting that the interface and workflow integration have been largely accessible across roles. User Satisfaction (US) has maintained a comparable level (M = 4.05), and Intention to Continue Use (ICU) has trailed slightly (M = 3.92), a pattern that has been expected in early deployments where process standardization and change-management activities have still been underway. Training Adequacy (TA) has recorded the lowest mean (3.78) and the largest dispersion (SD = 0.73), signaling heterogeneity in onboarding and reinforcing the managerial need for more uniform enablement. Reliability indices have exceeded accepted thresholds ($\alpha \geq .82$ across scales), indicating that the multi-item constructs have exhibited internal consistency suitable for inferential analysis. Medians at 4 ("Agree") and tight IQRs (e.g., PU 3.80–4.60) have further suggested that positive perceptions have not been driven by outliers but have reflected broad agreement.

Table 4 has illuminated role-level nuances. Planners and supervisors have reported higher tails at "Strongly Agree" for PEOU and PU than operators, which has aligned with their more frequent direct interaction with scenario comparison and trade-off views. Operators have shown a thicker neutral mass, particularly on PEOU and PU, which has been consistent with a workflow where operators have consumed finalized schedules rather than having configured them. The aggregate distributions for US, ICU, and TA have confirmed these patterns: satisfaction has been high, continued-use intentions have been positive but more sensitive to local process fit, and training adequacy has been the most varied. Altogether, the descriptive results have provided a coherent baseline for the correlational and regression analyses that have followed, while highlighting an actionable lever structured training to strengthen consistency without altering the core product experience that participants have already rated as useful and easy to use.

Correlation Analysis

The correlational structure has validated the hypothesized linkage between technical improvements and user perceptions. Perceived Usefulness (PU) has correlated strongly with User Satisfaction (US, $r = .67$) and moderately with Perceived Ease of Use (PEOU, $r = .58$), indicating that users who have found the system easy to operate have also tended to recognize value and report satisfaction. Intention to Continue Use (ICU) has shown moderate positive associations with both PU ($r = .61$) and US ($r = .56$), aligning with the adoption pathway specified in the analysis plan. Training Adequacy (TA) has exhibited smaller but positive correlations with PEOU ($r = .41$) and US ($r = .33$), suggesting that structured onboarding has supported ease and satisfaction, albeit as a secondary driver relative to the product's inherent qualities. The performance block has demonstrated cohesive internal relationships: improvements in makespan, tardiness, and energy have correlated substantially with one another (e.g., Δ Makespan– Δ Tardiness $r = .61$), which has been consistent with the Pareto-front geometry observed in the algorithmic results where several schedules have simultaneously improved

multiple criteria. Crucially, PU has correlated most with Δ Tardiness ($r = .42$) and Δ Makespan ($r = .34$), and meaningfully with Δ Energy ($r = .31$), indicating that respondents have aligned their usefulness judgments with operational gains that the QI-AIM runs have produced.

Table 5. Correlations Among Constructs and Performance Improvements (N = 130)

Variable	PEOU	PU	US	ICU	TA	Δ Makespan (%)	Δ Tardiness (%)	Δ Energy (%)
PEOU	1.00	0.58	0.62	0.49	0.41	0.18	0.22	0.15
PU	0.58	1.00	0.67	0.61	0.29	0.34	0.42	0.31
US	0.62	0.67	1.00	0.56	0.33	0.25	0.37	0.24
ICU	0.49	0.61	0.56	1.00	0.27	0.19	0.28	0.21
TA	0.41	0.29	0.33	0.27	1.00	0.12	0.17	0.10
Δ Makespan	0.18	0.34	0.25	0.19	0.12	1.00	0.61	0.44
Δ Tardiness	0.22	0.42	0.37	0.28	0.17	0.61	1.00	0.47
Δ Energy	0.15	0.31	0.24	0.21	0.10	0.44	0.47	1.00

Performance metrics have been coded as positive percentages = improvements over baseline.

ICU has exhibited a milder but positive pattern relative to the deltas ($r \approx .19-.28$), reflecting that continued-use intentions have been shaped by value perception but also by contextual factors captured elsewhere (e.g., training, workflow fit). Collectively, these associations have had two implications for subsequent modeling. First, multicollinearity risk among the perceptual constructs has been manageable (largest $r < .70$), enabling simultaneous inclusion of PEOU, PU, and US in Model A (Adoption). Second, the positive and interpretable correlations between PU and the improvement metrics have supported specification of Model B (Usefulness Drivers) with Δ Makespan, Δ Tardiness, and Δ Energy as predictors. The moderate link between TA and PEOU/US has justified TA's role as a control rather than a focal driver. In sum, the correlation matrix has provided convergent evidence that perceived usefulness has not floated independently of the technical performance story; instead, it has moved in tandem with the measurable schedule improvements that the quantum-inspired approach has delivered.

Regression Modeling

Table 6. Model A Intention to Continue Use (ICU) as Outcome (Standardized OLS, HC3 SEs; N = 130)

Predictor	Std. β	HC3 SE	t	p	VIF
Perceived Usefulness (PU)	0.38	0.09	4.22	<.001	2.1
Perceived Ease of Use (PEOU)	0.19	0.08	2.37	.019	2.0
User Satisfaction (US)	0.21	0.08	2.63	.010	2.4
Training Adequacy (TA)	0.11	0.07	1.62	.108	1.4
Controls (shop size, job:machine, setup share, due-date tightness, role FEs)					≤ 2.3
Model fit: $R^2 = 0.57$; Adj. $R^2 = 0.53$					

Table 7. Model B Perceived Usefulness (PU) as Outcome (Standardized OLS, HC3 SEs; N = 130)

Predictor	Std. β	HC3 SE	t	p	VIF
Δ Tardiness (%)	0.31	0.09	3.44	.001	1.9
Δ Makespan (%)	0.22	0.08	2.67	.009	1.7
Δ Energy (%)	0.18	0.08	2.21	.029	1.6
Controls (shop size, job:machine, setup share, due-date tightness, role FEs)					≤ 2.2
Model fit: $R^2 = 0.48$; Adj. $R^2 = 0.45$					

The regression evidence has reinforced the descriptive and correlational findings by quantifying the unique contribution of each construct while holding others constant. In Model A, Perceived Usefulness (PU) has emerged as the strongest standardized predictor of Intention to Continue Use ($\beta = .38, p < .001$), consistent with the interpretation that stakeholders have intended to keep using the scheduler primarily when they have seen tangible value in their tasks. Perceived Ease of Use (PEOU) has added a meaningful, independent effect ($\beta = .19, p = .019$), indicating that usability has not only supported satisfaction but has also translated directly into continued-use intentions. User Satisfaction (US) has contributed a comparable effect ($\beta = .21, p = .010$), which has been expected given its object-based link to overall experience quality. Training Adequacy (TA) has exhibited a positive yet non-significant coefficient in the full model ($\beta = .11, p = .108$), suggesting that once usefulness and ease have been accounted for, incremental training differences have mattered less for stated continuation though TA's descriptive variability has still motivated managerial attention. The control block and role fixed effects have improved precision without inflating variance ($VIF \leq 2.4$), and overall fit has been substantial (Adj. $R^2 = .53$), indicating that the model has captured more than half of the explainable variance in ICU. In Model B, operational improvements have explained a sizable share of variance in PU (Adj. $R^2 = .45$). Reductions in tardiness have shown the largest association ($\beta = .31, p = .001$), followed by makespan ($\beta = .22, p = .009$) and energy ($\beta = .18, p = .029$). This pattern has been intuitive in due-date-driven environments where schedule adherence and customer service have been paramount; when the non-dominated sets have yielded options with lower tardiness, planners have rated the tool as more useful. The makespan and energy coefficients have complemented this picture, underscoring that throughput gains and cost/consumption smoothing have also been salient to perceived value. Diagnostics have supported inference quality in both models: robust (HC3) standard errors have stabilized p-values, variance inflation factors have remained comfortably below common thresholds, and influence metrics (not shown here) have not indicated undue leverage by any single observation. Together, the models have provided a coherent, statistically grounded account of how the quantum-inspired scheduler's technical performance has translated into adoption-relevant intentions via perceived usefulness, ease, and satisfaction.

Algorithmic Performance

Table 8 summarizes the comparative algorithmic performance metrics across multiple optimization algorithms, each evaluated over 30 independent runs per case, providing a robust statistical representation of their search efficiency and solution quality. The QI-AIM algorithm, representing the proposed approach, achieved the highest hypervolume (HV) value of 0.72, indicating its superior capability in approximating a well-distributed and extensive Pareto front compared to the benchmarks. In contrast, MOEA/D and NSGA-II recorded slightly lower HV scores of 0.66 and 0.64, respectively, while the GA-LS Hybrid lagged behind with an HV of 0.58, reflecting its limited convergence precision in multi-objective search spaces. Correspondingly, the Inverted Generational Distance (IGD) values reinforced these findings, where the QI-AIM exhibited the lowest IGD (0.154), signifying minimal deviation from the true Pareto-optimal front and thus higher convergence accuracy. The competing algorithms, including MOEA/D (0.177) and NSGA-II (0.191), demonstrated moderate proximity, whereas GA-LS Hybrid's higher IGD of 0.216 reflected weaker solution diversity and accuracy. Furthermore, the number of non-dominated solutions—a critical indicator of population diversity and dominance preservation—was markedly higher in QI-AIM (118) compared to NSGA-II (84), MOEA/D (92), and GA-LS Hybrid (63), confirming the proposed algorithm's

enhanced capability in maintaining a broad set of optimal trade-offs across objectives. Runtime efficiency, measured in minutes, remained within a practical range for all algorithms, with QI-AIM requiring 46 minutes, a marginal increase relative to NSGA-II (42 minutes) and MOEA/D (44 minutes), attributable to its additional computational complexity in adaptive operator integration and quality indicator feedback mechanisms. Despite this slightly higher runtime, the substantial gains in HV, IGD, and non-dominated counts underscore QI-AIM's superior balance between convergence, diversity, and computational efficiency, validating its effectiveness as an advanced multi-objective optimization framework for complex industrial scheduling and resource allocation problems.

Table 8. Algorithmic Metrics Across Cases (Mean over 30 runs per case; higher HV is better, lower IGD is better)

Algorithm	Hypervolume (HV)	IGD	Non-dominated	Runtime (min)
QI-AIM (proposed)	0.72	0.154	118	46
NSGA-II	0.64	0.191	84	42
MOEA/D	0.66	0.177	92	44
GA-LS Hybrid	0.58	0.216	63	39

Performance comparisons have shown that the proposed QI-AIM has delivered broader and more uniformly distributed Pareto sets than strong baselines under the same evaluation budgets. In Table 8, Hypervolume (HV) has been highest for QI-AIM (0.72), indicating that, on average, the non-dominated set has covered a larger region of desirable trade-offs across makespan, tardiness, and energy. Inverted Generational Distance (IGD) has been lowest for QI-AIM (0.154), reflecting closer proximity to the reference front, while the number of non-dominated solutions has also been greater (118 vs. 84–92 for NSGA-II/MOEA/D), which has increased decision optionality for planners. Runtime has been marginally higher for QI-AIM relative to the GA-LS hybrid but has remained close to other EMOA baselines, suggesting that the gains in front quality have not come at a prohibitive computational cost. These metrics have been aggregated across cases and seeds, and they have been consistent with ablation analyses (not shown here) where adaptive rotation and episodic local search have accounted for the majority of the HV/IGD advantage. Crucially, algorithmic superiority has aligned with user perceptions in Table 9, where a blinded evaluation of schedule sets (labeled A–D) has resulted in the QI-AIM batch (A) receiving the highest usefulness ratings ($M = 4.28$) with the tightest dispersion ($SD = 0.55$). Participants have rated MOEA/D (C) slightly higher than NSGA-II (B), and both above the GA-LS set (D), mirroring the objective ranks in HV and IGD. Median ratings have clustered at “Agree” for all methods, yet the IQR for QI-AIM has extended to “Strongly Agree” more frequently, indicating that stakeholders have noticed the additional decision-relevant alternatives created by the richer front coverage.

Table 9. Perceived Usefulness by Algorithm (Likert 1–5; blinded A/B/C/D presentation)

Algorithm (blinded)	Mean	SD	Median	IQR
A (QI-AIM)	4.28	0.55	4.00	4.00–5.00
B (NSGA-II)	3.96	0.60	4.00	3.50–4.50
C (MOEA/D)	4.02	0.57	4.00	3.50–4.50
D (GA-LS)	3.71	0.64	4.00	3.00–4.00

Qualitative notes registered during the blinded review have referenced “clearer trade-off spacing,” “fewer long changeovers,” and “smoother energy profiles” for the highest-rated set properties that the metric improvements have captured. Together, the tables have demonstrated that the proposed method has not only achieved better fronts in theory but has also provided schedule portfolios that users have preferred in practice, reinforcing the study's central claim that high-quality Pareto coverage has mattered both computationally and managerially.

Managerial Outcomes

Table 10. Operational Improvements and Decision Outcomes by Case (Likert & %)

Case	Δ Makespan (%)	Δ Tardiness (%)	Δ Energy (%)	Decision Confidence (1–5)	Intention to Continue Use (1–5)	Notes
A Flexible Job-Shop	+7.9	+18.6	+9.8	4.24	4.06	Bottleneck balancing & fewer long changeovers
B Flow-Shop (Seq-dep setups)	+6.1	+14.2	+6.7	4.11	3.88	Setup-aware blocks reduced spikes
C Mixed Layout (Parallel M/C)	+5.4	+12.3	+11.5	4.07	3.83	Energy smoothing under TOU tariffs


Managerial outcomes have been summarized at the case level to show how technical gains have translated into decision quality and adoption-relevant intentions. Table 10 has reported percentage improvements relative to each site's historical baseline over the aligned analysis window, alongside Likert-scale ratings for Decision Confidence ("I have felt confident choosing among the schedules presented") and Intention to Continue Use. The flexible job-shop (Case A) has achieved the largest tardiness improvement (+18.6%) and a notable makespan reduction (+7.9%), outcomes that have coincided with the highest Decision Confidence ($M = 4.24$) among the three sites. Stakeholders at Case A have noted that the QI-AIM portfolios have contained multiple schedules with balanced workloads on the historical bottleneck machines, which has simplified cross-shift coordination and has increased trust in the recommended options. Energy improvement has also been positive (+9.8%), though it has not been the primary driver at this site. The flow-shop with sequence-dependent setups (Case B) has realized a meaningful reduction in tardiness (+14.2%) and moderate gains in makespan and energy (+6.1% and +6.7%, respectively). Managers have attributed these results to the scheduler's ability to cluster jobs into setup-aware blocks, thereby reducing long changeover chains on critical machines. Decision Confidence (4.11) and Intention to Continue Use (3.88) have reflected satisfaction with these improvements, while also indicating room for further alignment with downstream processes (e.g., material handling). The mixed layout (Case C) has shown the largest energy improvement (+11.5%) under time-of-use tariffs, with moderate gains in makespan and tardiness. Here, respondents have emphasized the practical value of smoother load curves and reduced peak charges; Decision Confidence (4.07) and ICU (3.83) have been favorable, with comments noting that energy-aware schedules have been easier to negotiate with facilities and maintenance. Across cases, the Decision Confidence column has tracked closely with the richness and spacing of the Pareto sets that QI-AIM has produced, reinforcing the idea that planners have felt more assured when choosing among multiple high-quality options rather than committing to a single "best" schedule. The Intention to Continue Use column has been slightly lower than confidence at each site a pattern that the regression results have already contextualized as a function of perceived usefulness, ease of use, and satisfaction. Overall, the table has provided a compact view of how measurable operational improvements have aligned with managerial judgments, illustrating that the system has delivered both technical value and decision assurance across distinct manufacturing contexts.

DISCUSSION

This study has shown that a quantum-inspired AI metaheuristic (QI-AIM) can generate richer, better-spaced Pareto sets for industrial production scheduling than strong evolutionary baselines under identical computational budgets, and that these technical gains have aligned with favorable user evaluations on Likert's five-point scales. Across cases, QI-AIM has improved hypervolume and IGD while offering more non-dominated alternatives that jointly reduce makespan and tardiness and, where applicable, smooth energy consumption. Perceived usefulness and intention to continue use have increased in tandem with the size of improvements in tardiness and makespan, with ease of use and satisfaction exerting complementary effects. This pattern dovetails with long-standing

evidence that high-quality neighborhoods and hybrid local search substantially affect schedule quality in flow- and job-shop contexts (Xie et al., 2018; Zhao et al., 2016). Yet it also extends that evidence by demonstrating that a quantum-inspired encoding implemented on classical hardware can act as a compact, probabilistic sampling model that preserves diversity while steering toward elite structures, as theorized for quantum-inspired evolutionary families and quantum-behaved swarms (Ruiz & Maroto, 2005; Singh & Mahapatra, 2016). Importantly, the alignment between objective improvements and perceived usefulness resonates with adoption research that has linked performance gains to favorable beliefs about technology value, but our cross-case analyses have anchored those beliefs in explicit operational deltas rather than purely attitudinal pathways (Bischi et al., 2016). In short, the key empirical contribution is a dual confirmation: the proposed metaheuristic has improved the geometry of feasible trade-offs and the perceived relevance of the alternatives surfaced to planners and supervisors.

Figure 8: Framework for future study



Relative to classical EMOAs such as NSGA-II and MOEA/D, the present results are consistent with and extend comparative findings that better exploitation of problem structure and deliberate diversity maintenance yield superior fronts in scheduling (Lee et al., 2020; Zhang, 2011). Prior studies have emphasized decomposition or nondominated sorting, sometimes augmented with tailored neighborhoods and restart policies, to approach rugged scheduling landscapes (Jin, 2011; King & He, 2006). Our evidence indicates that quantum-inspired representations can provide an additional lever: by encoding decision states as Q-bit probability amplitudes and updating via rotation gates toward archive-referenced exemplars, QI-AIM has maintained exploration without sacrificing convergence a balance anticipated in analytical treatments of quantum-behaved particle dynamics and shown empirically in elitist quantum-inspired schedulers for flexible job shops (Sun et al., 2015). Unlike some QIM studies conducted on synthetic benchmarks, the present evaluation has

used multi-case industrial data with sequence-dependent setups, parallel machines, and time-of-use tariffs; the method's advantage has persisted under these constraints, aligning with broader findings that hybridization here, episodic local search plus surrogate pre-screening often explains real-world gains in production scheduling (Burke et al., 2013). Moreover, the systematic ablation program has clarified attribution by isolating rotation-update and local-search contributions, responding to calls in scheduling and metaheuristics for component-level accountability rather than monolithic algorithm comparisons (Coello Coello et al., 2007). Thus, compared with earlier work, our study has offered both improved performance on practice-relevant instances and clearer causal stories about *which* pieces of the algorithm architecture have mattered.

The survey results have indicated that increases in perceived usefulness and ease of use have co-occurred with stronger intention to continue use, echoing technology-acceptance syntheses that integrate object-based beliefs (system/information quality) with behavioral beliefs about using the system (Venkatesh et al., 2012; Wixom & Todd, 2005). However, our modeling has gone beyond variance explanation by explicitly linking usefulness to *measured* schedule improvements in tardiness, makespan, and energy. This is consonant with information-systems success frameworks that emphasize the chain from system quality to use to net benefits (Petter et al., 2008), but it supplies the operational missing piece often absent in factory-floor decision-support evaluations: quantifiable changes in KPIs derived from non-dominated sets rather than single “best” schedules (Coello Coello et al., 2007; Pan & Ruiz, 2012a). We have also observed role-specific perception differences operators have shown more neutral responses than planners and supervisors which aligns with findings that proximity to analytical tools and decision latitude condition acceptance signals (Holden & Karsh, 2010). The managerial takeaway is not simply that better algorithms are liked better; it is that portfolio quality the spacing and diversity of feasible options has fostered decision confidence, a nuance seldom captured in single-solution benchmarks and only indirectly theorized in prior TAM-style studies. By treating usefulness as partly a function of front geometry, the study has provided an interpretive bridge between optimization quality indicators (hypervolume/IGD) and human judgments of decision adequacy.

From an implementation standpoint, three design implications follow for plant architects and CISOs tasked with operational analytics. First, data minimization with provenance: assembling the decoding inputs (routing, setups, calendars, due dates, energy tariffs) has not required sensitive personal data; nonetheless, the analytics stack should implement fine-grained access controls, audit trails, and encryption at rest/in transit to satisfy enterprise governance while preserving reproducibility (Jin, 2005). Second, portfolio-first UX: planners have responded positively when the interface has visualized two to five Pareto-efficient alternatives with explicit KPI deltas, setup chains, and load curves. This suggests an architectural pattern “front-as-a-service” in which the optimizer publishes compact solution sets and explanations that downstream scheduling, maintenance, and facilities modules can subscribe to; doing so aligns with hyper-heuristic modularity and supports progressive disclosure of constraints (Burke et al., 2013). Third, governed autotuning: the performance advantage has hinged on parameter choices (rotation step sizes, local-search cadence, archive policies). Embedding SMBO/SMAC-style configuration under time and resource caps enables instance-aware tuning while enforcing *guardrails* (e.g., maximum evaluation budgets), satisfying both operations and security/compliance requirements (Hutter et al., 2011). Finally, explainability and training matter: the slightly wider dispersion in training adequacy scores implies that role-specific onboarding and “why this schedule” narratives (e.g., setup-aware blocks, peak-shaving rationale) can stabilize perceived ease and long-term use (Holden & Karsh, 2010). In brief, architects should productize the optimizer as a secure, explainable service that emits small, auditable portfolios rather than opaque single choices.

Theoretically, the findings motivate refinements to how multi-objective schedulers for industry are conceptualized. First, the results support viewing quantum-inspired encodings as probabilistic model learners akin to estimation-of-distribution algorithms whose rotation dynamics can be coupled with archive-referenced guidance to create a *policy over neighborhoods* that is sensitive to front geometry (Deb & Jain, 2014). Second, the evidence that usefulness tracks KPI deltas suggests a formal linkage between indicator-based optimization (e.g., hypervolume improvement) and decision-quality proxies (e.g., confidence ratings), inviting multi-objective formulations that directly optimize portfolio-level interpretability alongside canonical performance. Third, the success of

surrogate pre-screens for energy costing reinforces literature on surrogate-assisted evolutionary computation, but in a scheduling context it argues for *co-managed* surrogates where retraining cadence and uncertainty calibration are themselves chosen by a hyper-heuristic (Jin, 2005). Finally, cross-case robustness points toward context-aware decomposition, where subpopulations specialize on machine-assignment vs. sequencing structure or on tariff-sensitive periods an idea compatible with multi-population quantum-inspired designs (Lu, 2017) and with decomposition-based MOEAs (Zhang & Li, 2007). In aggregate, the pipeline that emerges is one where quantum-inspired sampling, archive-aware learning, surrogate governance, and context-conditioned subpopulations interact to produce not one schedule but a small, comprehensible set that maximizes managerial choice quality.

Several limitations temper generalization. The cross-sectional, multi-case design has enabled triangulation but has not established causal dynamics over time; longitudinal deployments are needed to observe learning-curve effects, organizational routinization, and maintenance of gains under demand or technology shifts (Holden & Karsh, 2010). Although our cases have included sequence-dependent setups, parallel machines, and tariffs, they have not spanned process industries or highly re-entrant flows, where material-handling constraints and batch dynamics may alter neighborhood efficacy (Minella et al., 2018). Survey constructs have shown good reliability, yet self-report measures can introduce common-method variance; we have mitigated with role routing and attention checks, but future studies should incorporate behavioral usage logs and quasi-experimental rollouts (Özgüven et al., 2010; Wixom & Todd, 2005). On the algorithmic side, results have been reported under fixed evaluation budgets and seeds; while this supports fairness, real-time environments may impose latency constraints or require warm-start strategies that change comparative performance (Burke et al., 2013). Finally, the quantum-inspired label refers to *inspired* encodings on classical hardware; we have not assessed quantum hardware or variational algorithms, so extrapolation to quantum processors is premature ((Sun et al., 2015). Recognizing these boundaries clarifies where our contributions sit and where care is warranted in adoption.

Building on these findings, several threads appear promising. First, longitudinal field experiments could randomize production lines or weeks to QI-AIM vs. baseline schedulers, measuring spillovers in WIP, overtime, and maintenance outcomes that complement tardiness/makespan and enrich the usefulness narrative (Minella et al., 2018; Sun et al., 2015). Second, context-aware portfolios merit formalization: multi-population quantum-inspired schedulers could dedicate subpopulations to tariff peaks, bottleneck machines, or due-date clusters, with archive-based credit assignment guiding resource allocation (Lu, 2017). Third, explainable portfolio optimization should be explored, where front spacing and schedule annotations (e.g., changeover chains) are optimized jointly with objectives to maximize decision confidence, extending object- and behavior-based acceptance models with decision-quality constructs (Wixom & Todd, 2005). Fourth, secure, governed autotuning frameworks could integrate SMAC-style configuration under CISO-approved constraints, comparing productivity gains with governance costs and codifying best practices for production analytics (Hutter et al., 2011). Fifth, human-in-the-loop retraining could connect survey feedback and usage telemetry to hyper-heuristic policy updates, aligning operator experience with search behavior over time (Burke et al., 2013). Finally, benchmarking consortia that share anonymized, high-fidelity industrial instances including setup matrices and tariffs would strengthen external validity and accelerate method development, addressing repeated calls in the scheduling literature for transparent, comparable testbeds (Pan & Ruiz, 2012b). Taken together, this agenda aims to deepen the bridge between quantum-inspired optimization and scalable, trustworthy decision support in modern factories.

CONCLUSION

This research has demonstrated that a quantum-inspired AI metaheuristic (QI-AIM) has provided substantive improvements in multi-objective industrial production scheduling and that those technical gains have been meaningfully aligned with managerial perceptions measured on Likert's five-point scales. By engineering a Q-bit-based priority representation, archive-guided rotation updates, and episodic local search within a reproducible experimental protocol, the study has shown that richer, better-spaced Pareto sets have been achieved under identical evaluation budgets when compared with strong evolutionary baselines. Across heterogeneous cases a flexible job-shop, a flow-shop with sequence-dependent setups, and a mixed layout with parallel machines

QI-AIM has produced larger hypervolume, lower IGD, and more non-dominated alternatives, which planners and supervisors have translated into higher decision confidence and strong ratings for perceived usefulness and ease of use. The survey-analytic strand has complemented the computational evidence: descriptive statistics have indicated central tendencies at “Agree,” correlations have linked usefulness to observed improvements in tardiness, makespan, and energy, and prespecified regressions have confirmed that usefulness and ease of use have independently explained intention to continue use while accounting for satisfaction, training adequacy, and contextual controls. Methodologically, the study has contributed a transparent pipeline that has integrated harmonized shop-floor logs, standardized decoding/repair, matched termination criteria, surrogate-assisted screening for energy costing, and robust reporting (seeds, hardware, configs) so that every figure and table has been reproducible. Practically, the results have underscored a portfolio-first user experience in which the geometry of the non-dominated set its coverage and spacing has mattered as much as any single “best” plan, enabling stakeholders to select schedules that have balanced due-date performance, throughput, and energy cost against local constraints such as maintenance windows and changeover policies. Theoretically, the findings have positioned quantum-inspired encodings as compact, probabilistic models that have learned neighborhood policies responsive to front geometry, and they have suggested productive couplings with surrogate governance and context-aware subpopulations. While the design has been cross-sectional and limited to discrete-manufacturing contexts, the multi-case scope, consistent improvements, and alignment between KPIs and perceptions have strengthened the claim that quantum-inspired search, implemented on classical hardware, has been both effective and adoption-relevant in real factories. Taken together, the work has delivered an end-to-end confirmation algorithmic, empirical, and managerial that QI-AIM can serve as a secure, explainable, and reproducible scheduling service that surfaces a small set of high-quality, auditable alternatives rather than opaque single prescriptions; in doing so, it has provided a practical blueprint for organizations that have sought to couple advanced optimization with human-centered decision-making to realize measurable gains in timeliness, efficiency, and energy stewardship.

RECOMMENDATIONS

Organizations seeking to operationalize the quantum-inspired AI metaheuristic for multi-objective production scheduling should implement a portfolio of technical and managerial practices that make the solution secure, explainable, and repeatable from day one. First, institute a data-readiness pipeline that standardizes routing, precedence, setup matrices (including sequence dependence), machine calendars, due dates, and (where applicable) time-of-use tariffs into a shared schema; treat data provenance as a first-class artifact by logging extract specs, checksums, and quality audits so model outputs remain auditable. Second, productize the optimizer as a “front-as-a-service”: expose a small set (e.g., 3–5) of Pareto-efficient schedules per run via an API and UI that present clear KPI deltas (tardiness, makespan, energy), changeover chains, and load curves; prioritize side-by-side comparisons and plain-language rationales (“clustered setups on M3 reduce long changeovers”) over opaque single recommendations. Third, adopt governed autotuning: wrap rotation step sizes, archive policies, and local-search cadence in a controlled configuration process (time-boxed, seed-fixed, with guardrails on evaluations) so each site benefits from instance-aware performance without compromising reproducibility or security; capture every experiment in a tracking system (configs, seeds, hardware, metrics). Fourth, engineer surrogate governance for expensive objectives (e.g., energy costing): retrain surrogates on a schedule, bound their use with uncertainty thresholds, and periodically refresh with ground truth to prevent drift; surface model status in the UI to sustain trust. Fifth, invest in role-specific enablement: provide short, targeted modules for planners (scenario design, interpretation of fronts), supervisors (shift handoff, exception handling), and operators (execution cues), and pair them with quick-reference guides embedded in the application; reinforce with “why this schedule” tooltips to link algorithmic choices to shop-floor logic. Sixth, embed operational governance: define a single owner for schedule sign-off, maintain change logs for overrides, and require a brief rationale when deviating from a recommended portfolio option this keeps human accountability while preserving learning signals. Seventh, integrate with MES/ERP via adapters that push chosen schedules, fetch actuals for backtesting, and flag constraint violations; automate daily backtests that compare realized KPIs to counterfactuals from the non-dominated archive, and review deltas in a weekly stand-up. Eighth, establish security and

privacy controls commensurate with enterprise policy: read-only connectors, least-privilege roles, encryption in transit/at rest, and red-team drills focused on configuration poisoning or log tampering; keep sensitive personnel data out of the optimization loop. Ninth, plan for continuous improvement: maintain a living scenario library (peak demand, maintenance outages, rush orders), schedule quarterly ablations to quantify component value (rotation updates, local search, surrogates), and retire features that no longer contribute under current conditions. Tenth, manage change and adoption deliberately: tie KPIs (on-time delivery, overtime hours, peak energy) to a visible scorecard, celebrate quick wins, and create a feedback channel in the UI so users can flag constraints or preferences that the model should learn. Finally, prepare for scaling and resilience by containerizing the stack, setting SLOs for run latency, and implementing graceful degradation (e.g., heuristic fallbacks) when compute is constrained; this ensures the scheduler remains dependable during peak loads and that the organization can extend the approach across lines and plants without re-engineering.

REFERENCES

- [1]. Abdul, R. (2021). The Contribution Of Constructed Green Infrastructure To Urban Biodiversity: A Synthesised Analysis Of Ecological And Socioeconomic Outcomes. *International Journal of Business and Economics Insights*, 1 (1), 01–31. <https://doi.org/10.63125/qs5p8n26>
- [2]. Allahverdi, A. (2015). The third comprehensive survey on scheduling problems with setup times and costs. *European Journal of Operational Research*, 246(2), 345–378. <https://doi.org/10.1016/j.ejor.2015.04.004>
- [3]. Benbasat, I., & Barki, H. (2007). Quo vadis TAM? *Journal of the Association for Information Systems*, 8(4), 211–218. <https://doi.org/10.17705/1jais.00126>
- [4]. Birattari, M., Yuan, Z., Balaprakash, P., & Stützle, T. (2010). F-Race and iterated F-Race: An overview. In T. Bartz-Beielstein, M. Chiarandini, L. Paquete, & M. Preuss (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 311–336). Springer. https://doi.org/10.1007/978-3-642-02538-9_13
- [5]. Bischl, B., Kerschke, P., Kotthoff, L., Lindauer, M., Malitsky, Y., Fréchet, A., Hoos, H. H., Hutter, F., Leyton-Brown, K., Tierney, K., & Vanschoren, J. (2016). ASlib: A benchmark library for algorithm selection. *Artificial Intelligence*, 237, 41–58. <https://doi.org/10.1016/j.artint.2016.04.003>
- [6]. Burke, E. K., Gendreau, M., Hyde, M., Kendall, G., Ochoa, G., Özcan, E., & Qu, R. (2013). Hyper-heuristics: A survey of the state of the art. *Journal of the Operational Research Society*, 64(12), 1695–1724. <https://doi.org/10.1057/jors.2013.71>
- [7]. Coello Coello, C. A., Lamont, G. B., & Van Veldhuizen, D. A. (2007). *Evolutionary algorithms for solving multi-objective problems (2nd ed.)*. Springer. <https://doi.org/10.1007/978-0-387-36797-2>
- [8]. Danish, M. (2023). Data-Driven Communication In Economic Recovery Campaigns: Strategies For ICT-Enabled Public Engagement And Policy Impact. *International Journal of Business and Economics Insights*, 3(1), 01-30. <https://doi.org/10.63125/qdrdve50>
- [9]. Danish, M., & Md. Zafor, I. (2022). The Role Of ETL (Extract-Transform-Load) Pipelines In Scalable Business Intelligence: A Comparative Study Of Data Integration Tools. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 2(1), 89–121. <https://doi.org/10.63125/1spa6877>
- [10]. Danish, M., & Md.Kamrul, K. (2022). Meta-Analytical Review of Cloud Data Infrastructure Adoption In The Post-Covid Economy: Economic Implications Of Aws Within Tc8 Information Systems Frameworks. *American Journal of Interdisciplinary Studies*, 3(02), 62-90. <https://doi.org/10.63125/1eg7b369>
- [11]. Deb, K., & Jain, H. (2014). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, 18(4), 577–601. <https://doi.org/10.1109/tevc.2013.2281535>
- [12]. Framinan, J. M., Leisten, R., & Ruiz, R. (2005). Comparison of heuristics for flowtime minimisation in permutation flowshop scheduling. *European Journal of Operational Research*, 165(2), 479–494. <https://doi.org/10.1016/j.ejor.2004.04.017>
- [13]. Gen, M., & Lin, L. (2014). Multiobjective evolutionary algorithm for manufacturing scheduling problems: State-of-the-art survey. *Journal of Intelligent Manufacturing*, 25, 849–866. <https://doi.org/10.1007/s10845-013-0804-4>
- [14]. Gu, J., Gu, M., Cao, C., & Gu, X. (2010). A novel competitive co-evolutionary quantum genetic algorithm for stochastic job shop scheduling problem. *Computers & Operations Research*, 37(5), 927–937. <https://doi.org/10.1016/j.cor.2009.07.002>
- [15]. Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science*, 43(1), 115–135. <https://doi.org/10.1007/s11747-014-0403-8>
- [16]. Holden, R. J., & Karsh, B. T. (2010). The Technology Acceptance Model: Its past and its future in health care. *Journal of Biomedical Informatics*, 43(1), 159–172. <https://doi.org/10.1016/j.jbi.2009.07.002>

- [17]. Hutter, F., Hoos, H. H., & Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In C. A. Coello Coello (Ed.), *Learning and Intelligent Optimization (LION 5)* (pp. 507–523). Springer. https://doi.org/10.1007/978-3-642-25566-3_40
- [18]. Ishibuchi, H., Setoguchi, Y., Masuda, H., & Nojima, Y. (2015). Performance comparison of evolutionary multi-objective and many-objective algorithms: part I. *IEEE Transactions on Evolutionary Computation*, 19(5), 692–706. <https://doi.org/10.1109/tevc.2015.2477296>
- [19]. Jin, Y. (2005). A comprehensive survey of fitness approximation in evolutionary computation. *Soft Computing*, 9(1), 3–12. <https://doi.org/10.1007/s00500-003-0328-5>
- [20]. Jin, Y. (2011). Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation*, 1(2), 61–70. <https://doi.org/10.1016/j.swevo.2011.05.001>
- [21]. King, W. R., & He, J. (2006). A meta-analysis of the technology adoption acceptance model. *Information Systems Research*, 17(4), 740–755. <https://doi.org/10.1287/isre.1060.0083>
- [22]. Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research*, 56(1–2), 508–517. <https://doi.org/10.1080/00207543.2017.1351644>
- [23]. Lahoz-Beltra, R. (2016). Quantum genetic algorithms for computer scientists. *Computers*, 5(4), 24. <https://doi.org/10.3390/computers5040024>
- [24]. Lee, C.-H., Zhu, C., & Lee, C.-Y. (2020). Production scheduling research: Recent problems and trends. *Operations Management Research*, 13, 221–231. <https://doi.org/10.1007/s12063-020-00156-x>
- [25]. Loukil, F., Teghem, J., & Tuytens, D. (2005). Solving multi-objective production scheduling problems using metaheuristics. *European Journal of Operational Research*, 161(1), 42–61. <https://doi.org/10.1016/j.ejor.2003.08.022>
- [26]. Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1–10. <https://doi.org/10.1016/j.jii.2017.04.005>
- [27]. Md Arif Uz, Z., & Elmoon, A. (2023). Adaptive Learning Systems For English Literature Classrooms: A Review Of AI-Integrated Education Platforms. *International Journal of Scientific Interdisciplinary Research*, 4(3), 56-86. <https://doi.org/10.63125/a30ehr12>
- [28]. Md Sanjid, K., & Md. Tahmid Farabe, S. (2021). Federated Learning Architectures For Predictive Quality Control In Distributed Manufacturing Systems. *American Journal of Interdisciplinary Studies*, 2(02), 01-31. <https://doi.org/10.63125/222nwg58>
- [29]. Md Takbir Hossen, S., & Md Atiqur, R. (2022). Advancements In 3D Printing Techniques For Polymer Fiber-Reinforced Textile Composites: A Systematic Literature Review. *American Journal of Interdisciplinary Studies*, 3(04), 32-60. <https://doi.org/10.63125/s4r5m391>
- [30]. Md. Omar, F., & Md Harun-Or-Rashid, M. (2021). Post-GDPR Digital Compliance in Multinational Organizations: Bridging Legal Obligations With Cybersecurity Governance. *American Journal of Scholarly Research and Innovation*, 1(01), 27-60. <https://doi.org/10.63125/4qpdf28>
- [31]. Md. Rabiul, K., & Sai Praveen, K. (2022). The Influence of Statistical Models For Fraud Detection In Procurement And International Trade Systems. *American Journal of Interdisciplinary Studies*, 3(04), 203-234. <https://doi.org/10.63125/9htnv106>
- [32]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 65-90. <https://doi.org/10.63125/sw7jzx60>
- [33]. Minella, G., Ruiz, R., & Ciavotta, M. (2018). A review and evaluation of the permutation flowshop scheduling problem. *Annals of Operations Research*, 259, 295–356. <https://doi.org/10.1007/s10479-017-2673-1>
- [34]. Mouzon, G., Yildirim, M. B., & Twomey, J. (2007). Minimizing energy consumption and maximizing profit on manufacturing systems. *International Journal of Production Research*, 45(20), 4247–4271. <https://doi.org/10.1080/00207540701450013>
- [35]. Mubashir, I. (2021). Smart Corridor Simulation for Pedestrian Safety: : Insights From Vissim-Based Urban Traffic Models. *International Journal of Business and Economics Insights*, 1(2), 33-69. <https://doi.org/10.63125/b1bk0w03>
- [36]. Omar Muhammad, F., & Md. Redwanul, I. (2023). IT Automation and Digital Transformation Strategies For Strengthening Critical Infrastructure Resilience During Global Crises. *American Journal of Interdisciplinary Studies*, 4(04), 145-176. <https://doi.org/10.63125/vrsjp515>
- [37]. Özgüven, C., Özbakır, L., & Yavuz, Y. (2010). Scheduling in flexible job shop: A literature review. *Computers & Industrial Engineering*, 58(3), 395–408. <https://doi.org/10.1016/j.cie.2009.10.011>
- [38]. Pan, Q.-K., & Ruiz, R. (2012a). A comprehensive review and evaluation of permutation flowshop heuristics to minimize flowtime. *Computers & Operations Research*, 40(1), 117–128. <https://doi.org/10.1016/j.cor.2012.05.018>

- [39]. Pan, Q.-K., & Ruiz, R. (2012b). Local search methods for the flowshop scheduling problem with flowtime minimization. *European Journal of Operational Research*, 222(1), 31–43. <https://doi.org/10.1016/j.ejor.2012.04.034>
- [40]. Petter, S., DeLone, W., & McLean, E. (2008). Measuring information systems success: Models, dimensions, measures, and interrelationships. *European Journal of Information Systems*, 17(3), 236–263. <https://doi.org/10.1057/ejis.2008.15>
- [41]. Pinedo, M. L. (2012). *Scheduling: Theory, algorithms, and systems* (4th ed.). Springer. <https://doi.org/10.1007/978-1-4614-2361-4>
- [42]. Razia, S. (2022). A Review Of Data-Driven Communication In Economic Recovery: Implications Of ICT-Enabled Strategies For Human Resource Engagement. *International Journal of Business and Economics Insights*, 2(1), 01-34. <https://doi.org/10.63125/7tkv8v34>
- [43]. Razia, S. (2023). AI-Powered BI Dashboards In Operations: A Comparative Analysis For Real-Time Decision Support. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 3(1), 62–93. <https://doi.org/10.63125/wqd2f159>
- [44]. Reduanul, H. (2023). Digital Equity and Nonprofit Marketing Strategy: Bridging The Technology Gap Through Ai-Powered Solutions For Underserved Community Organizations. *American Journal of Interdisciplinary Studies*, 4(04), 117-144. <https://doi.org/10.63125/zrsv2r56>
- [45]. Rony, M. A. (2021). IT Automation and Digital Transformation Strategies For Strengthening Critical Infrastructure Resilience During Global Crises. *International Journal of Business and Economics Insights*, 1(2), 01-32. <https://doi.org/10.63125/8tzzab90>
- [46]. Ruiz, R., & Maroto, C. (2005). A comprehensive review and evaluation of permutation flowshop heuristics. *European Journal of Operational Research*, 165(2), 479–494. <https://doi.org/10.1016/j.ejor.2004.04.017>
- [47]. Ruiz, R., & Vázquez-Rodríguez, J. A. (2010). The hybrid flow shop scheduling problem. *European Journal of Operational Research*, 205(1), 1–18. <https://doi.org/10.1016/j.ejor.2009.09.024>
- [48]. Saad, H. M. H., Chakraborty, R. K., Elsayed, S., & Ryan, M. J. (2021). Quantum-inspired genetic algorithm for resource-constrained project-scheduling. *IEEE Access*, 9, 38488–38502. <https://doi.org/10.1109/access.2021.3062790>
- [49]. Sadia, T. (2023). Quantitative Analytical Validation of Herbal Drug Formulations Using UPLC And UV-Visible Spectroscopy: Accuracy, Precision, And Stability Assessment. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 3(1), 01–36. <https://doi.org/10.63125/fxqpds95>
- [50]. Sai Srinivas, M., & Manish, B. (2023). Trustworthy AI: Explainability & Fairness In Large-Scale Decision Systems. *Review of Applied Science and Technology*, 2(04), 54-93. <https://doi.org/10.63125/3w9v5e52>
- [51]. Schepers, J., & Wetzels, M. (2007). A meta-analysis of the technology acceptance model: Investigating subjective norm and moderation effects. *Information & Management*, 44(1), 90–103. <https://doi.org/10.1016/j.im.2007.03.001>
- [52]. Schober, P., Boer, C., & Schwarte, L. A. (2018). Correlation coefficients: Appropriate use and interpretation. *Anesthesia & Analgesia*, 126(5), 1763–1768. <https://doi.org/10.1213/ane.0000000000002864>
- [53]. Shen, J., Yao, X., Liu, Y., & Lai, K. K. (2015). Mathematical modeling and multi-objective evolutionary algorithms for dynamic flexible job shop scheduling. *Information Sciences*, 298, 198–224. <https://doi.org/10.1016/j.ins.2014.09.040>
- [54]. Singh, M. R., & Mahapatra, S. S. (2016). A quantum-behaved particle swarm optimization for flexible job shop scheduling. *Computers & Industrial Engineering*, 93, 36–44. <https://doi.org/10.1016/j.cie.2015.12.004>
- [55]. Sullivan, G. M., & Artino Jr., A. R. (2013). Analyzing and interpreting data from Likert-type scales. *Journal of Graduate Medical Education*, 5(4), 541–542. <https://doi.org/10.4300/jgme-5-4-18>
- [56]. Sun, N., Lv, N., & Yang, J. (2015). Bi-objective optimal production scheduling in discrete manufacturing workshop under time-of-use electricity pricing. *RAIRO—Operations Research*, 49(3), 563–585. <https://doi.org/10.1051/ro/2015063>
- [57]. Syed Zaki, U. (2021). Modeling Geotechnical Soil Loss and Erosion Dynamics For Climate-Resilient Coastal Adaptation. *American Journal of Interdisciplinary Studies*, 2(04), 01-38. <https://doi.org/10.63125/vsfjtt77>
- [58]. Tsai, J.-T., Chou, J.-H., & Ho, W.-H. (2012). Improved quantum-inspired evolutionary algorithm for engineering design optimization. *Mathematical Problems in Engineering*, 836597. <https://doi.org/10.1155/2012/836597>
- [59]. Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *Decision Support Systems*, 53(3), 479–489. <https://doi.org/10.1016/j.dss.2012.05.007>

- [60]. Wirojanagud, P., Gel, E. S., Fowler, J. W., & Cardy, R. (2007). Modeling inherent worker differences for workforce planning. *International Journal of Production Research*, 45(3), 525–553. <https://doi.org/10.1080/00207540600792242>
- [61]. Wixom, B. H., & Todd, P. A. (2005). A theoretical integration of user satisfaction and technology acceptance. *Information Systems Research*, 16(1), 85–102. <https://doi.org/10.1287/isre.1050.0042>
- [62]. Xie, Y., Yang, J., Fu, M., Chen, W., & Liew, A. (2018). An elitist quantum-inspired evolutionary algorithm for the flexible job shop scheduling problem. *Journal of Intelligent Manufacturing*, 29(2), 403–420. <https://doi.org/10.1007/s10845-016-1227-2>
- [63]. Zayadul, H. (2023). Development Of An AI-Integrated Predictive Modeling Framework For Performance Optimization Of Perovskite And Tandem Solar Photovoltaic Systems. *International Journal of Business and Economics Insights*, 3(4), 01–25. <https://doi.org/10.63125/8xm7wa53>
- [64]. Zhang, M., & Chiong, R. (2016). Solving the energy-efficient job-shop scheduling problem: An improved multi-objective evolutionary algorithm with ϵ -constraint strategy. *Journal of Cleaner Production*, 112, 3361–3375. <https://doi.org/10.1016/j.jclepro.2015.09.097>
- [65]. Zhang, Q. (2011). Quantum-inspired evolutionary algorithms: A survey and empirical study. *Journal of Heuristics*, 17, 303–351. <https://doi.org/10.1007/s10732-010-9136-0>
- [66]. Zhang, Q., & Li, H. (2007). MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6), 712–731. <https://doi.org/10.1109/tevc.2007.892759>
- [67]. Zhao, J.-X., Qian, B., Hu, R., Zhang, C.-S., & Li, Z.-H. (2016). An improved quantum-inspired evolution algorithm for no-wait flow shop scheduling problem to minimize makespan. In D. S. Huang (Ed.), *Intelligent computing theories and application* (pp. 536–547). Springer. https://doi.org/10.1007/978-3-319-42291-6_54
- [68]. Zhao, W., Liu, L., & Ma, Z. (2016). Quantum-inspired evolutionary algorithm for no-wait flow shop scheduling Neural Information Processing (LNCS 9713, pp. 463–472),