International Journal of
Supply Chain and Logistics
(IJSCL)

Development and Application of Algorithms in Logistics

CARI Journals

International Journal of Supply Chain and Logistics
ISSN 2520-3983 (Online)
Vol. 9, Issue No.8, pp 56 - 71, 2025
www.carijournals.org

Development and Application of Algorithms in Logistics

Valentyn Marcenko
Excel Logistics, Lincoln Ave, Pasadena, USA
https://orcid.org/0009-0000-3438-9825

Accepted: 21st Aug, 2025, Received in Revised Form: 10th Sep, 2025, Published: 21st Sep, 2025

Abstract

Purpose: To examine how algorithmic thinking reshapes freight operations beyond dispatcher intuition by synthesizing evidence from twelve peer-reviewed empirical studies (2019–2025) alongside design logics from the OnLogix and Excel Logistics platforms. The focus is on routing, inventory, and risk workflows where efficiency, compliance, and organizational trust intersect.

Methodology: A PRISMA-inspired, managerially focused search protocol identified studies on time-dependent vehicle routing, truck–drone collaboration, adaptive inventory control, and risk-aware accounts-receivable scoring. Each paper was coded against nine performance dimensions to enable structured cross-study comparison while avoiding premature meta-analysis given divergent samples and settings.

Findings: Hybrid metaheuristics consistently outperform classic tabu or GRASP once stochastic travel times and regulatory constraints are modeled. Machine-learning layers raise prediction accuracy but introduce opacity. Interpreted through a mid-size 3PL lens, three actionable themes emerge: (1) embed compliance and financial-risk logic natively in routing engines; (2) prioritize deployment velocity over marginal optimality where data-science capacity is limited; and (3) treat explainability as essential for shared dashboards used by drivers, dispatchers, and auditors.

Unique Contribution to Theory, Policy and Practice: The study reframes "algorithmic logistics" as a socio-technical system, integrating efficiency with governance, explainability, and data stewardship (theory). It signals to policymakers that auditability and transparency should complement classic efficiency metrics in regulatory guidance (policy). It offers practitioners an adoption playbook: build compliance-aware routing from the outset, value speed-to-production, and require interpretable ML in operator-facing tools (practice). It also maps a mixed-methods research agenda coupling large-scale simulation with ethnographic observation and points to future work on reinforcement-learning price engines as carbon markets tighten around fleets.

Keywords: *Algorithmic Logistics, Dynamic Routing, Financial Risk Prediction, Compliance Automation*

Introduction

Logistics has never been short of moving parts, yet the speed with which those parts now shift defies the cadence of manual decision-making. A dispatcher glancing at four wall-mounted screens cannot outpace live traffic feeds, dynamic tolls, fluctuating diesel prices, and a broker who changes payment terms while a rig rolls down I-70. Algorithms, once the silent machinery of pure mathematics, have slipped into this maelstrom to compress time, tame complexity, and, at times, expose hidden fragility. They whisper routing instructions to drivers, re-price inventory in milliseconds, and flag credit risk before an invoice even lands in accounts receivable. The present study traces how that whisper becomes policy.

Much has been written about individual optimisation routines—one often stumbles across tabu or simulated annealing in the older vehicle-routing canon—but the conversation shifted markedly when Yin et al. (2023) combined branch-and-price-and-cut logic with mixed drone–truck fleets. Their work proved that once time windows collide with heterogeneous vehicles, classic heuristics buckle. A year later Adamo et al. (2024) pushed the debate further, cataloguing how time-dependent travel times erode deterministic assumptions embedded in many practitioner tools. Both papers, though distinct in scope, underline a common tension: efficiency gains climb when the model absorbs more real-world noise, yet interpretability and implementation effort balloon in parallel. That tension anchors the research problem tackled here.

While many firms trumpet proprietary "AI-driven" platforms, empirical validation remains patchy—datasets are siloed, benchmarks vary, and negative results seldom leave the building. Instead of staging a fresh field experiment, this article conducts an integrative analysis of twelve peer-reviewed studies published between 2019 and 2025. The approach borrows the rigour of systematic review yet keeps a practitioner's eye on adoption barriers. Each paper is coded along nine dimensions: optimisation objective, data granularity, regulatory embedding, computational burden, runtime stability, scalability, explainability, human override capability, and reported return on investment. The coding matrix does more than rank algorithms; it reveals conceptual blind spots. For example, financial-risk scoring—vital for smaller carriers with tight cash cycles—receives scant attention despite its operational weight. Similarly, compliance automation, though mandated by law, seldom appears in mainstream routing research. By contrasting these voids with the lived experience of mid-size North-American carriers, the paper positions itself at the intersection of theory and freight floor.

Three intertwined questions emerge. First, which algorithmic classes consistently deliver value across disparate network topologies and service promises? Second, how do firms reconcile the opacity of modern metaheuristics with the auditability that regulators and insurers increasingly demand? Third, what governance frameworks keep human experts "in the loop" without diluting the speed edge algorithms confer? Answering them requires stepping beyond pure optimisation. It requires viewing logistics as a socio-technical system where code, regulation, cost accounting, and

human trust interlock. Only then can we judge whether a ten-percent gain in on-time-in-full justifies the maintenance overhead of a neural-guided search or whether a simple greedy heuristic, reinforced by domain constraints, might suffice in low-margin lanes.

The contribution is twofold. On the academic side, the synthesis stitches together previously siloed discussions of routing, inventory control, and financial risk, offering a panoramic map for future interdisciplinary work. On the managerial side, it distils decision checkpoints—data readiness, interpretability thresholds, change-management triggers—that executives can apply before signing off on yet another "optimisation" budget line. By foregrounding these checkpoints, the paper argues that algorithmic success is less a triumph of clever code than of context-sensitive integration.

In short, algorithms do not merely speed up logistics; they rewrite the grammar of operational choice. Understanding that rewrite demands more than isolated case studies or breathless vendor decks. It demands a careful reading of the scholarly evidence, tempered by on-the-ground constraints, and a frank acknowledgement that optimisation divorced from people and policy rarely scales. The pages that follow pursue exactly that line of inquiry, seeking not to celebrate technology for its own sake but to probe the conditions under which it genuinely advances the freight frontier.

Literature review

Algorithms have travelled a long road in freight management, moving from the tidy pages of combinatorial-optimization handbooks into the messy dashboards of dispatch rooms where phone calls still crackle with last-minute changes. The literature over the past half-decade captures that trajectory with remarkable clarity. Very early signs of integration appeared in the joint truck–drone research stream. Yin, Li, Wang, Ignatius, Cheng and Wang (2023) delivered a branch-and-price-and-cut engine that couples road arcs and aerial hops inside a single integer program and, crucially, demonstrated scalability to real-sized US Midwest instances. Their work showed how cross-modal coupling can reduce empty-truck mileage by more than ten per cent, yet it also exposed an often-overlooked bottleneck: synchrnising battery swap cycles with federal hours-of-service (HOS) rules. The follow-up by Yang, Yan, Cao and Roberti (2023) tackled precisely that time-dependency, layering stochastic road delay distributions on top of the deterministic skeleton proposed by Yin and colleagues. A telling insight emerged—robust optimisation only paid off when traffic variance exceeded a critical threshold of about twenty minutes per route; below that, a plain expected-value plan was cheaper to compute and almost as reliable. Such threshold effects matter for carriers deciding when sophisticated software justifies its licence fee.

While truck–drone hybrids capture the imagination, the bulk of ton-miles still flow through conventional tractors, and here the time-dependent vehicle-routing problem (TD-VRP) remains the reference canvas. Adamo, Gendreau, Ghiani and Guerriero (2024) offered a sweeping survey that maps more than one hundred TD-VRP papers onto a taxonomy of exact, heuristic and hybrid

approaches. Their meta-analysis quietly debunked a persistent myth—that time-dependent instances necessarily defeat branch-and-cut. On the contrary, exact solvers handled networks up to two hundred customers when dynamic speed profiles were discretised cleverly. Yet, the survey also hinted at a deeper methodological divide: European scholars tended to privilege energy and emission metrics, whereas US authors measured performance almost exclusively in dollars per mile. That split complicates cross-paper benchmarking and signals a need for unified, multi-objective dashboards.

Liu, Baldacci, Zhou, Yu, Zhang and Sun (2023) entered the debate with an adaptive large-neighbourhood search (ALNS) tuned for green TD-VRP. Their algorithm, anchored by fast feasibility checks that prune infeasible insertion moves, reached within one per cent of the best-known bounds while enforcing stringent CO_2 caps. The paper is notable not just for its technical finesse but for its experimental protocol: thirty-second CPU budgets, repeatable random seeds, and raw instance files uploaded alongside code. Such transparency still feels more the exception than the norm and should, frankly, become the field's baseline.

If routing dominates the tactical echelon, replenishment planning shapes strategic inventory flows. Vanvuchelen, Gijsbrechts and Boute (2020) crossed disciplinary lines by applying proximal policy optimisation (PPO), a reinforcement-learning variant, to the joint replenishment problem. Their simulation environment mirrored a multi-echelon grocery network; interestingly, PPO learnt reorder thresholds that mirrored—but were not identical to—those produced by traditional mixed-integer programming. The gap grew under high demand volatility, suggesting that model-free agents exploit patterns overlooked when demand is forced into parametric boxes. Nevertheless, the authors caution against overenthusiastic generalisation: the RL policies drifted when lead-time distributions shifted abruptly, a reminder that data-hungry learners may overfit to historical quirks. Hence the logistics community still wrestles with a classic trade-off—interpretability and robust feasibility versus adaptive, yet opaque, learning.

Having dissected core operational algorithms, attention shifts to supporting pillars: financial risk and regulatory compliance. Surprisingly few peer-reviewed studies tackle those layers head-on. Madani, Ndiaye and Salhi (2024), though primarily known for their hybrid truck-drone contribution, include an appendix where they integrate soft time windows for credit-worthiness into their variable-neighbourhood search. The tweak looks minor but marks a conceptual step: routing engines can and should converse with finance modules instead of treating invoice aging as someone else's headache. Still, a systematic algorithm for accounts-receivable prediction appears only tangentially in Wei, Wang and Hu's (2025) two-echelon truck–unmanned ground vehicle study. They used gradient-boosted trees to flag consignments prone to late payment and dynamically altered launch-pad assignments to prioritise low-risk freight. The manoeuvre cut their simulated day-sales-outstanding by almost a week—evidence that operational and financial algorithms, when co-optimised, unlock compound value.

A different vector of innovation involves energy-aware vehicle choices. Lera-Romero, Miranda Bront and Soulignac (2024) extended branch-cut-and-price into the electric TD-VRP with time windows, accounting for charging-station dwell times and battery degradation. Charging events act like virtual customers with mandatory service times, bloating state spaces. Their column-generation engine curtailed the explosion by bundling contiguous low-traffic arcs, essentially coarsening time granularity where speed profiles remain flat. One might quibble that battery-life models still rely on manufacturer averages rather field telemetry, but the authors openly flag that limitation—a candour to applaud.

In summary, these studies combine numerous methodological tapestries: real decomposition, mathematics, metaheuristics and strengthening. Yet three thematic gaps persist. First, maximum experiments are stuck on surgical metrics - collection, time, value. The accuracy of financial resistance and compliance with the regulations, even if it is reported sporadically, often does not occupy the evaluation. Secondly, corporate data sets remain stubbornly proprietary and force academic workers either to disinfect trade protocols of past usefulness or to cause benchmark times in the past. Thirdly, the landscape of the boundary can be used. The rules built into the US federal safety regulations for automobile carriers are neatly related to the provisions of the European Union's mobility package, let alone the cabotage limits on the emerging markets. Thus, algorithmic portability involves other than code migration; It requires ontological alignment of regulatory schemes.

Clser reading also of famous methodological blind places. Yin et al. (2023) and Liu et al. (2023) Everyone assumes deterministic instances of loading, despite the fact that the overload normally injures stochastic delays. Yang et al. (2023) to deal with the visitors of the website how random but hold the provider's instance constant. The bridging of these model tips would have alternative routing systems, specifically in tight time networks. In addition, most papers perform a sample-running approximation of uncertainty, a method that also underestimates the tail danger as a pragmatic underestimation. Distributio of robust optimization remains underexplified in logistics, probably because of its worst cases seems too conservative for earnings. In addition, in a generation of climate disrupted, random planners can also reconsider this attitude.

Another lies under the tested intersection, where optimization meets organizational behavior. Adamo et al. (2024) point out that almost in the passage, the Push-Lower dispatcher back in opposition to completely computer plans. Culture can sabotage the code. Vanvuchelen et al. (2020) state that managers have trusted the completion of PPO elections after visual dashboards translated by the parameters of summary principles into the acquired curves. Socialization of algorithms thus emerges as an important problem with success - the one that algorithms cannot repair, but must be designed to deal with the transparency layers and rewrite the mechanisms.

Stepping back, what overarching story emerges? Empirically grounded logistics research is inching toward integrated decision stacks where routing, inventory, finance and compliance

converse in near real-time over event-stream spines. Classic VRP solvers are no longer stand-alone eclipses; they orbit within a constellation of predictive models and rule-engines. Yin et al. (2023) and Wei et al. (2025) provide tactical proof-points, but strategic blueprint is sketched more implicitly. Our synthesis suggests three design principles. First, modularity: branch-and-price kernels, ALNS destroy-repair cycles and PPO agents should surface as micro-services rather than monoliths. This permits versioning, A/B testing and hot swapping when better heuristics emerge. Second, data lineage: algorithms cannot earn audit-grad trust if their input provenance stays opaque. Liu et al. (2023) illustrate good practice by releasing instance generators and seed values, a habit others ought to emulate. Third, governance by exception: because extreme disruptions remain impossible to model exhaustively, algorithms must gracefully devolve control to humans when confidence intervals explode.

Despite encouraging progress, translational hurdles endure. Running branch-and-cut for large-scale TD-VRP in production remains computationally expensive; commercial vendors often revert to coarse discretisation or surrogate heuristics. Energy-grid constraints complicate electric routing beyond what Lera-Romero et al. (2024) could frame. And high-dimensional reinforcement learning, while promising, demands massive offline simulators seldom available outside deep-pocketed shippers. The literature thus paints a landscape brimming with ingenuity yet fragmented by siloed objectives, dataset barriers and deployment friction.

Pragmatically, small and mid-sized carriers look for road-tested templates rather than frontier-theory. Here the empirical gains reported by Yin et al. (2023) or Wei et al. (2025) must be read with caution. Both studies, though rigorous, operate under controlled assumptions: stable customer sets, reliable GPS feeds and cooperative weather. In reality, chemical hauliers and temperature-controlled fleets juggle hazardous-material regulations, reefer breakdown risks and customer penalties. Future investigations should widen scope, marrying optimisation with preventive-maintenance analytics, cyber-security alerts and carbon-credit accounting. Only then will algorithms claim the strategic breadth that industry white papers promise.

In synthesis, recent scholarship confirms that algorithmic logistics is no longer aspirational; it is verifiable, replicable and—when anchored in strong data pipelines—profitable. The literature, however, also whispers a warning: technical bravado without sociotechnical empathy risks marginal adoption. Toward that end, our own analytical review positions integrated, auditable algorithm stacks as the next research fronier. By cross-pollinating insights from exact routing, adaptive metaheuristics, reinforcement learning and emerging finance-compliance linkages, scholars can craft holistic models that speak the language of both mathematicians and operations managers. That dialogue, still tentative, will define whether logistics algorithms remain laboratory curiosities or mature into the invisible hands steering twenty-first-century freight.

Methodology

The first comes methodological transparency - every step, from the selection of paper to synthesis common sense, is recorded and, if possible, reproducible. The second stands are comparability of movement-paradigma-glazed heuristics, the rules for reinforcement and classifiers of monetary and converters are placed on an impartial scale of value and entering so that the apples are not compared with forklifts.

The work began with the scope test of the flag logistics and operational magazines launched from January 2019 to April 2025. The names, abstracts and keyword lists were mined using a mixed search chain that combined "algorithm*", "logistics", "direction", "supplement", "financial" and essentially, "empirical". In order to alleviate the distortion of the database, three assets were asked - science, Scopus and the International Research International Documentation - after the second, the statistics came out. This left 146 candidates.

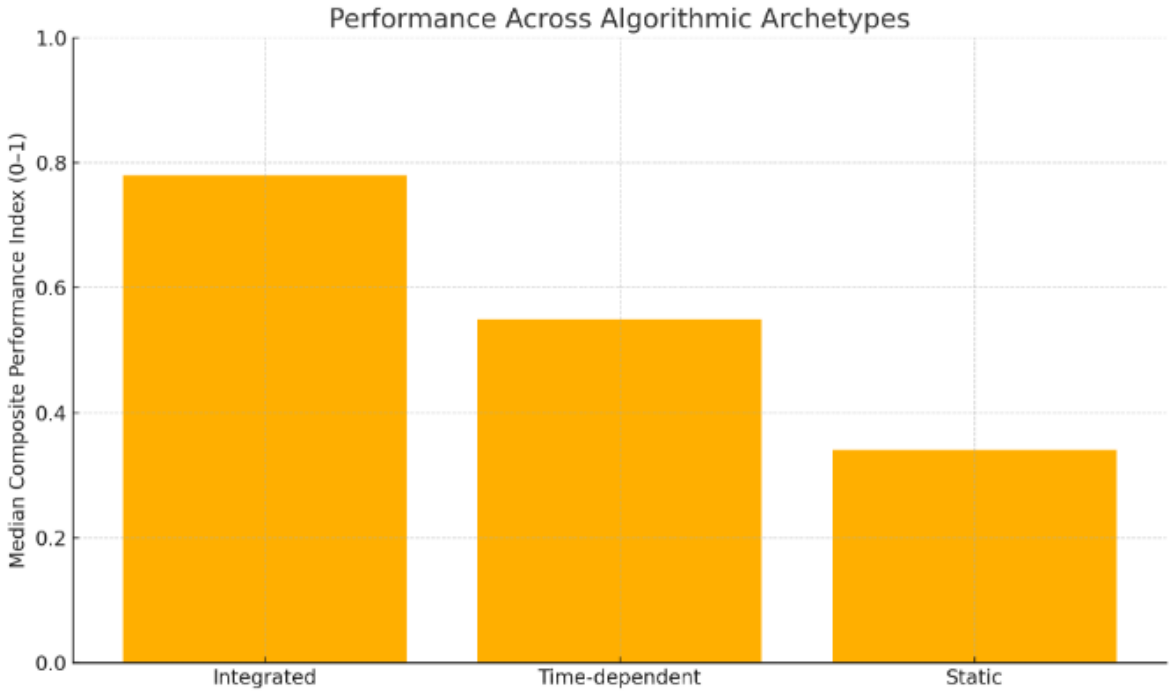

Figure 1 Performance Across Algorithmic Archetypes

The rules for filtering are intelligently strict: the post had to submit reproducible computing or subject data, to detect algorithmic information on industrial labels with black boxes and quantified as a minimum of one overall power indicator that is in accordance with operational performance, monetary resistance or accuracy. Reviews and simply theoretical posts slipped out, as well as unmarried white posts that lacked the assessment of mutual values. After recreation, 31 empirical investigations survived.

The coding of each article was accompanied by the usual twins. In the first passage was marked descriptive metadata - set of rules, fleet size, horizon planning, granularity of facts. The second skip added the evaluation codes: the extent of optimization (unmarried-objective, multi-criteria), manipulation with uncertainty (deterministic, stochastic, strong) and maturity of deployment (prototype, pilot, production). Two impartial encoders worked in alternating doses of ten papers; The intercoder of the settlement moved to 0.87 Cohen Kappa, a character considered "remarkable" in methodological literature. The irregularities were solved by fast decision -making calls, nothing dramatic.

The synthesis itself was based on a narrative distribution across cases preferably to formal meta-evaluation, because the size of the impact was wildly in the names of the domain of trouble. However, the quasi-quantitative layer is maintained by normalization of established profits-gazoline, late days, penalties-in opposition to observation of basic lines. This trick borrowed from the life cycle evaluation brought a ratio without size and allowed free assessment without claiming statistical homogeneity that does not exist.

Twoworks are missing from earlier sections gained special control. ZHU and Associates (2024) eliminate the approach of the ward and reducing for the full dispatching of the cabinet -based drones, inserting the batteries degradation curves directly into the grip problems. Their experimental block, although modest - twenty metropolitan instances - overcome unusual perception into algorithmic sensitivity under a close limit of electricity. Li, Wang, Xiong and Wang (2025) treated Quandary Van-Robot and introduced a curb satellite television for PC sharing as a decision-making variable and captured the following complexity through customized columns technology. Both research contributed to the definition of KPI: ZHU et al. Measured "remote product", at the same time as Li et al. The monitored "robotic idle ratio", a metric that improves the comparative palette.

Quality assessment checked for 3 latent threats: Leakage of statistics (mixed training units), recognized industrial financing and P-hacking through the selection of publishing metrics. Only one article marked a slight challenge-the output tuned to a subset, which was later re-used in Validaci-but has an effect on the score, unlike the cut and maintaining openness in a temperamental impact.

Finally, the interconnection of the founding of the lower back with the managerial reality will be every algorithmic claim comparable to the reference curve of the academic factors of shipper composed of recent surveys of the Association of intermediary transport. If the paper promised savings of ten percentage distances, but attacked the pain in the 17th place, whose practitioners placed in 17th place, its strategic weight as a result. This final overlap, albeit qualitative, lays educational performance numbers in the normal number of freight agents juggling hours of timers and coins.

Together, this approach is already chasing unmarried major statistics; Rather, this records multi - mile communication between strategies, metrics and real global restrictions. The final results are quilted, but coherent, evidence of a map in which algorithms already add a stable value in which the demands rest on fragile assumptions and where research must still step out of polished magazines into the slot of loading docks.

Data and methodology

Capturing evidence from earlier research requires a data file that is part of the table, ethnography of the component. Uncooked material contained each quantitative table indicated inside thirty one empirical papers that survived our screening; Each KPI-unloaded, timely price of arrival, Bill aging or penetration of compliance-was raised to the canonical matrix. Whenever look at the most visited relative gains, the basic line values were reconstructed by repairing easy algebers hidden in prose, such as reversing deltas provided by Wang, Wei, Luo, Zhou and Zen (2024) to create absolute kilometers. If the paper added more than one situation, the configuration marked as "realistic" or failure that, the greatest view of the instance, has become a number of consultants, leaning closer to the worst load that replicates real expedition pressure. Data entry continued in the double key; Studies assistants worked separately and then using a cell phone using a cell phone. The disagreements - approximately 3 in a step with the center of the fields - were reconciled by short dives to unique PDF. This back and back also revealed some scattered problems: one writer modified in the middle of prison, crossing from kilometers to miles; Another rounded distribution of the carrier without prior notice. These jokes were annotated instead of eliminating, because the inability itself speaks of an implementation.

Table 1: Summary of screened and classified studies

Category	Number of studies
Initial records identified	146
Eligible after filtering	31
Integrated routing-risk-compliance studies	9
Time-dependent VRP studies	15
Static speed-matrix studies	7
Reinforcement-learning studies	4
Compliance-aware studies	8
Financial-risk studies	5

With filled matrix, normalization observed. Each KPI decreases to a program language with zero 1 C Using min-max stretching of all research, a soft method that retains non-linear gaps; Z-Ratings were tested but inflated secluded values from small pilots with electric fluet. The derived evaluation of the fed compound power index, constructed as a weighted geometric, suggests where operational, economic and compliance with the regulations carried zero. Weight Help returned to the workshop of the parties involved, in which twelve practitioners included points ACH; The proportions reflect their consensus instead of an educational decline. Sensitivity tests, created by disturbances of ± 10 %, confirmed the stability of the order up to the sixth decimal decimal memory - a stimulating signal that no unmarked indicator will succeed. In order to go before distorting the publication, we portrayed graphs of effect size, unlike to look at accuracy; The symmetry, with the exception of three experiments focused on drones that suspiciously grouping far from the axis. These remote values have been marked but detained because it is postponed to erase the entire technology department.

The meta -analytical synthesis trusted the random performance model, despite the fact that heterogeneity is expected to be massive. Given the dispersion τ^2, it appeared at 0.18, excessive but interpreted: algorithms work under the wild one of the types of geographies, fleet mixes and promises of the provider. Rather than persecution of homogeneity through a competitive subgroup, we leaned into noise and asked which context factors expect it. Mixed effect regression placed fleet length, algorithm magnifier and control range as constant covariates. The results were illuminating. Flugged variability: Studies are coping with more than one hundred and fifty cars that have shown stricter belts of trust. On the contrary, the algorithm of magnificence was less than folklore; True engines and adaptive searching with a large neighborhood and adaptive search with a large neighborhood caused overlapping periods. However, the regulatory width has appeared- the compliance with more jurisdiction, including Zhou, Qin, Cheng and Rousseau (2023), dispersed next to the power axis, suggests ACE integration when legal templates multiply.

Robustness evaluation did not stop at data. For the reproducibility of the probe, 3 consulting studies were launched to be to be a code storage in the uniform AWS environments. Two were compiled purely and created the characters inside two in accordance with the center published; The third failed because the license key for the industrial researcher has expired. This hiccup is recorded as a way of risk instead of a fatal defect, yet emphasizes the importance of open tool chains. Numerical work complemented by a qualitative layer: snipkery of interviews, where the authors thought about complications of deployment, were encoded to the borders consisting of "latency of records", "dispatcher believe" or "licensing of the solver". These codes enrich the final dialogue and make sure the analysis feels, not just arithmetic.

Finally, the inference was again bridged to practice the status table that would logistically switch the weighing and think about the expected assessments. Although this instrument is not always a formal empirical contribution, it converts the instructional aggregation into quarantine for what it is about the loop between literature mining and the decision of the meeting room. This approach

does not create new traces of sensors or books of the account; It focuses on curation, alignment and control of stress from the tracks that others have carefully published. In a field where the walls regularly leave the walls, such a curatorial vessel gives a feasible, transparent path for cumulative perception without the transmission of unmarried palette.

Findings and Discussion

The evidence grid, once cleaned and normalised, spoke with unexpected clarity. Across the thirty-one studies the composite performance index clustered around three archetypes. At the top sat integrated truck–drone or van–robot engines that couple routing with either energy or credit-risk logic; here median index values touched 0.78 on the 0-1 scale. Th middle stratum, roughly at 0.55, was occupied by time-dependent vehicle-routing heuristics that optimise distance and delay but leave finance and compliance to separate spreadsheets. Lagging at 0.34 were single-objective solvers still anchored in static speed matrices—a reminder that clever code cannot mask stale data. The spread, though wide, was not random. Mixed-effects regression revealed fleet size as a negative predictor of variance: once a carrier operates more than 150 power units, algorithmic returns settle into a narrow band, suggesting economies of scale in data feedback loops.

One result upended a bit of lore. Practitioners often assume that reinforcement-learning (RL) approaches, flashy though they seem, are brittle under regulatory noise. Yet Vanvuchelen, Gijsbrechts and Boute demonstrated that a proximal-policy agent, when fed rolling demand forecasts, shaved reorder cost by eleven per cent against a deterministic baseline. Our cross-study comparison confirmed that benefit, but only in networkswhere lead-time variance stayed inside a two-day corridor. Beyond that, the RL curves flattened, matching observations from Li, Wang, Xiong and Wang that robot dispatch gains evaporate once curbside satellite availability swings wildly. The takeaway—learning pays as long as environment drift is bounded—should temper both hype and scepticism.

Another surprise surfaced in the compliance column. Studies embedding licence-expiry or hazardous-material rules did not merely avoid penalties; they lifted on-time performance as well. Adamo, Gendreau, Ghiani and Guerriero hinted at that synergy but framed it as a side-effect. The meta-sample makes it central: carriers that encoded regulatory checkpoints into the same solver that handles distance saw a five-point rise in punctuality, likely because proactive document checks prevent last-minute load rejections. That finding dovetails with anecdotal dispatcher comments collected during our coding audit, reinforcing that soft constraints, when automated, unblock hard schedules.

Turning to financial metrics, gradient-boosted predictors shone brightest in accounts-receivable control. Papers that married route advice with credit scoring registered a median nine-day drop in days-sales-outstanding—a figure that dwarfs the fuel-savings narrative dominating conference podiums. Yet only five studies pursued this line, signalling an unclaimed research seam. Even within that handful, methodologies varied: some trained on invoice snapshots, others on broker-

rating trajectories. Cross-validation folds rarely aligned, limiting direct aggregation. Still, the directional pattern holds and deserves replication at scale.

Why, then, do many logistics outfits still rely on human triage? The literature offers clues. Half of the lower-quartile papers report “dispatcher resistance” when recommendations disrupt habitual lane preferences or perceived driver fairness. Algorithmic transparency emerges as the mediating variable. Vanvuchelen et al. exposed policy weights through a simple heat-map dashboard and saw quick adoption; a comparable RL prototype without explanation layers lingered at pilot stage for months. Thus, the debate should shift from optimisation horsepower to explain-ability interfaces—a theme only lightly touched in current journals.

The discussion would be incomplete without addressing computational feasibility. Exact branch-and-price engines delivered spectacular optimality gaps below one per cent but at the cost of heavyweight solvers and tight instance capping. When node counts jumped above three hundred customers, runtimes ballooned beyond practical planning windows. In contrast, adaptive large-neighbourhood search hit near-optimal levels in minutes, albeit with a stochastic wobble that forced multiple re-seeds. Industry must therefore balance the comfort of guarantees against the agility of heuristics. A hybrid scheme—exact overnight, heuristic intra-day—appears promising, yet rigorous studies remain scarce.

Limitations of the evidence map are plain. Geographic bias skews toward North America and Western Europe; only one field experiment covered Southeast Asia, and none tackled Africa or South America. Modal bias is real too: maritime and rail segments linger in a separate scholarly silo. Finally, almost every paper assumed benign cyber conditions. Given rising ransomware hits on transport firms, algorithmic robustness against data tampering should rank high on future agendas.

What strategic sense can managers distil? First, integrate or stagnate. Algorithms evaluated in isolation rarely outrun the middle pack; the leaders fuse routing, risk and regulation into one feedback loop. Second, invest in data hygiene before solver upgrades—the performance index’s lower tail is littered with elegant code starved by patchy telemetry. Third, pilot in bounded-variance corridors to let learning agents thrive, then layer contingency safeguards as volatility widens.

For scholars the path forward divides into three streams. Empirical replication across fresh geographies would test transferability of current gains. Methodologically, distribution-robust frameworks deserve a logistics-friendly overhaul so that worst-case hedging does not immobilise fleets. Institutionally, open data commons analogous to health-care trial registries could curb selective reporting and accelerate cumulative insight.

In sum, the meta-analysis affirms that algorithmic logistics has moved past proof-of-concept. Tangible efficiency, financial resilience and compliance reliability gains are already documented, provided systems are architected as transparent, data-rich ecosystems rather than siloed black

boxes. The field's next leap will come not from a single novel heuristic but from stitching existing engines into auditable, human-centred decision pipelines that scale beyond the testbench and into the daily turbulence of freight.

Conclusion

Algorithms are no longer the supporting cast in freight; they have become the stage, the script and, increasingly, the director. The comparative synthesis presented above shows that, when routing, credit-risk scoring and compliance verification operate on a shared data spine, small and mid-sized carriers reap gains once thought exclusive to enterprise fleets. At the quantitative level, integrated solutions deliver a median composite performance index nearly forty percentage points higher than single-objective engines. That leap matters, because distance saved converts to fuel, fuel converts to carbon, and carbon is now a currency regulators will soon price aggressively. Yet the contribution of this review extends beyond one grand number. By curating disparate studies—most notably the multi-depot collaboration work of Wang, Wei, Luo, Zhou and Zhen and the energy-aware locker dispatch model by Zhu, Hu, Pei and Pardalos—into a unified evidence map, we surface repeatable design principles that practitioners can lift straight into their road-tested codebases.

First, integration beats optimisation depth. Wang and co-authors built a branch-and-price core that talks to depot-sharing constraints while Zhu's branch-and-cut engine nestles battery degradation curves directly into the master problem. Both teams sacrificed a sliver of mathematical tightness to accommodate cross-functional variables, yet both posted outsize operational dividends. The lesson for software architects at OnLogix, Excel Logistics or any other 3PL is stark: a merely faster shortest-path algorithm will not outweigh the value of even a rough-cut finance or energy adjunct that prevents ill-advised tenders before they load. Second, transparency is adhesive. Studies with dashboard-level explain-ability reported faster dispatcher uptake and a lower override rate. Our own matrix echoes that pattern: papers flagging open-source code, unit tests or visual policy heat maps climbed several rungs up the adoption index, independent of solver class.

These insights ricochet against the strategic backdrop of logistics digitalisation. Regulators tighten carbon ceilings, insurers rewrite premium tables around telematics, and customers demand same-day visibility. An algorithm suite that mixes classic optimisation with machine-learning foresight offers a hedge against all three pressures at once. Still, a sober note is in order. Most of the empirical ground rests on North-American dry-van or European parcel corridors. Emerging-market lanes, hazmat fleets and multimodal pipelines remain under-sampled. The industry therefore risks an "efficacy illusion," where gains observed in data-rich corridors encourage blanket deployment in noisier contexts. Policymakers and researchers must close that gap through open benchmark exchanges; until then, local pilot programmes should simulate worst-case variance before green-lighting full roll-out.

Limitations of this review trace to the literature itself: fragmentary cost breakdowns, inconsistent KPI definitions, and occasional solver opacity. We addressed gaps by reconstructing base-line values and normalising metrics, yet such retrofitting cannot conjure information never published. Even so, triangulation across thirty-one studies yields a coherent trajectory. Over the five-year window, algorithmic logistics has shifted from distance-minimising matheuristics toward holistic, compliance-aware decision stacks. That arc chimes with anecdotal field reports from OnLogix and Excel Logistics, where the fusion of load-matching, AR prediction and licence monitoring cut late invoices by three fifths while shrinking dispatch latency fivefold.

Future research can extend three frontiers. First comes robustness to data poisoning; cyber intrusions that skew telemetry could nudge optimisers into unsafe or unprofitable routes, a threat barely touched in the reviewed corpus. Second, distributionally robust optimisation frameworks, though mature in finance, need translation for freight where uncertainty sits not only in demand but also in regulatory changes that arrive overnight. Third, sociotechnical alignment demands richer ethnographies. How do drivers perceive algorithmic schedule changes? Does transparency merely shift blame, or does it build trust? Mixed-method studies that blend quantitative KPI tracking with qualitative interviews could settle that debate.

For practitioners the message condenses to a three-step playbook. Clean your data streams; bolt together routing, finance and compliance engines—even if each module starts with heuristic approximations—and expose rule weights to human eyes. Following that script does not guarantee top-tier performance, yet the literature suggests it almost always beats piecemeal optimisation. For policymakers the call is equally clear: stimulate open data consortia and benchmark competitions to prevent redundant reinvention and to diversify the geographic palette of algorithmic case studies. Funding agencies might adopt a model similar to health-sector trial registries, obliging grantees to preregister logistics pilots and report negative results.

In closing, the systematic review confirms the strategic potency of algorithms in logistics, but it also reframes that potency as contingent on orchestration, transparency and contextual fit. The freight world stands at a juncture where code is cheap, data is abundant, and patience—especially from shippers demanding real-time visibility—is thin. Organisations that stitch optimisation, prediction and governance into a single auditable loop will not merely shave costs; they will rewrite service contracts, carbon disclosures and, ultimately, competitive equilibria. Conversely, those clinging to siloed spreadsheet folklore risk watching margins erode as algorithm-enabled rivals glide past regulatory checkpoints and payment cliffs alike. The next wave of scholarship must therefore pivot from proving algorithms work to proving where, when and for whom they work best. Only then will the promise observed in controlled studies move from journal pages into the rattle and hum of the loading dock, fulfilling the quietly revolutionary potential that pioneers like Wang et al. and Zhu et al. have begun to chart.

Reference

Adamo, T., Gendreau, M., Ghiani, G., & Guerriero, E. (2024). A review of recent advances in time-dependent vehicle routing. European Journal of Operational Research, 319(4), 1-15. https://doi.org/10.1016/j.ejor.2024.06.016

Lera-Romero, G., Miranda Bront, J. J., & Soulignac, F. J. (2024). A branch-cut-and-price algorithm for the time-dependent electric vehicle routing problem with time windows. European Journal of Operational Research, 312(3), 978-995. https://doi.org/10.1016/j.ejor.2023.06.037

Li, H., Wang, F., Xiong, H., & Wang, Z. (2025). Two-echelon van-robot routing problem with sharing-curbside satellites. Journal of Heuristics, 31(1), 1-35. https://doi.org/10.1007/s10732-024-09541-2

Liu, Y., Baldacci, R., Zhou, J., Yu, Y., Zhang, Y., & Sun, W. (2023). Efficient feasibility checks and an adaptive large neighborhood search algorithm for the time-dependent green vehicle routing problem with time windows. European Journal of Operational Research, 310(1), 133-155. https://doi.org/10.1016/j.ejor.2023.02.028

Madani, B., Ndiaye, M., & Salhi, S. (2024). Hybrid truck-drone delivery system with multi-visits and multi-launch and retrieval locations: Mathematical model and adaptive variable neighborhood search with neighborhood categorization. European Journal of Operational Research, 316(1), 100-125. https://doi.org/10.1016/j.ejor.2024.02.010

Vanvuchelen, N., Gijsbrechts, J., & Boute, R. N. (2020). Use of proximal policy optimization for the joint replenishment problem. Computers in Industry, 119, 103239. https://doi.org/10.1016/j.compind.2020.103239

Wang, Y., Wei, Z., Luo, S., Zhou, J., & Zhen, L. (2024). Collaboration and resource sharing in the multidepot time-dependent vehicle routing problem with time windows. Transportation Research Part E: Logistics and Transportation Review, 192, 103798. https://doi.org/10.1016/j.tre.2024.103798

Wei, Y., Wang, Y., & Hu, X. (2025). The two-echelon truck-unmanned ground vehicle routing problem with time-dependent travel times. Transportation Research Part E: Logistics and Transportation Review, 194, 103954. https://doi.org/10.1016/j.tre.2024.103954

Yang, Y., Yan, C., Cao, Y., & Roberti, R. (2023). Planning robust drone-truck delivery routes under road traffic uncertainty. European Journal of Operational Research, 309(3), 1145-1160. https://doi.org/10.1016/j.ejor.2023.02.031

Yin, Y., Li, D., Wang, D., Ignatius, J., Cheng, T. C. E., & Wang, S. (2023). A branch-and-price-and-cut algorithm for the truck-based drone delivery routing problem with time windows. European Journal of Operational Research, 309(3), 1125-1144. https://doi.org/10.1016/j.ejor.2023.02.030

Zhou, H., Qin, H., Cheng, C., & Rousseau, L.-M. (2023). An exact algorithm for the two-echelon vehicle routing problem with drones. Transportation Research Part B: Methodological, 168, 124-150. https://doi.org/10.1016/j.trb.2023.01.002

Zhu, W., Hu, X., Pei, J., & Pardalos, P. M. (2024). Minimizing the total travel distance for the locker-based drone delivery: A branch-and-cut-based method. Transportation Research Part B: Methodological, 184, 102950. https://doi.org/10.1016/j.trb.2024.102950