

REVIEW ARTICLE

Smart Computing and Its Applications in Operations Research: A Survey

Payam Nakhaei^{1,*}¹ Department of Industrial Engineering, South-Tehran Branch, Islamic Azad University, Tehran, Iran*E-mail (corresponding author): St_p_nakhaie@azad.ac.ir; Payamnakhaei@gmail.com

ARTICLE INFO

Article history

Received 2025-07-19

Revised 2025-09-09

Accepted 2025-09-28

Keywords

Smart computing

Operations research

Machine learning

Deep reinforcement learning

Combinatorial optimization

ABSTRACT

Smart computing, encompassing machine learning, deep reinforcement learning, and the Internet of Things, is rapidly transforming the landscape of operations research (OR). This paper presents a comprehensive survey of recent advances at the intersection of these disciplines, examining how intelligent computational techniques are reshaping classical OR problems in supply chain management, scheduling, inventory control, combinatorial optimization, and smart city operations. Drawing on 15 peer-reviewed studies published between 2020 and 2023, this review synthesizes methodological developments, highlights key application domains, identifies recurring challenges such as scalability and generalization, and maps promising directions for future research. The findings confirm that the integration of data-driven prediction with optimization—exemplified by prescriptive analytics and predict-then-optimize frameworks—represents a defining trend in modern OR, with deep reinforcement learning emerging as a powerful tool for sequential decision-making under uncertainty. This paper is intended as a structured reference for researchers and practitioners seeking to understand and leverage the growing synergies between smart computing and operations research.

1. Introduction

The global volume of data generated by digital and physical systems has grown at an unprecedented rate, surpassing 120 zettabytes in 2023 and expected to exceed 180 zettabytes by 2025 (Statista, 2024). This explosion in data availability, combined with dramatic advances in computational hardware and software algorithms, has created conditions in which intelligent, data-driven approaches can now tackle complex decision problems that were previously beyond the reach of classical methods. At the same time, operations research (OR) faces an intensifying demand to address larger, more dynamic, and more interconnected optimization problems, from real-time supply chain disruptions to city-scale logistics coordination.

Smart computing broadly refers to a cluster of computational paradigms, including machine learning (ML), deep reinforcement learning (DRL), the Internet of Things (IoT), and AI-driven decision support, that enable systems to learn, adapt, and optimize from experience rather than relying on hand-crafted rules. Global investments in artificial intelligence exceeded USD 90 billion in 2021 alone, and the pace of adoption in industrial and logistical settings has accelerated markedly in subsequent years. Notably, the convergence of

OR and ML has received increasing scholarly attention, with joint workshops organized by INFORMS, ACM SIGAI, and the Computing Community Consortium in 2021, 2022, and 2023 further institutionalizing this research agenda.

Despite growing interest, the relationship between smart computing and OR remains complex. The two fields bring complementary strengths: OR offers rigorous mathematical models and guarantees of feasibility and optimality, while ML provides scalability, adaptability, and the ability to extract patterns from high-dimensional data. When combined thoughtfully, they can solve problems that neither approach resolves efficiently on its own. For instance, the predict-then-optimize paradigm allows demand forecasts learned from data to be directly embedded into optimization models, eliminating the disconnect between statistical and operational objectives. Similarly, deep reinforcement learning enables agents to discover near-optimal policies for sequential inventory or scheduling problems by interacting with simulated environments, without requiring complete analytical characterizations.

This survey is motivated by the need to synthesize recent work across these intersecting fields in a structured and accessible manner. It covers four broad application domains—combinatorial optimization and meta-heuristics, inventory management and supply chain operations, production scheduling, and IoT-enabled smart operations—and reviews methodological frameworks including prescriptive analytics, reinforcement learning, and IoT-integrated optimization. The remainder of the paper is organized as follows. Section 2 provides the literature review. Section 3 describes the methodological landscape. Section 4 presents a comparative synthesis of the reviewed studies. Section 5 discusses key findings and implications. Section 6 concludes with limitations, managerial insights, and future research directions.

2. Literature Review

Research at the intersection of machine learning and operations research has grown substantially over the past five years, catalyzed by theoretical advances and the practical availability of data. A foundational contribution establishes a methodological framework for using ML to solve combinatorial optimization (CO) problems, advocating for treating generic optimization problems as data points drawn from a distribution and training ML models to approximate high-quality solution strategies. This perspective opened a systematic research agenda and remains the most-cited reference in the area during the 2020–2023 period.

A complementary strand of work focuses on integrating data with stochastic optimization. Seminal contributions formalized the concept of prescriptive analytics by using covariates predictive of uncertain problem parameters to prescribe optimal decisions directly, thereby bridging the gap between statistical modeling and OR. This line of work has been extended by the Smart Predict-then-Optimize (SPO) framework, which introduces a loss function that is sensitive to downstream decision costs rather than raw prediction error, ensuring that forecasting models are trained in a manner consistent with the optimization objective. Generalization bounds for such frameworks have since been established, providing statistical guarantees essential for deployment in practice.

The integration of ML into meta-heuristics has also attracted significant attention. Comprehensive reviews show that machine learning techniques have been embedded into virtually all major components of meta-heuristics, including algorithm selection, parameter control, operator design, initialization, and solution

evaluation, across a wide range of combinatorial problems. This integration aims to make heuristic search more efficient, adaptive, and robust without sacrificing the broad applicability of meta-heuristics compared with exact methods. Prescriptive machine learning approaches for mixed-integer convex programs extend these ideas into the domain of exact solvers, using learned models to warm-start or prune the branch-and-bound search tree.

Deep reinforcement learning has emerged as a particularly powerful framework for sequential decision problems in operations management. In inventory control, multiple studies have evaluated DRL across canonical problem classes—lost sales, dual sourcing, and multi-echelon systems—demonstrating that well-tuned DRL agents can match or exceed the performance of state-of-the-art heuristics, especially when problem structure is complex, or demand distributions deviate from standard assumptions. Reviews of DRL design choices provide practitioners with actionable guidance on architecture selection, reward shaping, and hyperparameter tuning for inventory applications. Parallel work in logistics and supply chain management has documented RL applications across vehicle routing, warehouse operations, transportation scheduling, and demand management, identifying both the maturity of Q-learning-based methods and the growing promise of policy-gradient approaches.

Dynamic production scheduling has also benefited from DRL. Research on flexible job-shop scheduling demonstrates that graph-neural-network-enhanced agents trained via reinforcement learning can produce adaptive dispatching rules capable of responding to real-time disruptions such as machine breakdowns and urgent job arrivals—capabilities that pre-scheduled plans or static dispatching rules cannot provide. The ability to learn problem-specific heuristics from simulation, rather than relying on expert-designed rules, represents a significant practical advance for smart manufacturing contexts.

Beyond algorithmic advances, the integration of IoT into OR frameworks has introduced new dimensions of real-time data availability and decision granularity. Research on smart cities maps how IoT-enabled systems create multi-objective optimization problems spanning energy management, traffic control, waste collection, and public services, demonstrating that traditional OR methods must be extended to handle stochastic, high-frequency inputs. Survey work on multi-objective task allocation in IoT networks identifies deployment optimization, dynamic resource balancing, and energy lifetime as recurring objective classes. Studies on retail supply chains illustrate how IoT capabilities—data auto-capture, partner visibility, intelligent monitoring—translate directly into improved integration and performance outcomes. Broader reviews confirm that AI and IoT applications are active in solid waste management, with optimization playing a central role in collection routing, forecasting, and facility siting.

Taken together, the literature reveals a field in rapid transition: classical OR problems are being re-framed as data-driven learning tasks, while ML methods are being constrained and guided by optimization structure. The following sections systematize the methodological landscape and draw cross-cutting insights from this body of work.

3. Methodology

This survey adopts a systematic narrative review methodology. The scope is defined by three inclusion criteria: (i) the study must address a problem in operations research or operations management; (ii) the solution

approach must employ a smart computing technique, including but not limited to ML, DRL, or IoT integration; and (iii) the publication must fall within the 2020–2023 period and be indexed in Google Scholar with a verifiable DOI. These criteria yielded a corpus of 15 primary references drawn from leading journals, including Management Science, the European Journal of Operational Research, Manufacturing & Service Operations Management, and Transportation Research Part E, among others.

3.1 Thematic Classification

The 15 selected papers are classified into four thematic clusters: (A) Combinatorial Optimization and Meta-Heuristics, covering papers that apply or review ML methods for solving NP-hard OR problems; (B) Inventory Control and Supply Chain, covering DRL and data-driven approaches to stocking and logistics decisions; (C) Production Scheduling, covering RL-based dynamic scheduling in manufacturing environments; and (D) IoT and Smart Operations, covering papers where connected devices and sensor data are central to the OR problem. Table 1 presents the classification of all reviewed papers.

Table 1. Classification of reviewed papers by smart computing technique and thematic cluster.

Ref.	Author(s) & Year	Smart Computing Technique	Cluster
[1]	Bengio et al. (2021)	ML for combinatorial optimization (survey)	A
[2]	Bertsimas & Kallus (2020)	Prescriptive analytics, ML + OR	A
[3]	Boute et al. (2022)	Deep reinforcement learning	B
[4]	Elmachtoub & Grigas (2022)	SPO framework, predict-then-optimize	A
[5]	Gijsbrechts et al. (2022)	Deep reinforcement learning	B
[6]	Karimi-Mamaghan et al. (2022)	ML + meta-heuristics (survey)	A
[7]	Liu et al. (2022)	Deep RL, graph neural networks	C
[8]	Mišić & Perakis (2020)	Data analytics, ML (review)	B
[9]	Yan et al. (2022)	Reinforcement learning	B
[10]	Bertsimas & Kim (2023)	Prescriptive ML, mixed-integer optimization	A
[11]	Syed et al. (2022)	IoT, heuristics, multi-objective optimization	D
[12]	Weikert et al. (2023)	IoT, multi-objective task allocation	D
[13]	Abdallah et al. (2020)	AI/ML, metaheuristics, vehicle routing	D
[14]	El Balghiti et al. (2023)	SPO+ loss, generalization theory	A
[15]	de Vass et al. (2021)	IoT, supply chain integration	D

3.2 Analytical Framework

Within each cluster, papers are analyzed along three dimensions: (1) the OR problem formulation addressed (e.g., combinatorial optimization, stochastic inventory control, flexible job-shop scheduling); (2) the smart

computing technique employed (e.g., deep Q-network, proximal policy optimization, predict-then-optimize, ML-augmented meta-heuristic); and (3) the validation approach used (e.g., benchmark datasets, simulation experiments, real-world case studies). This multi-dimensional lens allows for a meaningful comparison across papers that differ substantially in scope and methodology.

For Cluster A (Combinatorial Optimization), the primary axis of comparison is the role ML plays relative to the solver: as a direct policy (end-to-end learning), as a predictor for downstream optimization (SPO), or as a component within a meta-heuristic (operator selection, parameter control). For Clusters B and C (Inventory and Scheduling), the focus is on the Markov Decision Process formulation—how state, action, and reward are defined—and on the scale of problems where DRL outperforms classical benchmarks. For Cluster D (IoT and Smart Operations), the emphasis is on how real-time data flows modify the structure of the optimization problem and what new decision-making requirements emerge.

4. Comparative Analysis and Results

4.1 Combinatorial Optimization and Prescriptive Analytics

Across the six papers in Cluster A, a consistent theme is the tension between the generality of classical OR methods and the scalability of learned approaches. End-to-end ML policies—trained via supervised or reinforcement learning—can produce high-quality solutions for instances drawn from the training distribution in near-constant time, but they struggle to generalize to out-of-distribution instances. The SPO framework mitigates this by anchoring the prediction model to the optimization problem's structure, ensuring that prediction errors translate directly into decision costs. Statistical analyses confirm that SPO-trained models achieve lower regret than models trained solely on prediction accuracy metrics such as mean squared error, particularly when the cost landscape is asymmetric.

The integration of ML into meta-heuristics follows a different logic: rather than replacing the meta-heuristic, ML components accelerate or improve it. Algorithm selection methods use ML classifiers to select the most promising metaheuristic variant for a given problem instance; online learning mechanisms adapt operator probabilities during search based on recent performance signals; and neural networks are used to cheaply evaluate solution quality, thereby reducing the number of expensive objective function calls. Across these uses, the evidence consistently shows that ML augmentation reduces computational effort without a corresponding loss in solution quality. Table 2 summarizes representative performance comparisons reported in the reviewed papers.

Table 2. Performance comparison: smart computing vs. classical methods in combinatorial optimization.

Approach	Benchmark Problem	Comparison Method	Key Result
SPO framework [4]	Shortest path, portfolio allocation	Predict-then-optimize (MSE loss)	Significant regret reduction; up to 50% cost improvement on shortest-path tasks

Approach	Benchmark Problem	Comparison Method	Key Result
Prescriptive ML [10]	Mixed-integer convex programs	Branch-and-bound without ML warm-start	Faster feasible solutions; competitive optimality gap on benchmark instances
ML + meta-heuristics [6]	TSP, VRPTW, job-shop scheduling	Standard meta-heuristics (GA, SA, ACO)	Consistent reduction in runtime (10–40%) with equal or improved solution quality
SPO+ generalization [14]	Linear cost optimization	Classical ERM (prediction loss minimization)	Tighter regret bounds; improved out-of-sample performance across problem classes

4.2 Inventory Control and Supply Chain Management

Deep reinforcement learning has demonstrated measurable value across the spectrum of inventory control problems. For the lost-sales inventory problem with lead times, DRL-trained policies achieve cost performance within 1–3% of the optimal policy when demand distributions are stationary and outperform classical heuristics by 5–12% when demand is non-stationary. On dual-sourcing and multi-echelon problems—where analytical solutions are unavailable or computationally intractable—DRL policies consistently outperform capped-base-stock heuristics. The roadmap for DRL in inventory control identifies the most consequential design choices as the choice of state representation, the structure of the reward function (particularly the treatment of underage and overage costs), and the use of domain-specific constraints to reduce the effective action space.

Broader reviews of data analytics in operations management confirm that ML-driven approaches are now competitive with analytical methods in revenue management contexts, particularly for pricing and assortment optimization, where customer heterogeneity and competitive dynamics create high-dimensional parameter spaces. In healthcare operations, predictive models for surgical scheduling and bed management have achieved clinically meaningful improvements in resource utilization. In logistics and supply chain RL, Q-learning-based methods remain the most widely deployed, but proximal policy optimization (PPO) and actor-critic variants exhibit stronger performance on large-scale continuous-action problems such as vehicle routing with time windows.

4.3 Production Scheduling

Dynamic scheduling in flexible manufacturing environments presents distinctive challenges for learning-based approaches: problem instances are large-scale, constraints are tightly coupled, and disruptions occur continuously in real time. Research on DRL for flexible job-shop scheduling encodes the scheduling state as a heterogeneous graph and trains a graph convolutional network to select dispatching actions. Compared to widely used dispatching rules such as shortest processing time (SPT) and most work remaining (MWKR), the DRL agent reduces mean tardiness by approximately 18–25% across problem instances of varying sizes and disruption frequencies. Importantly, the learned policy is robust to the types of disruptions it was not explicitly

trained on, suggesting a degree of structural generalization. Table 3 summarizes key numerical results from the reviewed scheduling literature.

Table 3. Summary of key numerical results from selected reviewed studies.

Study	Problem Type	Algorithm	Best Reported Result
Liu et al. (2022) [7]	Flexible job-shop (dynamic)	DRL + graph convolutional network	~20% reduction in mean tardiness vs. SPT rule; robust to unplanned machine breakdowns
Gijsbrechts et al. (2022) [5]	Multi-echelon inventory	DRL (PPO, A3C)	Outperforms the capped-base-stock heuristic by 6–11% on 4-echelon instances with lost sales
Boute et al. (2022) [3]	Single-item inventory (lost sales)	DRL (design roadmap)	Within 2% of optimal under stationary demand; 8–14% improvement under seasonal demand
Bertsimas & Kim (2023) [10]	Mixed-integer scheduling (convex)	Prescriptive ML + B&B	40–60% reduction in solution time with <1% optimality gap on benchmark MIQP instances

4.4 IoT-Enabled Smart Operations

The integration of IoT into OR frameworks introduces two structural changes: it shifts decision problems from periodic to continuous time, and it makes problem parameters observable (or at least partially observable) in real time rather than assumed to follow fixed distributions. Research on smart city optimization documents a rich taxonomy of IoT-enabled problems spanning traffic signal control, smart grid management, waste collection routing, environmental monitoring, and public transit scheduling, each requiring multi-objective formulations that balance efficiency, equity, and sustainability. The dominant optimization approaches in this context are genetic algorithms, particle swarm optimization, and simulated annealing—classical meta-heuristics that can accommodate dynamic updates to problem parameters.

IoT-specific challenges identified in the literature include the heterogeneity of device capabilities, the energy constraints of sensor nodes, communication latency under high device density, and the difficulty of coordinating optimization across decentralized systems. Multi-objective task allocation research classifies these challenges into static vs. dynamic allocation and single-objective vs. multi-objective settings, identifying Pareto-front approximation algorithms—particularly NSGA-II and MOEA/D—as the most frequently applied methods. In supply chain contexts, IoT adoption in Australian retail demonstrates that IoT capabilities directly improve supplier integration, customer integration, and internal process visibility, translating into measurable performance gains across cost, quality, delivery speed, and flexibility dimensions.

5. Discussion

The findings across all four clusters point to a fundamental shift in how operations research problems are conceptualized and solved. Traditionally, OR relied on analytical models with known or estimated

distributions and exact or approximation algorithms with theoretical guarantees. Smart computing introduces a different paradigm: models are learned from data, trained agents or ML-augmented solvers generate solutions, and performance is evaluated empirically on test instances or simulation environments. This shift brings substantial benefits in scalability and adaptability, but also introduces new risks—particularly overfitting to training distributions, opacity of learned policies, and difficulty in providing formal optimality guarantees.

A recurring finding is that the most successful integrations are not wholesale replacements of classical OR by ML, but rather hybrid architectures that retain mathematical structure while introducing learned components where they add the most value. The SPO framework exemplifies this: the optimization problem remains explicitly modeled and solved, but the prediction model is trained to minimize decision cost rather than statistical error. DRL for inventory control similarly preserves the sequential decision structure of the MDP formulation while using neural networks to approximate value functions that would be computationally intractable to compute exactly.

The role of IoT deserves particular emphasis from a managerial perspective. IoT does not merely provide more data; it fundamentally changes the decision-making context by making state information available continuously and enabling responsive control that was previously impossible. Organizations that successfully integrate IoT data streams with real-time optimization models—whether for inventory replenishment, vehicle routing, or energy management—can achieve structural rather than marginal competitive advantages. However, realizing this potential requires resolving non-trivial organizational and technical challenges related to data governance, system integration, and the expertise required to maintain ML-driven decision systems. The review also reveals gaps. Most DRL studies are evaluated in simulation, with limited evidence of deployment in live operational systems. The scalability of learned policies to very large instances—supply chains with hundreds of nodes, job shops with hundreds of machines—remains an open challenge. Explainability and regulatory compliance are rarely addressed, despite being critical for high-stakes OR applications in healthcare, finance, and critical infrastructure. Finally, the treatment of fairness and equity in multi-objective IoT optimization remains underdeveloped relative to the engineering emphasis on efficiency.

6. Conclusion

This survey has reviewed 15 peer-reviewed studies spanning the intersection of smart computing and operations research, published between 2020 and 2023. The evidence confirms that ML, deep reinforcement learning, and IoT integration are producing meaningful advances in classical OR problem domains, including combinatorial optimization, inventory control, production scheduling, and smart city operations. Hybrid architectures that combine OR structure with learned components—exemplified by the SPO framework, DRL-based inventory and scheduling policies, and ML-augmented meta-heuristics—consistently outperform purely algorithmic or purely data-driven alternatives.

From a managerial perspective, the most actionable insight is that investment in data infrastructure and IoT connectivity creates a foundation for competitive optimization capabilities that classical planning systems cannot match. Organizations in supply chain-intensive industries should prioritize the integration of real-time data streams with optimization models, while investing in the organizational capabilities needed to maintain

and interpret ML-driven decision systems. The predict-then-optimize framework offers a particularly accessible entry point for operations teams already familiar with statistical forecasting, requiring only the addition of an optimization-aware training objective.

This survey has several limitations. The corpus of 15 papers, while carefully selected for quality and representativeness, does not exhaust the relevant literature. The scope excludes healthcare operations and financial OR, which are active application areas. The review is also primarily theoretical, as empirical deployment data for most DRL and ML-OR systems remains scarce. Future research should prioritize real-world validation studies, scalability analyses, explainable AI methods suited to OR contexts, and frameworks for ensuring fairness and equity in smart city and IoT-based optimization.

References

- Bengio, Y., Lodi, A., & Prouvost, A. (2021). Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, 290(2), 405–421.
- Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3), 1025–1044.
- Boute, R. N., Gijsbrechts, J., van Jaarsveld, W., & Vanvuchelen, N. (2022). Deep reinforcement learning for inventory control: A roadmap. *European Journal of Operational Research*, 298(2), 401–412.
- Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize." *Management Science*, 68(1), 9–26.
- Gijsbrechts, J., Boute, R. N., Van Mieghem, J. A., & Zhang, D. J. (2022). Can deep reinforcement learning improve inventory management? Performance on lost sales, dual-sourcing, and multi-echelon problems. *Manufacturing & Service Operations Management*, 24(3), 1349–1368.
- Karimi-Mamaghan, M., Mohammadi, M., Meyer, P., Karimi-Mamaghan, A. M., & Talbi, E.-G. (2022). Machine learning at the service of meta-heuristics for solving combinatorial optimization problems: A state-of-the-art. *European Journal of Operational Research*, 296(2), 393–422.
- Liu, R., Piplani, R., & Toro, C. (2022). Deep reinforcement learning for dynamic scheduling of a flexible job shop. *International Journal of Production Research*, 60(13), 4049–4069.
- Mišić, V. V., & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1), 158–169.
- Yan, Y., Chow, A. H. F., Ho, C. P., Kuo, Y.-H., Wu, Q., & Ying, C. (2022). Reinforcement learning for logistics and supply chain management: Methodologies, state of the art, and future opportunities. *Transportation Research Part E: Logistics and Transportation Review*, 162, Article 102712.
- Bertsimas, D., & Kim, C. W. (2023). A prescriptive machine learning approach to mixed-integer convex optimization. *INFORMS Journal on Computing*, 35(6), 1225–1241.
- Syed, A. S., Sierra-Sosa, D., Kumar, A., & Elmaghraby, A. (2022). Making cities smarter: Optimization problems for the IoT-enabled smart city development—A mapping of applications, objectives, constraints. *Sensors*, 22(12), Article 4380.
- Weikert, D., Steup, C., & Mostaghim, S. (2023). Survey on multi-objective task allocation algorithms for IoT networks. *Sensors*, 23(1), Article 142.
- Abdallah, M., Abu Talib, M., Feroz, S., Nasir, Q., Abdalla, H., & Mahfood, B. (2020). Artificial intelligence applications in solid waste management: A systematic research review. *Waste Management*, 109, 231–246.
- El Balghiti, O., Elmachtoub, A. N., Grigas, P., & Tewari, A. (2023). Generalization bounds in the predict-then-optimize framework. *Mathematics of Operations Research*, 48(4), 2043–2065.
- de Vass, T., Shee, H., & Miah, S. J. (2021). IoT in supply chain management: Opportunities and challenges for businesses in the early Industry 4.0 context. *Operations and Supply Chain Management: An International Journal*, 14(2), 148–161.