Choi, Tsan-Ming, Basu, Preetam, Tiwari, Sunil, Biswas, Indranil and Kunc, Martin (2026) *Exploring complementarities and differences between Operational Research and Generative AI to transform operations.* Journal of the Operational Research Society, 77 (7). ISSN 0160-5682.

Exploring Complementarities and Differences between Operational Research and Generative AI to Transform Operations

Authors:

Tsan-Ming Choi, University of Liverpool, UK

Preetam Basu, Kent Business School, UK

Sunil Tiwari, University of Bristol, UK

Indranil Biswas, NEOMA Business School, France

Martin Kunc, University of Southampton, UK

1. Introduction

The rapid ascent of Generative Artificial Intelligence (GenAI) encompassing Large Language Models (LLMs), Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and diffusion models has prompted a fundamental reassessment of how quantitative methods are conceived and applied in Operational Research (OR). This special issue of *Journal of the Operational Research Society* (JORS) was motivated by a simple but important question: what happens when a discipline built on rigorous mathematical modelling, stochastic analysis, and algorithmic optimisation encounters a technology that can generate text, code, data, and decisions at scale?

The call for papers for this JORS special issue attracted considerable scholarly interest. A total of 44 manuscripts were submitted, of which 19 were accepted following a rigorous double-blind peer-review process and 25 were rejected. The accepted papers, covering areas such as mathematical programming, stochastic systems, simulation, supply chains, strategy, and transport logistics, show how widely GenAI applies to operational research.

Early discussions of GenAI in OR were mostly speculative, e.g., could LLMs write optimisation models? The papers in this special issue go further. They offer new analytical frameworks, computational experiments, game-theoretic models, and case studies. Together, they show where GenAI strengthens the OR toolkit, where it creates new decision challenges, and where it produces unintended consequences that deserve careful attention.

This SI shows that the relationship between GenAI and OR is not about one replacing the other or blindly celebrating GenAI. The selected set of papers reveal that it is more like a partnership with some non-trivial challenges. On one side, GenAI greatly expands what OR can do. It helps turn unstructured data (like text) into useful insights, rebuild models from written descriptions, generate solutions for complex problems, improve demand forecasting, and support large groups in making decisions more efficiently. On the other side, GenAI itself becomes something that OR needs to study. Organizations must decide when and how to use it, manage risks (like errors or unreliable outputs), consider its environmental impact, and ensure its results are trustworthy. So, while GenAI strengthens OR, it also introduces new problems that OR must help solve.

We introduce the featured papers around four thematic clusters. The first cluster, *GenAI for Optimisation and Forecasting*, covers papers that embed generative models within OR's traditional problem-solving workflow. The second, *GenAI for Supply Chain Intelligence*, groups papers that use LLMs and RAG architectures to enhance supply chain visibility, simulation, and resilience. The third, *GenAI Adoption: Strategic and Behavioural Dimensions*, focusses on papers that model GenAI adoption as an endogenous decision, examining its equilibria, externalities, and firm-level consequences. The fourth, *Human-AI Collaboration in Operations*, addresses the interface between GenAI systems and human decision-makers in scheduling, logistics, group decisions, and service operations.

We introduce these thematic clusters (called "themes") one by one in Section 2.

2. Themes

Table 1 shows a structured table mapping all 19 accepted papers to their themes, with key methodological highlights.

Table 1. The four themes and features of the accepted papers

Theme	No. of Papers	Dominant GenAI Methods	Dominant OR Methods
GenAI for Optimisation & Forecasting	4	LLMs, GANs, LSTM, DRL	MIP, Metaheuristics, CTMC
GenAI for Supply Chain Intelligence & Simulation	4	RAG, BERT, Prompt Engineering	DES, DEA, Network Science, ABM
GenAI Adoption: Strategic & Behavioural	5	Scenario Generation, Algorithmic AI	Game Theory, Evolutionary Games
Human-AI Collaboration in Operations	6	Chained LLMs, MARL, GenAI Agents	MCC, MARL, CLMM, Platform Design
Total	**19**		

2.1. Theme 1: GenAI for Optimisation and Forecasting

Theme 1 brings together papers that place GenAI at the heart of OR's classical mission: solving optimisation problems better, faster, and with greater accuracy. The papers collectively chart a spectrum from automated problem formulation and algorithm generation to AI-augmented demand forecasting and reward-function design. A summary of them is shown in Table 2 and we also introduce them one by one.

Table 2. Summary of papers of Theme 1

Paper No.	Paper Title	Key GenAI Method	OR Method	Key Contribution
1A	The Use of GenAI in Operations Research: Review and Future Research Agenda	GANs, VAEs, LLMs	Mathematical programming, stochastic systems, simulation	Dual-role framework (augmentation vs. automation); maps GenAI across OR pillars
1B	GenAI–Enhanced Sales Forecasting and Optimization of Multi-period and Multi-product	BERT-AttenSeqNet	Genetic Algorithm–Ant Colony Optimisation (GA-ACO), rolling horizon	End-to-end forecast-then-optimise pipeline for e-commerce; reduces total operational costs

	Inventory Routing Problem			
1C	Optimizing Retailer-Supplier Inventory Systems with Defective Returns	LSTM + GenAI scaling layer	Continuous-Time Markov Chain (CTMC), nonlinear integer programming	GenAI as external calibration layer for stochastic models; counterintuitive reorder-point insights
1D	Generative Heuristics-Driven for Blockchain-Enhanced Reputation Systems in Short Food Supply Chains	Deep Reinforcement Learning Hyperheuristic (DRLH)	Rolling-horizon optimisation	AI-generated heuristic operators + blockchain trust signals reduce contract termination by ~50%

Paper 1A: The use of Generative Artificial Intelligence (GenAI) in operations research: review and future research agenda (Zhou and Sheu 2025)

This systematic review provides the conceptual scaffolding for the entire special issue. The papers span the three core pillars of operational research: mathematical programming and optimisation, stochastic systems, and simulation with strategic analysis. They examine how GANs, VAEs, diffusion models, and LLMs are being used across each area. A central theoretical contribution is the articulation of a *dual role* framework: GenAI either augments human intelligence (as in NL2OPT — Natural Language to Optimisation where LLMs translate problem descriptions into formal models) or automates tasks entirely (as in GAN-generated optimal decisions for portfolio optimisation and logistics scheduling). The review establishes that LLMs and data-generative models occupy distinct OR niches: LLMs dominate model-level and strategy-level tasks, while GANs and VAEs excel at data-level challenges such as synthetic population generation and distributional learning for queues and traffic systems. The authors identify three most pressing unsolved challenges of GenAI usage in OR: (i) verification and trustworthiness, (ii) computational scalability, and (iii) ethical governance.

Paper 1B: Generative AI–enhanced sales forecasting and optimization for the multi-period multi-product inventory routing problem with procurement (Li et al. 2026)

In this paper, the authors demonstrate end-to-end integration of GenAI into a classical OR workflow. The authors deploy a BERT-AttenSeqNet architecture combining transformer-based contextual encoding with attention-weighted sequence modelling to forecast multi-product demand in a self-operated e-commerce setting. These AI-generated demand signals feed directly into a multi-period, multi-product Inventory Routing Problem with Procurement Decisions (MIRP-PD), solved via a hybrid Genetic Algorithm–Ant Colony Optimisation (GA-ACO) framework operating under a rolling horizon. Real-world validation across fifteen operational scenarios confirms that AI-based forecasting substantially reduces total operational costs, including holding, transportation, and stockout components, compared to deterministic and classical time-series benchmarks. The paper exemplifies the *forecast-then-optimise* paradigm: GenAI serves as the informational substrate for classical OR solvers rather than displacing them, achieving measurable efficiency gains precisely through their integration. The rolling-horizon architecture also enables dynamic adaptation to demand shocks, a feature of growing practical importance in volatile e-commerce environments.

Paper 1C: Optimizing Retailer-Supplier Inventory Systems with Defective Returns: A Continuous-Time Markov Chain and Neural Network-Based Approach (Aghsami et al. 2026)

In this paper, the authors address a compound inventory challenge combining defective product returns, price-dependent demand, uncertain delivery times, and customer dissatisfaction-driven returns through a hybrid framework that integrates LSTM-based demand forecasting with a Continuous-Time Markov Chain (CTMC) model. The GenAI module operates as a contextual scaling layer: qualitative signals (promotional intensity, market sentiment, seasonal narratives) are converted into multiplicative adjustment coefficients that modify the LSTM baseline forecast without disturbing the CTMC's analytical structure. This elegant design preserves the tractability and interpretability of Markov analysis while allowing narrative intelligence to enrich underlying parameters. A nonlinear integer programming model then determines optimal reorder points and order quantities under the adjusted stochastic parameters. The principal managerial insight is striking and counterintuitive: a 50% increase in supplier delivery rate reduces optimal reorder levels by approximately 20% while *increasing* expected profit by 2%, because faster replenishment improves inventory turnover more than it exposes the retailer to lost-sales risk. The paper's formal treatment of GenAI as an adjustable external calibration layer offers a design template with broad applicability across OR stochastic models.

Paper 1D: Generative Heuristics-Driven for Blockchain-Enhanced Reputation Systems and Dynamic Optimization in Short Food Supply Chains (Dahbi,et al. 2025)

Short food supply chains (SFSCs) present a distinctive optimisation challenge: perishability, quality heterogeneity, trust deficits, and volatile seasonal demand simultaneously complicate farmer-customer matching decisions. This paper proposes an AI-driven Deep Reinforcement Learning Hyperheuristic (DRLH) that uses GenAI to generate and refine heuristic operators dynamically. The GenAI module predicts the most contextually appropriate insertion and removal operators based on historical optimisation patterns and real-time system state. The framework is integrated with a blockchain-based reputation system that tracks farmer transparency, organic certification, and freshness compliance, incorporating these trust signals directly into the objective function and heuristic selection. Experimental results across three scenarios (High Seasonal Demand, Logistical Disruption, and Negative Reputation Impact) confirm that AI-driven DRLH consistently outperforms classical rolling-horizon and non-GenAI DRLH baselines, achieving average gap values below 2.2% and demonstrating the fastest convergence. The blockchain reputation mechanism reduces contract termination rates by up to half relative to reputation-free settings, underscoring how trust-aware AI optimisation can simultaneously improve operational efficiency and supply chain integrity.

Remarks: Together, these four papers of Theme 1 establish that GenAI's contribution to OR optimisation is most potent when it operates in concert with classical solvers, stochastic models, and heuristic frameworks enriching the informational substrate and adaptive capacity of established methods rather than supplanting their mathematical rigour.

2.2. Theme 2: GenAI for Supply Chain Intelligence and Simulation

The second theme addresses how GenAI, particularly LLM-based Retrieval-Augmented Generation (RAG) pipelines, can transform what OR practitioners know about supply chains and how they build models of them. These papers share a conviction that the most consequential data for supply chain OR is often unstructured: buried in regulatory filings, news archives, earnings calls, and healthcare reports. Table 3 summarizes the featured papers in Theme 2.

Table 3. Summary of papers of Theme 2

Paper No.	Paper Title	Key GenAI Method	OR Method	Key Contribution
2A	Integrating Digitally Enhanced Data Extraction and Simulation Modelling for Supply Chain Resilience	RAG + LLM extraction pipeline	Agent-based simulation	LLM-Integrated Simulation Data Framework; 668 validated insights from 972 documents for India's lithium stockpiling
2B	Supply Chain Mapping through Retrieval-Augmented Generation: Electronics Industry	RAG + GPT-3.5 Turbo (entity/relation extraction)	Network science (centrality, modularity)	4,644-node, 8,341-edge directed supply chain graph for Foxconn, Flex, Jabil; 82.7% precision
2C	Unlocking the Potential of Past Research: Using GenAI to Reconstruct Healthcare Simulation Models	Prompt engineering (Perplexity.AI)	Discrete-Event Simulation (DES)	GenAI reconstructs published DES models; exposes gaps in OR reporting standards
2D	From Voices to Value: Benchmarking Airline Efficiency with GenAI-Assisted Text Analytics and DEA	BERT-based sentiment/topic extraction	Data Envelopment Analysis (DEA)	Unstructured reviews converted to DEA outputs; reveals efficiency gaps invisible to operational data alone

Paper 2A: Integrating Digitally Enhanced Data Extraction and Simulation Modelling for AI-Driven Supply Chain Resilience: An Operational Research Framework for Strategic Stockpiling of Critical Minerals (Nie et al. 2026)

This paper from the University of Cambridge's Industrial Resilience Research Group proposes the LLM-Integrated Simulation Data Framework, a domain-agnostic, four-stage methodology for formalising the extraction of narrative intelligence from unstructured sources and integrating it with structured datasets to enhance agent-based simulation models. Applied to India's strategic lithium carbonate stockpiling challenge, the framework processed 972 documents (news archives, policy reports, corporate disclosures) through a RAG pipeline to generate 668 validated structured insights, achieving a 63% high-confidence extraction rate. These insights were embedded into a multi-echelon agent-based model whose parameters and decision rules were thereby grounded in empirical textual evidence rather than researcher assumptions. The simulation yields a striking policy finding that contradicts conventional supply-chain wisdom: under tightening markets, concentrated sourcing from Chile outperforms geographic diversification, because binding logistics and processing capacity diminish the diversification premium. The paper's methodological contribution, a four-stage, auditable pipeline from narrative to simulation-ready parameter, represents a significant advance for OR in data-sparse, geopolitically sensitive domains and establishes a reproducible template for evidence-based model design.

Paper 2B: Supply Chain Mapping through Retrieval-Augmented Generation: Applications to the Electronics Industry (Jackson et al. 2026)

This paper from MIT's Center for Transportation and Logistics proposes an automated, RAG-based framework for constructing multi-tier supply chain network maps from unstructured public data, specifically SEC 10-K filings and earnings call transcripts. Applied to three of the world's largest electronics contract manufacturers (Foxconn, Flex, and Jabil), the framework constructed a directed supply chain graph encompassing 4,644 nodes and 8,341 edges. Network science analytics including centrality, modularity, average path length, and diameter then reveal structural properties with direct resilience implications: all three focal firms exhibit modularity scores above the industry average (0.49–0.61 vs. 0.38), with Jabil's high modularity (0.61) suggesting a particularly compartmentalised network structure that should limit cascading disruption propagation. Validation against SP Capital IQ commercial data yields a precision of 82.7% overall and over 90% for supplier relationships, demonstrating that the approach is reliable enough for practical risk assessment and compliance applications. The paper makes the case that automated, AI-driven supply chain mapping can serve as a continuous upstream data generator for digital twin systems and strategic resilience planning delivering multi-tier visibility that manual mapping cannot achieve at comparable scale.

Paper 2C: Unlocking the Potential of Past Research: Using Generative AI to Reconstruct Healthcare Simulation Models (Monks et al. 2025)

The vast majority of Discrete-Event Simulation (DES) models built for healthcare OR are never publicly shared; the coded artefact remains with the original authors and is effectively lost to the community over time. This paper investigates whether GenAI can reconstruct such models from the textual descriptions published in academic journals, using SimPy for coding and Streamlit for browser-based interfaces. Working from a structured methodology of prompt engineering with Perplexity.AI, the team successfully generated, tested, and internally reproduced two healthcare DES models of substantially greater complexity than any previously reported AI-generated simulation. Full numerical result replication was achieved for one model but not the other, likely due to missing distributional information in the original publication. The authors conclude that iterative prompt engineering, systematic testing, and domain expertise were all essential, GenAI could not substitute for any of these.

The paper makes an important contribution to OR's open science agenda: it demonstrates that GenAI can dramatically reduce the cost of model reuse while simultaneously exposing the inadequacy of current reporting standards, creating a compelling case for future richer conceptual models.

Paper 2D: From voices to value: Leveraging online reviews to benchmark airline efficiency with GenAI-assisted text analytics and DEA (Golghamat et al. 2025)

This paper bridges qualitative consumer intelligence and quantitative efficiency measurement by developing a GenAI-assisted text analytics pipeline for airline performance benchmarking. Passenger reviews from online platforms are processed through a BERT-based sentiment and topic extraction system, generating structured performance scores across service dimensions that are otherwise invisible to traditional productivity metrics. These scores are then incorporated as outputs in a Data Envelopment Analysis (DEA) efficiency model, enabling a comprehensive airline benchmarking exercise that accounts for experiential quality alongside conventional operational variables such as capacity utilisation and on-time performance. The methodology reveals significant efficiency differentials between carriers that conventional operational data alone cannot explain, with service quality dimensions, particularly cabin comfort and staff responsiveness. emerging as key discriminators. By demonstrating how GenAI transforms unstructured voice-of-customer data into structured DEA inputs, the paper offers a replicable template for GenAI-enhanced performance measurement across any service industry where experiential quality is both commercially critical and difficult to quantify through administrative records.

Remarks: The four featured papers of Theme 2 collectively establish a new paradigm for OR model construction and system assessment: unstructured textual data, when processed through LLM-based pipelines, can supply the contextual intelligence that structured databases alone cannot provide whether for simulation parameterisation, network mapping, model reuse, or service benchmarking.

2.3. Theme 3: GenAI Adoption — Strategic and Behavioural Dimensions

The third theme addresses GenAI not as a methodological tool but as a strategic phenomenon: a technology whose adoption, pricing, and societal consequences are proper objects of OR analysis. These papers deploy game theory, evolutionary dynamics, and analytical economic modelling to understand the incentives, equilibria, and externalities associated with GenAI integration in markets and organisations. A summary of the featured papers of Theme 3 is shown in Table 4.

Table 4. Summary of papers of Theme 3

Paper No.	Paper Title	Key GenAI Method	OR Method	Key Contribution
3A	Investigating the Manufacturer's Decision to Invest in GenAI: Emission-Sensitivity and Quality	GenAI as product quality/process lever	Analytical economic modelling	Non-monotone adoption thresholds; high emission-sensitivity can *reduce* carbon emissions under GenAI
3B	AI Adoption in the Presence of Consumer Resistance: Is "Win-Win" Possible?	AI as cost-reduction and personalisation tool	Stackelberg game theory	Consumer resistance is endogenous; can enhance consumer surplus even while undermining supplier-manufacturer value chain
3C	GenAI-Enhanced Risk Response in Project Portfolios: A Complex Network Evolutionary Game	GenAI as coordination enabler	Complex Network Evolutionary Game (CNEG)	GenAI + reward-penalty mechanisms eliminate asymmetric equilibria; high-capability managers show complacency effect
3D	Generative AI for Market Entry: Entrant Strategies and Uncertainty in New Drug Development	LLM-driven scenario generation	Real-options / sequential decision modelling	GenAI populates scenario space for pharmaceutical market entry; identifies robust strategies under regulatory uncertainty
3E	When AI Bites Back: The Hidden Consequences of GenAI in Food Delivery	Algorithmic personalisation / demand prediction	Game-theoretic modelling	GenAI efficiency gains distributed asymmetrically; platform operators benefit at expense of restaurant partners

Paper 3A: Investigating the manufacturer's decision to invest in generative AI: role of emission-sensitivity and quality enhancements (Chakraborty et al. 2025)

In this paper, the authors interrogate GenAI adoption through the lens of environmental economics and consumer behaviour. They develop a formal analytical model of a monopolist manufacturer's GenAI investment decision in a market segmented between emission-sensitive and emission-agnostic consumers. The model yields non-monotone propositions of significant policy relevance: GenAI adoption increases prices and manufacturer profit only when consumer emission-sensitivity falls below endogenously determined thresholds, and sufficiently high emission-sensitivity can, paradoxically, *reduce* total carbon emissions under GenAI adoption, because consumer resistance suppresses production and production is itself the primary emissions source. The paper is grounded in striking empirical context: data centres now account for 21% of national electricity consumption in Ireland, and Google's carbon emissions rose 51% between 2019 and 2024, directly undermining its sustainability commitments. The theoretical implication for policymakers is profound: consumer environmental preferences can function as a form of market discipline that formal carbon regulation supplements rather than replaces, and GenAI investment incentives are fundamentally shaped by the distribution of those preferences.

Paper 3B: AI adoption in the presence of consumer resistance: is "win-win" possible for an AI supplier and a manufacturer? (Zhao et al. 2026)

In this paper, the authors model the strategic dyad between an AI supplier (setting development level and licence fee) and a downstream manufacturer (choosing whether to adopt), using a Stackelberg game framework. The authors decompose AI's value creation into two levers: a Cost Reduction Effect (CRE) and a Consumer Personalisation Effect (CPE). The analysis reveals asymmetric strategic logic, CPE matters more when consumer quality sensitivity is low, while CRE dominates when it is high, and demonstrates that consumer resistance is an endogenous variable with complex welfare implications. Crucially, while resistance invariably impedes win-win outcomes for the AI supplier and manufacturer, it does not necessarily reduce consumer utility: when privacy valuations are sufficiently high, resistance can enhance consumer surplus even as it undermines the value chain between supplier and adopter. The paper introduces consumer resistance as a structural variable in AI adoption modelling a contribution of growing practical relevance as regulatory and public scrutiny of AI systems intensifies globally.

Paper 3C: GenAI-enhanced risk response in project portfolios: a complex network evolutionary game approach (Zhang et al. 2026)

In this paper, the authors examine how GenAI can promote Synergistic Response Behaviours (SRBs), coordinated risk management actions, across large, interdependent project portfolios. Using a Complex Network Evolutionary Game (CNEG) framework, the authors model the diffusion of SRBs across portfolio networks, comparing dynamics with and without GenAI support. The key theoretical result is that GenAI, when combined with well-designed reward-penalty mechanisms, can create a unique cooperative Evolutionary Stable Strategy that eliminates the asymmetric equilibria characterising conventional portfolio risk management. Applied to a case calibrated on China State Construction Engineering Group's portfolio (validated against the Boeing 737 MAX programme), the authors report a counterintuitive finding: portfolios with *stronger* inherent risk-response capabilities exhibit *slower* SRB diffusion, because high-capability project managers believe they can manage risks independently. This organisational complacency effect means that GenAI's cooperative value is greatest precisely where self-confidence is most entrenched. This is a finding with immediate implications for programme directors and portfolio risk officers seeking to design GenAI-augmented governance systems.

Paper 3D: Generative AI for market entry: entrant strategies and uncertainty in new drug development (Xu et al. 2025)

New drug development involves a compounding sequence of regulatory, clinical, and commercial uncertainties that render market entry strategy notoriously difficult to optimise. In this paper, the authors apply GenAI tools, specifically LLM-driven scenario generation and Monte Carlo–style probability assessment, to model strategic options for pharmaceutical entrants under multiple stages of uncertainty. The framework enables entrants to simulate competitive responses, regulatory pathways, and reimbursement scenarios that are difficult to specify analytically with traditional decision trees or real-options models. The authors make a methodological contribution by showing how GenAI can populate the scenario space of a complex sequential decision problem with contextually plausible trajectories derived from regulatory and market intelligence, rather than relying on researcher-specified probability distributions. The practical implication is that pharmaceutical firms can use LLM-augmented modelling to identify robust entry strategies, those that perform acceptably across a wide range of scenarios, and to target clinical trial investments more efficiently by anticipating competitive and regulatory dynamics early in the development pipeline.

Paper 3E: When AI bites back: the hidden consequences of generative AI in food delivery (Zhang and Zhang 2026)

In this paper, the authors take a critical stance to examine how GenAI adoption in food delivery platforms generates systemic consequences that are invisible to conventional efficiency analyses. Using a combination of game-theoretic modelling and empirical analysis, the authors demonstrate that GenAI-driven personalisation and demand prediction, while increasing platform revenue and consumer convenience, can simultaneously disadvantage restaurant partners through algorithmic repricing, demand concentration on a small number of algorithmically favoured merchants, and erosion of pricing autonomy. The paper's most provocative finding is that the efficiency gains attributable to GenAI may be distributed asymmetrically across platform stakeholders accruing disproportionately to the platform operator at the expense of the supply side of the marketplace. This distributional concern is analytically tractable within the paper's framework and suggests that regulatory oversight of AI-mediated marketplaces should attend not only to consumer welfare and privacy but to the market power dynamics that GenAI adoption can entrench. The paper joins a growing literature on the unintended consequences of algorithmic management in platform economies.

Remarks: Theme 3 reveals that the strategic dimensions of GenAI adoption are at least as analytically rich as its operational dimensions. Game theory, evolutionary dynamics, and economic modelling have distinctive and irreplaceable contributions to make in understanding when GenAI creates value, for whom, and at whose expense.

2.4. Theme 4: Human-AI Collaboration in Operations

The final cluster collects papers that address the interface between GenAI systems and human agents in operational settings: scheduling, logistics platform design, group decision-making, and service operations. These papers share a concern with how human-AI collaboration can be structured to capture the complementary strengths of each. A summary of the features papers of Theme 4 is shown in Table 5.

Table 5. Summary of papers of Theme 4

Paper No.	Paper Title	Key GenAI Method	OR Method	Key Contribution
4A	Human-AI Coordination for Large-Scale Group Decision Making with Heterogeneous Feedback Strategies	GenAI as decision agent (decisive / equal / auxiliary)	Minimum Cost Consensus (MCC), intuitionistic fuzzy sets	Heterogeneous adjustment costs aligned to AI acceptance; reduces consensus cost across all three scenarios
4B	Aircraft Push-Back Planning with Stochastic Requests and Service Times	GenAI-assisted reward shaping	Multi-Agent Reinforcement Learning (MARL)	Cooperative MARL outperforms rule-based scheduling; validated at major hub airport
4C	Incorporating Gen AI in Logistics: Optimal Cloud and Edge Platform Design for ADAS	GenAI cloud analytics	Analytical platform design (B2B / B2C / edge-only)	Profitability and customer welfare are strategic complements; co-investment by governments and automakers recommended
4D	From Trip Metrics to Corporate Disclosures: Predicting Ride-Hailing Prices	GenAI extraction from corporate disclosures	Cumulative Link Mixed Model (CLMM), longitudinal panel data	Qualitative corporate narrative improves ordinal price prediction; generalisable to longitudinal pricing in any platform market
4E	The GenAI LLM Product Design Multi-Model Framework for Manufacturing Operations	Chained specialist LLM agents	Multi-criteria design optimisation	Modular, auditable multi-agent pipeline for product design; matches expert quality at fraction of time/cost
4F	GenAI Agents are the Premise of AGI? GenAI Agents-based Integration in Enterprises' Operations	Autonomous GenAI agents (perceive–reason–plan–act)	Capability maturity framework	Positions current agents on AGI continuum; proposes governance framework for increasing agent autonomy

Paper 4A: Human-AI coordination for large-scale group decision making with heterogeneous feedback strategies (Zhang et al. 2025)

Large-Scale Group Decision-Making (LSGDM) involves many experts with different backgrounds working together to evaluate options. This can lead to problems like information overload, bias within subgroups, and difficulty reaching agreement. In this study, the authors introduce a new approach where Generative AI (GenAI) joins human experts as a decision-making participant. GenAI can play different roles: it can lead decisions, act equally with humans, or simply support them. The model groups experts based on how similar their opinions are, how much they trust each other, and how they feel about using AI. It also adjusts the importance (weights) of each group and the AI using advanced methods. The model was tested in a medical case where 20 experts evaluated treatment options for COPD. Results show that combining human expertise with AI helps reduce disagreements and reach consensus faster—even in complex situations. Importantly, when the cost of changing opinions is adjusted based on how comfortable experts are with AI, the system performs better than when all experts are treated the same.

This has strong implications for designing effective human-AI decision systems, especially in critical fields like healthcare.

Paper 4B: Aircraft push-back planning with stochastic requests and service times: a multi-agent reinforcement learning approach (Wang et al. 2026)

Airport ground operations are highly time-sensitive and unpredictable, making scheduling very challenging. In this study, the authors focus on aircraft pushback planning, i.e. how planes move from gates to the runway, when requests and service times are uncertain. The authors propose using a Multi-Agent Reinforcement Learning (MARL) approach, where multiple tug vehicles (agents) learn how to coordinate with each other by interacting with a simulated airport environment. Generative AI (GenAI) helps improve this system by shaping rewards and guiding decision-making, allowing the system to respond better to sudden demand changes compared to traditional rule-based methods. The key finding is that when these agents cooperate and share information, they perform much better than both independent learning systems and traditional scheduling methods, especially in reducing delays. Tests using real data from a major airport show that this approach works well in practice. Overall, the study demonstrates how combining AI and learning-based methods can solve complex, uncertain, and time-critical scheduling problems that are difficult for traditional techniques.

Paper 4C: Incorporating Gen AI in logistics: Optimal cloud and edge platform design for advanced driver assistance systems (Chowdhury and Singh 2026)

Advanced Driver Assistance Systems (ADAS) represent a growing frontier for logistics technology, combining edge-device sensing with cloud-based GenAI analytics to improve driving safety and fleet efficiency. In this paper, the authors develop an analytical model of a Mobility Services Solutions Provider's (MSSP's) optimal investment decisions across edge device features and cloud-service capabilities, recognising that these are complementary products that must be designed jointly. Three market scenarios are examined: B2B with exogenous pricing, B2C with endogenous pricing, and edge-device-only provision. The paper's key finding is that MSSP profitability and customer welfare are strategic complements highlighting any initiative that improves customer well-being simultaneously increases firm profit. This important result establishes MSSPs face no inherent trade-off between efficiency and customer focus in this domain. Furthermore, lower-level ADAS systems deliver proportionally greater consumer benefit from small improvements in edge device features, while higher-level systems require larger capability gains. The policy implication is that governments and large automakers should consider co-investing in ADAS technology, since all parties share common interests. The paper broadens GenAI's OR context into hardware-software co-design, a methodological territory largely unexplored in the operations literature.

Paper 4D: From trip metrics to corporate disclosures: predicting ride-hailing prices for longitudinal data using an integrated Cumulative Link Mixed Model and generative AI-based methodology (Sengupta et al. 2026)

Ride-hailing prices are influenced by many factors, such as trip distance, time, demand surges, company strategy, and broader economic conditions. In this study, the authors build a model that combines two types of data: (1) structured trip data (like distance and demand), and (2) insights extracted by Generative AI (GenAI) from company reports, earnings calls, and regulatory documents. The GenAI part helps identify early signals—such as expansion plans, regulatory changes, or competitive actions—that can affect prices even before they show up in operational data. By combining these qualitative insights with traditional data, the model predicts price categories more accurately over time. The results show that this combined approach performs much better than models using only operational data or only text analysis. The key takeaway is that adding strategic and qualitative information improves price prediction, and this method can be applied to other situations where company decisions influence outcomes before they are visible in numbers.

Paper 4E: The Generative Artificial Intelligence large language product design multi-model framework for manufacturing operations (Leung et al. 2025)

Product design in manufacturing increasingly requires the integration of engineering constraints, customer preference data, sustainability requirements, and market intelligence, a multi-criteria synthesis that is difficult for any single modelling approach to handle. In this paper, the authors proposes a multi-model LLM framework that chains together specialist GenAI models, each responsible for a distinct phase of the product design process: customer requirement extraction from unstructured text, constraint formalisation, design alternative generation, and sustainability scoring. The framework demonstrates how prompt engineering and model orchestration can coordinate a pipeline of GenAI agents to handle what would otherwise require multi-disciplinary human teams. Case studies in consumer electronics and industrial equipment manufacturing validate that the framework generates design alternatives of comparable quality to expert-led processes at a fraction of the time and cost. The paper develops a replicable architecture for LLM-based multi-model coordination in manufacturing OR contexts — an architecture that is modular, auditable, and adaptable to different product categories through corpus and schema adjustment.

Paper 4F: Generative AI agents are the premise of artificial general intelligence? The exploration of GenAI agents-based integration in enterprises' operations (Wamba-Taguimdje and Kamdjoug 2026)

In this paper, the authors ask whether GenAI agents, autonomous AI systems capable of perceiving, reasoning, planning, and acting, represent a pathway toward Artificial General Intelligence (AGI) in enterprise settings. Drawing on a systematic analysis of GenAI agent deployments across manufacturing, logistics, finance, and service operations, the authors develop a capability maturity framework that positions current GenAI agents on a continuum from narrow task automation to context-aware multi-step reasoning. The authors argue that while current systems exhibit impressive performance on bounded tasks, they remain far from AGI: they lack robust causal reasoning, genuine transfer learning across domains, and the meta-cognitive capacity to recognise the limits of their own competence. However, the paper identifies enterprise operations as a domain where the incremental deployment of GenAI agents, even in their current form, can generate substantial value through coordinated automation of repetitive, high-volume, rules-intensive tasks. The paper closes with a research agenda for GenAI agent governance in enterprise operations, emphasising the importance of human oversight mechanisms, failure mode transparency, and accountability structures as agent autonomy increases.

Remarks: The six papers of Theme 4 collectively demonstrate that human-AI collaboration in operations is neither a solved problem nor a distant aspiration; it is an active and urgent research frontier, demanding formal models of interaction structure, incentive alignment, and oversight architecture that OR is uniquely positioned to provide.

3. Conclusion and Future Research

The nineteen papers featured in this special issue collectively demonstrate that OR and GenAI could be a good match. GenAI provides OR with new capabilities, to extract structured knowledge from unstructured text, to generate heuristics adaptively, to populate simulation models with contextual intelligence, to forecast demand with richer informational bases, and to assist human decision-makers in reaching consensus at lower cost. OR provides GenAI with new rigour formal verification frameworks, game-theoretic models for behavioural and incentive analyses, stochastic models of uncertain environments, and optimisation architectures that transform generative outputs into actionable decisions.

Several cross-cutting areas deserve further and deeper investigation in the future as the field moves forward.

First, **the augmentation–automation balance** remains unresolved. The majority of papers in this special issue achieve their best results when GenAI augments rather than replaces classical OR methods, enriching simulation data, scaling heuristic search, or calibrating stochastic parameters, while full automation of OR workflows remain constrained by verification challenges. Future research should develop formal criteria for determining when and under what conditions each paradigm is appropriate.

Second, **trust, interpretability, and hallucinations**[1] remain critical barriers to operational deployment. In this SI, the healthcare simulation paper (Monks et al. 2025) documents instances where LLMs produce plausible but logically flawed code; the group decision-making paper (Zhang et al. 2025) proposes heterogeneous feedback costs calibrated to expert trust in AI as a partial solution. It is crystal clear that the related fields of studies urgently need to establish scalable, semi-automated verification frameworks for GenAI-generated OR outputs, particularly in high-stakes domains such as healthcare, critical infrastructure, and financial markets.

Third, **distributional consequences of GenAI adoption** deserve sustained attention. In this SI, the food delivery paper (Zhang and Zhang 2026) and the consumer resistance paper (Zhao et al. 2026) both demonstrate that GenAI's efficiency gains can be distributed asymmetrically, benefiting platform operators and AI suppliers while imposing extra costs on suppliers and workers, and hurting consumer welfare. The traditional OR multi-objective analysis and mechanism design would be useful to help design adoption frameworks that are simultaneously efficient and equitable.

Fourth, **GenAI for OR methodology,** as opposed to GenAI applied to operational problems, is an emerging and underexplored research frontier. Our featured papers using RAG for simulation parameterisation (Nie et al. 2026), RAG for supply chain mapping (Jackson et al. 2026), and GenAI for healthcare DES reconstruction (Monks et al. 2025) all suggest that the most significant methodological contributions may not lie in using GenAI to solve known OR problems faster, but in enabling entirely new categories of evidence-based modelling, which were previously infeasible.

Finally, **open science and reproducibility** warrant special attention. This is also related to "responsible AI" - developing and using AI in a trustworthy and ethical way. As GenAI becomes embedded in OR frameworks, establishing clear standards for documenting the prompts, models, validation protocols, and computational pipelines used is essential for enabling independent replication and scholarly accumulation. We hence encourage future researchers in OR to treat GenAI component documentation seriously (e.g., with the same level of rigour currently expected of mathematical proof and computational experiment design in standard OR studies.)

We are grateful to all authors, reviewers and the *JORS* editorial board members for their kind contribution to make this collection possible. We hope these nineteen papers serve as a foundation for the next generation of research at the frontier of OR and GenAI.

Disclosure statement: No potential conflict of interest was reported by the authors.

[1] In the domain of GenAI, hallucinations commonly refer to some AI generated false or misleading responses, which sound correct and confident.

References

1. Aghsami, A., et al. (2026). Optimizing retailer-supplier inventory systems with defective returns: A continuous-time Markov chain and neural network-based approach. Journal of the Operational Research Society.
2. Chakraborty, A., Mandal, P., & Kumar, R. R. (2025). Investigating the manufacturer's decision to invest in generative AI: role of emission-sensitivity and quality enhancements. *Journal of the Operational Research Society*, 1–15. https://doi.org/10.1080/01605682.2025.2579853.
3. Chowdhury, S., & Singh, N. K. (2026). Incorporating Gen AI in logistics: Optimal cloud and edge platform design for advanced driver assistance systems. *Journal of the Operational Research Society*. https://doi.org/10.1080/01605682.2026.2618516.
4. Dahbi, H., Chaouni Benabdellah, A., Belhadi, A., Kamble, S., & Zouadi, T. (2025). Generative heuristics-driven for blockchain-enhanced reputation systems and dynamic optimization in short food supply chains. *Journal of the Operational Research Society*, 1–23. https://doi.org/10.1080/01605682.2025.2573824.
5. Golghamat Raad, N., Fick, J., Shah, N., & Srinivas, S. (2025). From voices to value: Leveraging online reviews to benchmark airline efficiency with GenAI-assisted text analytics and DEA. *Journal of the Operational Research Society*, 1–24. https://doi.org/10.1080/01605682.2025.2568162.
6. Jackson, I., Saenz, M. J., Ivanov, D., & Ma, B. J. (2026). Supply chain mapping through retrieval-augmented generation: Applications to the electronics industry. *Journal of the Operational Research Society*. https://doi.org/10.1080/01605682.2025.2608868.
7. Leung, E. K. H., Balan, N., Lee, C. K. H., & Xie, S. (2025). The Generative Artificial Intelligence large language product design multi-model framework for manufacturing operations. *Journal of the Operational Research Society*, 1-27. https://doi.org/10.1080/01605682.2025.2570407.
8. Li, P., Qu, T., Wu, N., Zu, Y., Wen, F., & Tang, Z. (2026). Generative AI–enhanced sales forecasting and optimization for the multi-period multi-product inventory routing problem with procurement. *Journal of the Operational Research Society*, 1–24. https://doi.org/10.1080/01605682.2026.2618519.
9. Monks, T., Harper, A., & Heather, A. (2025). Unlocking the potential of past research: Using generative AI to reconstruct healthcare simulation models. *Journal of the Operational Research Society*. https://doi.org/10.1080/01605682.2025.2554751.
10. Nie, W., Li, F., Tsolakis, N., & Kumar, M. (2026). Integrating digitally enhanced data extraction and simulation modelling for AI-driven supply Chain resilience: an operational research framework for strategic stockpiling of critical minerals. *Journal of the Operational Research Society*, 1-28. https://doi.org/10.1080/01605682.2025.2612145.
11. Sengupta, P., Adhikari, A., & Pal Singh, S. (2026). From trip metrics to corporate disclosures: predicting ride-hailing prices for longitudinal data using an integrated Cumulative Link Mixed

Model and generative AI-based methodology. *Journal of the Operational Research Society*, 1-22. https://doi.org/10.1080/01605682.2026.2640139.

12. Thomassey, S., Tran, K.P., Vanderhaegen, F., & Choi, T.M. (2026) *Responsible Artificial Intelligence and Data Science. Springer*. https://doi.org/10.1007/978-3-032-14055-5.
13. Wamba-Taguimdje, S. L., & Kala Kamdjoug, J. R. (2026). Generative AI agents are the premise of artificial general intelligence? The exploration of GenAI agents-based integration in enterprises' operations. *Journal of the Operational Research Society*, 1–30. https://doi.org/10.1080/01605682.2025.2612146.
14. Wang, Q., Chung, T. H. C., & Wen, X. (2026). Aircraft push-back planning with stochastic requests and service times: a multi-agent reinforcement learning approach. *Journal of the Operational Research Society*, 1-17. https://doi.org/10.1080/01605682.2026.2618517.
15. Xu, X., Nie, J., & Zhao, Y. (2025). Generative AI for market entry: entrant strategies and uncertainty in new drug development. *Journal of the Operational Research Society*, 1–17. https://doi.org/10.1080/01605682.2025.2605228.
16. Zhang, J., Wang, N., & Tang, M. (2025). Human-AI coordination for large-scale group decision making with heterogeneous feedback strategies. *Journal of the Operational Research Society*. https://doi.org/10.1080/01605682.2025.2466677.
17. Zhang, X., Arisian, S., Li, H., Bai, L., & Shen, W. (2026). GenAI-enhanced risk response in project portfolios: a complex network evolutionary game approach. *Journal of the Operational Research Society*, 1-26. https://doi.org/10.1080/01605682.2026.2633436.
18. Zhang, Y., & Zhang, L. (2026). When AI bites back: the hidden consequences of generative AI in food delivery. *Journal of the Operational Research Society*, 1-19. https://doi.org/10.1080/01605682.2026.2630922.
19. Zhao, Q., Fan, J., & Wang, Z. (2026). AI adoption in the presence of consumer resistance: is "win-win" possible for an AI supplier and a manufacturer? *Journal of the Operational Research Society*, 1–18. https://doi.org/10.1080/01605682.2026.2633437.
20. Zhou, Q., & Sheu, J. B. (2025). The use of Generative Artificial Intelligence (GenAI) in operations research: Review and future research agenda. *Journal of the Operational Research Society*, 1-21. https://doi.org/10.1080/01605682.2025.2561762.