MACHINE LEARNING

Machine Learning-Based Route-to-Market Optimization for Smart Retail Distribution Networks in Saudi Arabia

Mohamed Saad Hassan Ismail

Published November 2026

Abstract

Retail distribution networks in Saudi Arabia must coordinate variable outlet demand, constrained transport capacity, retailer receiving windows and geographically diverse delivery conditions. This critical narrative review evaluates machine-learning-assisted route-to-market optimization in the context of Saudi Vision 2030 and the National Transport and Logistics Strategy. It synthesizes research on vehicle routing, neural combinatorial optimization, predictive demand modeling and learning-assisted local search, and proposes an integrated architecture connecting commercial data, predictive models, constraint-aware optimization and human-supervised dispatch. The review distinguishes operational feasibility from algorithmic benchmark performance and develops a framework for evaluating distance, on-time-in-full service, fleet utilization, cost, reliability and emissions. Two conceptual figures and comparison tables support the discussion; all numerical scenario values are explicitly illustrative rather than empirical findings. The evidence suggests that hybrid approaches are plausible candidates for deployment when prediction quality, governance, commercial service policies and optimization robustness are jointly addressed. Future research should prioritize Saudi-specific datasets, controlled evaluation, heterogeneous fleet constraints, multimodal interfaces and transparent reporting of uncertainty.

Keywords: machine learning; route-to-market; Saudi Arabia; Vision 2030; FMCG; vehicle routing; logistics platforms; demand forecasting; retail distribution

1 Introduction

Saudi retail distribution is a complex coordination problem rather than a simple task of moving goods from a depot to a customer. National and regional fast-moving consumer goods (FMCG) networks must reconcile the competing requirements of availability, delivery reliability, inventory productivity, service frequency, vehicle utilization and distribution cost. A route that minimizes distance can still perform poorly when it misses a retailer's receiving window, exposes perishable food to excessive dwell time or visits the wrong outlets at the wrong frequency. This paper examines how machine learning can improve route-to-market decisions when integrated with formal transport optimization and operational governance. Its setting is Saudi Arabia, where metropolitan concentration, widely dispersed settlements and diverse retailer formats create distinct planning environments.

The analysis is framed by Saudi Vision 2030 and the National Transport and Logistics Strategy, which place emphasis on logistical connectivity, operational efficiency, innovation and integration across transport services. These national ambitions do not, by themselves, demonstrate that a particular machine-learning technique will reduce costs in a specific distribution company. The relevant research question is operational: which predictions improve actionable decisions, under which

constraints, and against which credible conventional baselines? A meaningful answer requires distinguishing prediction quality from route quality, and route quality from wider commercial performance.

The term route-to-market denotes the complete commercial and physical pathway through which products reach retail outlets. It encompasses territory design, channel assignment, visit frequency, order capture, inventory replenishment, fleet scheduling and execution feedback. Vehicle routing is an important component, but it is not synonymous with the entire route-to-market system. A workable model therefore couples outlet-level demand information to decisions about delivery quantities, sequence, travel times and capacity. This broader interpretation is particularly relevant to sales-led distribution businesses where field sales teams and logistics operators share decisions but report against different targets.

Existing research in vehicle-routing problems (VRPs) has advanced from exact optimization and human-designed heuristics towards learning-assisted search and neural construction methods. Surveys by Bogyrbayeva et al. (2024) and Zhou et al. (2025) make clear that learning methods can be competitive on particular problem families, yet operational deployment remains difficult when constraints, data distributions and real-world disruptions differ from benchmark assumptions. The contribution of this review is to connect that technical literature to commercially relevant Saudi FMCG operations, identify a practical hybrid architecture and propose a defensible evaluation design without claiming to have conducted a company field trial.

2 Aim, objectives and research questions

The aim is to critically evaluate the potential of machine-learning-supported route-to-market optimization to improve efficiency, resilience and retailer service in Saudi Arabia's FMCG distribution networks. The first objective is to map route-to-market decisions and their data requirements. The second is to compare predictive models and optimization approaches, including their limitations. The third is to define measurable logistics, commercial and sustainability outcomes. The fourth is to propose a phased implementation architecture appropriate for data-mature and data-constrained organizations. The fifth is to identify Saudi-specific research gaps and governance requirements.

Four research questions structure the synthesis. RQ1 asks which features of Saudi retail distribution create the largest opportunities and obstacles for predictive routing. RQ2 asks how demand forecasting, arrival-time prediction and dynamic routing should interact. RQ3 asks whether learning-assisted systems can outperform carefully tuned conventional routing systems on a risk-adjusted, out-of-sample basis. RQ4 asks what organizational capabilities and data safeguards are needed before such systems can be responsibly deployed. These questions deliberately avoid assuming a positive causal impact before an evaluation is completed.

3 Review methodology

This is a structured critical narrative review supplemented by a conceptual system design and a transparent illustrative scenario. It is not presented as an exhaustive systematic review or meta-analysis. Searches were directed to scholarly publisher pages, recognized conference proceedings, institutional repositories and official Saudi government sources. Search concepts included vehicle routing, machine learning, deep reinforcement learning, attention models, learning-assisted local search, demand forecasting, last-mile freight, FMCG distribution, logistics platforms and Saudi Vision

2030. Searches emphasized work published between 2020 and 2026 while retaining selected foundational routing studies necessary to explain the methods.

Inclusion favored peer-reviewed surveys and methodological studies addressing combinatorial optimization, routing under operational constraints, logistics decision support or governmental policy framing. Foundational publications before 2020 were included only when they provided indispensable architectural or algorithmic context. Exclusions covered marketing claims with no methodological disclosure, sources with unclear provenance, unrelated passenger transport applications and papers whose performance claims could not be distinguished from their experimental assumptions. Where publisher abstracts were accessible but full experimental details were not, conclusions were stated cautiously.

The synthesis proceeded by mapping each source to a decision layer: forecasting, travel-time estimation, optimization, execution, evaluation or governance. Methodological quality was judged qualitatively by the transparency of benchmark datasets, comparators, constraint definitions, generalization tests and computation budgets. No numerical synthesis of effect sizes was attempted because different studies solve non-equivalent instances and report inconsistent outcome measures. The primary limitations are selection bias inherent in a narrative review, restricted access to some full texts, transferability from synthetic benchmarks and the absence of primary Saudi retail transaction or route telemetry data.

Table 1. Literature synthesis and methodological implications.

Source / approach	Main contribution	Practical qualification
Bengio et al. (2021)	ML–optimization integration	Learned components require careful coupling to solvers
Hottung & Tierney (2022)	Neural improvement search	Feasibility and comparison with strong baselines matter
Bogyrbayeva et al. (2024)	Taxonomy of learned VRP methods	Results vary by problem variant and dataset
Zhou et al. (2025)	Recent routing research directions	Generalization and operational deployment remain open
Saudi transport strategy	Policy alignment	Policy targets are not company-level effect sizes

4 Saudi route-to-market context

Saudi Arabia combines dense urban retail demand with long intercity corridors, spatially dispersed towns and variable outlet accessibility. Urban deliveries may face tight access windows, parking constraints and recurrent congestion, whereas remote routes may be limited by total travel time, sparse order density and low return-load opportunity. These are operational hypotheses to be tested with local data rather than universal characterizations of every delivery corridor. Model design should consequently distinguish urban van routes, regional truck replenishment and low-density rural service patterns instead of enforcing a single national optimization template.

FMCG distribution also contains cross-functional incentives that can undermine technically optimal routes. A sales manager may prefer frequent visits to protect shelf availability and relationships, while fleet managers prioritize consolidated drops and utilization. Retailers may place orders late or require deliveries during narrow receiving periods. Promotions can create temporary spikes that

historic averages underpredict. A network-wide optimizer should represent agreed service policies explicitly, because an apparently low-cost plan that violates commercial commitments is not executable.

The Vision 2030 framework calls for the Kingdom to become a global logistics hub, while the transport ministry identifies transport integration, innovation and business intelligence as strategic priorities. At the scale of a distribution company, these principles translate into interoperable data exchange, transparent operational measurements, resilient transport decisions and stronger connections between sales, warehousing and delivery functions. Multi-modal networks become relevant when intercity freight uses combinations of road, rail or port-linked movements; most daily retail outlet visits, however, remain a last-mile routing problem. The study therefore treats multimodal coordination as an adjacent research extension rather than pretending that every FMCG delivery uses multiple transport modes.

5 Core decision problem and mathematical representation

At its simplest, a capacitated VRP seeks vehicle routes from one or more depots to retail outlets while respecting vehicle capacities and minimizing a chosen objective. Real FMCG applications usually require a richer vehicle-routing problem with time windows, heterogeneous fleets, split deliveries, delivery priorities, uncertain demand and variable travel times. Network decisions may also incorporate outlet assignment to depots, visit-day selection and sales-territory policies. The scale and structure of these constraints materially influence which computational methods are realistic.

Let each outlet i have a predicted demand d_{*i*}, required service duration s_{*i*} and acceptable receiving window $[a_{*i*}, b_{*i*}]$. Let t_{*ij*} represent predicted time between nodes i and j, c_{*ij*} a generalized travel cost and x_{*ijk*} a binary variable indicating vehicle k traverses arc i to j. A baseline objective minimizes the sum of vehicle movement costs, lateness penalties, shortage penalties and vehicle activation costs. Feasible solutions satisfy visit, flow-conservation, vehicle-capacity and time-propagation constraints. A second objective may penalize estimated emissions, although emissions must be estimated using a documented physical or empirical model rather than distance alone.

An integrated solution may be expressed conceptually as

$$\min \sum_{i,j,k} c_{ij} x_{ijk} + \lambda_1 \sum_{i} \text{lateness}_i + \lambda_2 \sum_{i} \text{shortage}_i + \lambda_3 \sum_{k} \text{emissions}_k \tag{1}$$

The weights λ reflect organizational preferences and should not be chosen merely to produce favorable research results. Hard legal, safety and product-integrity constraints remain hard constraints rather than negotiable penalty terms. Prediction errors enter the optimization through d_{*i*} and t_{*ij*}; robust optimization or scenario analysis can protect against those errors when failure costs are asymmetric.

Commercial strategy adds a higher decision layer. Outlet segmentation, sales potential, service frequency and product-category replenishment rules determine which visits are required. These inputs should be governed jointly by sales, finance and logistics. It is misleading to attribute all improvement to a routing algorithm if the intervention also changes customer segmentation, reduces visit frequency or alters delivery promises. An evaluation should therefore document each concurrent policy change and report both operational and commercial consequences.

6 Machine learning in demand and service prediction

Demand forecasting supplies the quantities used by order planning and capacity constraints. Relevant inputs include lagged sales by outlet and stock-keeping unit, day-of-week effects, promotional calendars, holidays, outlet format, geographic region, product shelf-life and available stock. Model families range from simple seasonal baselines and gradient-boosted trees to sequence models. Complex architectures are not automatically preferable: intermittent demand, sparse stores and shifting promotional policies may make simple hierarchical approaches more stable and explainable. Predictions must be generated using only information genuinely available at the dispatch decision time.

Demand is frequently censored by stockouts: observed sales do not equal unconstrained consumer demand when shelves are empty. A model trained on uncorrected sales can learn to predict the consequences of past supply failures and perpetuate under-allocation. Data preparation therefore needs stock availability indicators, returns, cancellations, promotion flags and manual override records. Evaluation should separate ordinary days, promotion periods, new outlets and low-volume territories. A single national forecast accuracy metric can hide severe failure in the locations where service quality matters most.

Travel-time estimation converts geographic distance into estimated arrival times. Features may include departure time, corridor, stop density, vehicle class, observed route histories and recurring congestion patterns. Predicted travel-time distributions are more useful than isolated point estimates when service windows are tight. Quantile predictions, such as the 90th percentile travel time, allow the route planner to balance on-time performance against excess buffers. Reliable GPS traces and stop-event timestamps are essential, but telemetry can contain location errors, missing segments and driver-specific behavior.

Service duration also matters. A small convenience outlet may require a short unload, while a supermarket receiving dock can involve queuing, paperwork and a longer unloading sequence. Estimating stop service time from outlet type, delivery quantity and historical dwell time can improve feasibility. Human overrides should be captured as structured events so future learning distinguishes genuine operational constraints from informal habits. A controlled monitoring system needs to watch error drift over time, because forecast gains measured during development may decay after seasonal shifts or business restructuring.

7 Learning-assisted vehicle routing approaches

Conventional routing techniques include mixed-integer programming for smaller or decomposable instances, constructive heuristics, local search and metaheuristics. These methods provide an important benchmark because they explicitly encode constraints and often perform reliably in production. Learning can contribute in three ways: it can estimate uncertain inputs, choose promising actions within an optimizer or directly construct routes. Each contribution must be evaluated independently. A claim that a deep model is innovative does not establish that its final route is feasible, cost-effective or easily maintained.

Neural constructive models learn policies that select the next stop using representations of nodes and routing states. The attention-based architecture of Kool et al. (2019) illustrated how learned representations could address several routing variants. POMO (Kwon et al., 2020) exploited multiple equivalent optimal solutions to stabilize policy learning and search. These methods are scientifically

significant but often evaluated on synthetic problem distributions. Generalizing from a standard-capacity benchmark to Saudi operations with multiple depots, different vehicle types and retailer receiving constraints requires careful revalidation.

Learning-assisted improvement methods start with a feasible route and iteratively modify it. Hottung and Tierney (2022) describe neural large-neighborhood search, in which learned components guide modifications of candidate solutions. This hybrid direction is especially attractive to logistics practitioners because a conventional solver can retain responsibility for hard constraint enforcement while learning focuses computational effort where improvements may be found. The broader reviews of Bengio et al. (2021), Bogyrbayeva et al. (2024) and Zhou et al. (2025) support the importance of integrating machine-learning methods with established operations research techniques.

Reinforcement learning frames routing as sequential decision-making, in which an agent is rewarded for building routes with desirable outcomes. The reward specification is critical. Rewarding only short distance can produce late deliveries, overloaded vehicles or systematic neglect of remote outlets unless feasibility is enforced. Offline training on historical data may also reproduce legacy biases. For a first Saudi FMCG deployment, a defensible architecture uses supervised predictive models, a constraint-aware optimization engine and human dispatch approval before exploring end-to-end reinforcement-learning dispatch.

8 Proposed hybrid architecture for smart distribution networks

The proposed architecture consists of six linked layers. The first is a controlled data foundation joining orders, products, stock, retailers, depots, vehicle capabilities, approved geospatial coordinates and event timestamps. The second generates forecasts for demand, travel time, unload duration and exceptions. The third represents service policies and validates mandatory constraints. The fourth solves routing and territory-assignment problems using a conventional optimization engine informed by learned estimates. The fifth provides dispatch oversight and execution visibility. The sixth collects outcomes for monitoring, feedback and model recalibration.

Integration should be implemented through well-defined interfaces rather than a monolithic AI application. The sales system should remain the source of truth for retailer orders and service agreements, while warehouse and transport systems control inventory and dispatch execution. The optimizer should receive a versioned snapshot with clear timestamps and return the proposed routes, expected arrival windows, predicted cost and reasons for infeasibility. Planners need the ability to approve, adjust or reject a plan, and every override should be logged for audit and later analysis.

A sensible operational workflow begins with an overnight plan using confirmed orders and probabilistic next-day demand. Before loading, the system reconciles vehicle capacity and available stock. During delivery, new events may trigger selective re-optimization, but not every delayed vehicle warrants network-wide rerouting. Continuous optimization can destabilize driver schedules and customer commitments if changes are too frequent. A threshold-based policy should distinguish minor deviations from severe disruptions, while preserving a clear single version of the approved route.

The architecture should include a fallback mode. If a prediction service fails, the system must still calculate a feasible route from conventional travel-time assumptions and predefined service policies. If geolocation quality is inadequate, dispatchers should be warned rather than shown spurious

precision. Model transparency can be achieved with feature attribution for predictors, constraint diagnostics for optimizers and side-by-side cost-service trade-offs for route alternatives. This is more actionable to management than treating a neural-network score as a complete explanation.

The practical data model should identify stable entities and their relationships before selecting technology. Retailer master records need a unique outlet identifier, approved geographic point, receiving-window agreement and channel classification. Order lines need product identifiers, quantities, planned delivery date, cancellation reason and confirmation timestamp. Vehicle records should include legal payload capacity, operating cost, product suitability and actual availability. Driver and vehicle identifiers should be separated from analytic datasets where direct identification is unnecessary. A central event dictionary is needed so a recorded arrival means the same operational milestone across depots. Without these foundations, apparently advanced forecasts can become elaborate transformations of inconsistent historical records.

An example daily optimization cycle illustrates the value of that information architecture. At an agreed order cutoff, the system captures accepted retailer orders and reconciles available stock. Expected demand is used primarily for capacity buffers and next-cycle positioning, not to quietly substitute predicted orders for contractual orders. The system estimates service durations and time-dependent travel, then optimizes assignments and visit sequences. Dispatchers review capacity, time-window violations and unusual changes from established routes. Approved plans become execution records in the transport management system. After delivery, actual arrival, unloading, rejections and shortages are recorded and linked to the specific planning snapshot. This traceability allows retrospective assessment of which predicted features or route decisions contributed to a failure.

A secondary objective is organizational learning. If dispatchers frequently reject routes serving a particular district, that pattern may indicate missing road restrictions, unrealistic retailer access windows or a legitimate commercial policy. It may also indicate reluctance to change routines. The system should not automatically treat every override as ground truth; instead, it should classify reasons, verify recurring constraints and expose disagreements to joint review. Model retraining schedules should be event-driven where possible, particularly after territory redesign, a new distribution centre, an ERP migration or a significant policy change. Recalibration should include rechecking the conventional baseline rather than comparing only with an obsolete manual process.

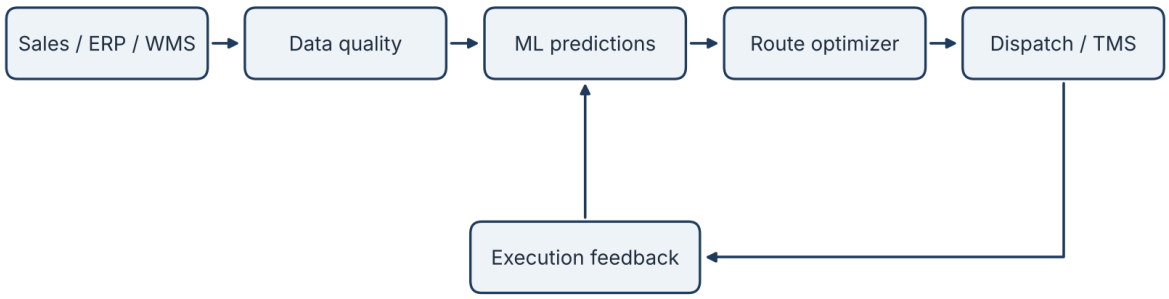


Figure 1. Proposed predictive and constraint-aware retail logistics workflow.

9 Evaluation framework and illustrative scenarios

A credible research design compares at least three approaches: the existing manual or rule-based planning process; a well-tuned conventional optimization baseline; and the proposed machine-learning-assisted optimizer. The same days, orders, vehicle capacities, depot locations, receiving windows and constraint definitions must be used in matched tests. Historical replay can identify feasibility and forecast problems, but a live controlled pilot is needed to assess driver behavior, retailer satisfaction and actual cost. Randomized territory-period assignment or a carefully designed stepped-wedge rollout may reduce the influence of seasonality and operational change.

Core metrics include kilometers per delivered case, on-time-in-full delivery percentage, vehicle utilization, cost per stop, route planning time, failed deliveries, retailer stockout exposure and estimated greenhouse-gas emissions. Financial measures should account for software cost, data maintenance, dispatch labor and exception handling, rather than assuming that shorter routes equal net profit. Operational and commercial metrics need equal attention; cutting service frequency may reduce kilometers while damaging sales. Disaggregated analysis should cover urban, suburban and low-density territories, fleet type and outlet segment.

Table 2 and Figure 2 deliberately present a hypothetical planning example, not observational results. The illustrative values represent a possible pattern in which manual planning is compared with a traditional solver and a hybrid predictive system. They are included to explain metric reporting and sensitivity analysis only. No simulation was executed and no real-world performance improvement is claimed. In a publishable empirical study, these values must be replaced with audited observations or results from a fully documented reproducible computational experiment.

Evaluation must include uncertainty and robustness. Bootstrap confidence intervals can be used for mean route-cost differences when sampling supports them; repeated trials and matched calendar blocks can address day-to-day variation. Pre-specified subgroup tests can identify whether gains disappear during peaks, promotions or remote deliveries. Sensitivity analyses should vary demand error, travel-time uncertainty, fuel price, vehicle availability and penalties for missed windows. A system that wins only under its favored assumptions is not operationally robust.

A detailed pilot protocol should freeze the candidate models and comparator settings before performance measurement. The pilot population should include representative outlet classes and route densities, with exclusion criteria documented in advance. Where randomized territories are impractical, matched territories can be selected based on historical stop counts, distances, weekly demand, retailer format and service windows. A pre-intervention observation period is important because planners may need time to correct location records before any algorithm can perform fairly. The primary outcome should be specified explicitly, for example total controlled distribution cost subject to an on-time-in-full floor, instead of selecting whichever metric improves most after observing the data.

Model ablation tests can reveal which elements create measurable value. One configuration should retain the conventional route engine with historical-average travel times; a second should add supervised travel-time estimates; a third should also add demand forecasts; and a fourth may incorporate learning-assisted search. These comparisons make it possible to separate predictive benefit from solver improvement. They also protect against a common evaluation mistake: comparing a sophisticated system using clean, current data with a manual process using old and incomplete

data. Each experimental condition should receive the same data cleansing, product rules and retailer constraints, with computation times and hardware stated.

Commercial impact requires careful causal interpretation. On-time deliveries may improve without increasing product sales if shelves are already well stocked; conversely, sales may rise because of a concurrent promotion rather than a routing intervention. Measuring retailer inventory availability, returns, lost delivery opportunities and service complaints can help interpret pathways. Costs should be measured as direct operational accounting quantities where possible, including fuel, overtime, vehicle rental, failed re-delivery and incremental system expense. Estimated emissions should use declared fleet and fuel assumptions. Where emissions and cost objectives conflict, report the trade-off transparently instead of collapsing it into a single favorable number.

Table 2. Illustrative scenario (hypothetical values only).

Indicator	Manual planning	Conventional solver	Hybrid ML + solver
Relative vehicle-kilometres	100	88	82
On-time delivery (%)	90	94	96
Vehicle utilization (%)	68	75	79
Planning time (minutes)	120	35	25

Note: All values in Table 2 are invented solely to demonstrate how pilot outcomes might be displayed. They are not measurements or outputs from an implemented model.

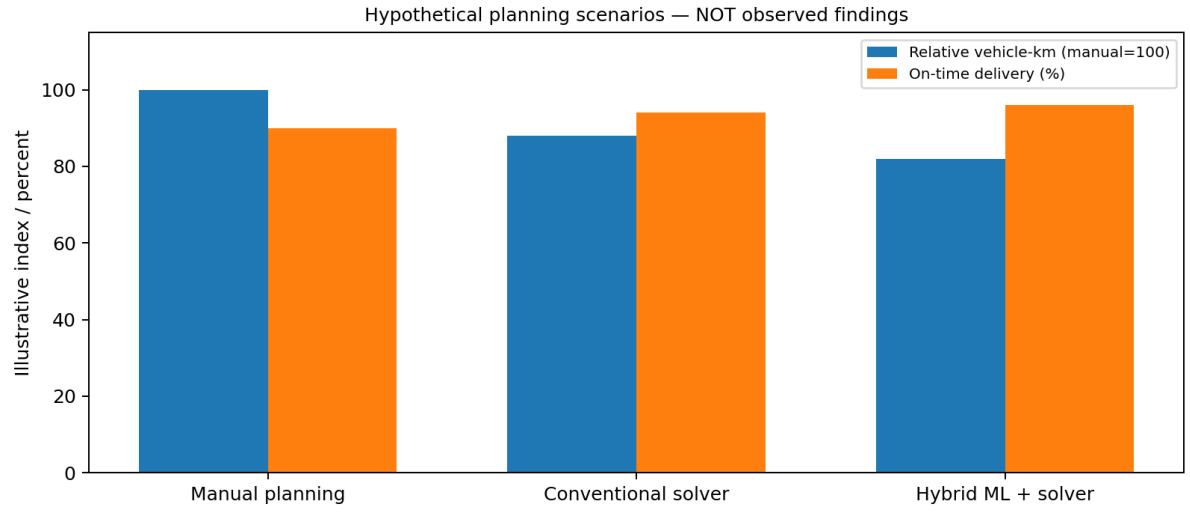


Figure 2. Illustrative scenario comparison (hypothetical, not experimental or observed data).

10 Critical discussion

The most compelling value proposition is not the replacement of route planners by autonomous algorithms, but the combination of better predictions, explicit constraints and rapid alternative generation. Machine learning estimates where uncertainty is likely; optimization turns those estimates into feasible resource assignments; and dispatch staff evaluate operational credibility. In the Saudi setting, this separation can help organizations modernize progressively without abandoning useful staff knowledge or existing transport systems. The design also avoids attributing all operational improvement to a single algorithmic component.

Nevertheless, major limitations remain. Historical routes are selected decisions rather than random observations: travel times and demand patterns reflect how companies previously served customers. Training models on these records introduces selection effects. Low-volume outlets may be underrepresented, recent geographic expansions lack sufficient history, and historic data may not contain the disruptions most relevant to resilience. Neural approaches can additionally be sensitive to changes in instance size and topology. These issues justify conservative deployment thresholds and strong comparisons against conventional algorithms.

There is also a governance risk when objectives are set without stakeholder participation. An optimizer that prioritizes the cheapest outlets may repeatedly deprioritize rural retailers or smaller customers. Explicit service floors, territorial equity constraints and management review are therefore appropriate. Likewise, driver telemetry and customer information should be collected for defined purposes under applicable Saudi personal-data governance. The Saudi Data and Artificial Intelligence Authority provides policies and regulations relevant to personal data protection and data management; legal review remains necessary before operating a specific processing arrangement.

Logistics platforms are the mechanism for connecting these decisions. A platform creates value when it reliably integrates retailer demand, warehouse capability, transport constraints and execution events. Multimodal integration is a possible longer-term development: regional replenishment through ports or rail-connected hubs could pass inventory to road distribution fleets. As there is insufficient evidence of a deployed multimodal platform, this paper treats that connection as a future research direction rather than an existing organizational capability.

11 Research gaps and future research agenda

First, Saudi-specific, publicly documented benchmark datasets remain a priority. Studies should describe anonymized route geography, receiving constraints, heterogeneous vehicles, delivery volumes and uncertainty in ways that allow fair comparisons without disclosing commercially sensitive retailer data. A reusable benchmark suite could include separate urban and intercity scenarios, consistent units and standard feasibility checks. Without such benchmarks, reported gains from different models cannot be compared meaningfully.

Second, researchers should jointly optimize commercial and logistics objectives rather than minimizing fleet distance alone. This could involve service frequency, expected lost sales from stockouts, customer importance, expiration risk and long-term retailer relationships. Trade-offs should be reported with a Pareto frontier rather than hidden inside unexplained penalty weights. Third, models should be tested under distribution shift: festival demand, unexpected closures, fuel price changes and expansion to new areas. Robustness under such conditions can be more valuable than small improvements on average test days.

Fourth, explainability and operational trust require systematic evaluation. A useful research design would assess whether dispatchers understand why the system changed a route, whether they overrule sensible recommendations and whether explanations change adoption rates. Fifth, sustainable logistics calls for more precise emissions modeling that accounts for vehicle type, load, congestion and idling. Future studies can investigate electric-vehicle readiness and low-emission distribution zones, but such scenarios should not be presented as current universal operational capabilities.

Sixth, the effect of multimodal freight integration should be tested at the network-design level. Long-haul inbound movements, depot selection and last-mile trips operate over different timescales and

may require hierarchical models. A research program could evaluate whether better forecasts at distribution depots improve replenishment scheduling for rail or port-linked freight. Such work is especially aligned with the Kingdom's logistics connectivity agenda, provided it is based on documented transport choices rather than assumed multimodal activities.

An implementable research program could proceed in three stages. Stage one establishes an anonymized Saudi retail routing benchmark with urban and regional scenarios and deliberately stressful cases. Stage two tests forecasting, optimization and hybrid combinations under a reproducible compute budget and with identical feasibility rules. Stage three evaluates the strongest candidate through a managed real-world trial with retailer-service safeguards. Each stage has a clear failure condition: insufficient data reliability, lack of out-of-sample improvement or unacceptable operational disruption. This staged approach is important because a technical paper can report attractive benchmark performance while leaving unanswered questions about deployment maintenance, exception management and adoption.

Policy-relevant evaluation should also consider who benefits from route consolidation. Large retailers can receive more reliable service while small stores experience longer replenishment intervals if objectives are not balanced. A minimum service constraint can limit this effect, but it may reduce achievable distance savings. Rural access and low-frequency routes deserve explicit sensitivity testing. Similarly, vehicle optimization can reduce unnecessary kilometers while increasing delivery concentration at particular hours; the latter could move congestion to loading docks or neighborhoods. A broad assessment must therefore consider both fleet efficiency and impacts on service patterns.

The methodological contribution of future Saudi field studies could be stronger if researchers publish negative results and boundary conditions. Examples include districts in which learned demand estimates are too unreliable, situations where a conventional heuristic is computationally sufficient or promotion periods when manual knowledge remains essential. Such findings would identify deployment limits and prevent wasteful investments. Transparent reporting of training leakage, changing test populations and statistical uncertainty would also make cross-company results more credible. The goal is not to prove that artificial intelligence always wins, but to find where its incremental value justifies its operational complexity.

12 Practical, managerial and policy implications

For distribution executives, the first decision is readiness rather than model selection. Auditing master data, outlet coordinates, actual stop times and service-level definitions is often a better initial investment than acquiring a highly complex model. Management should identify a small number of territories with stable event capture and realistic implementation support. A phased pilot can then isolate data defects, quantify achievable benefits and establish the circumstances in which a conventional optimizer is already sufficient.

For sales leadership, route-to-market optimization must protect the customer proposition. Service frequency and visit timing affect availability, promotion execution and retailer relationships. Joint governance between sales and transport functions is needed to define the trade-offs and authorize policy changes. For finance, a credible business case includes recurring licensing and maintenance costs, organizational training, exception handling and potential benefits from fewer failed deliveries and better utilization. For technology teams, monitoring, security and rollback capability are mandatory rather than optional enhancements.

For public policy and sector development, interoperable logistics data standards, workforce skills and responsible AI governance are practical enablers of wider efficiency improvements. Saudi Vision 2030 provides a strategic direction, but individual digitalization projects require transparent implementation metrics and sector-specific evidence. Collaboration among retailers, distributors, universities and logistics providers could produce privacy-preserving benchmark data and shared evaluation protocols while allowing firms to protect commercially sensitive records.

A final managerial consideration is the difference between a research-ready prototype and a maintainable service. The first may operate on cleaned historical extracts with researcher supervision; the second must manage late orders, missing addresses, new drivers, changing vehicle availability and exception escalation every day. Production readiness should therefore be measured with operational reliability indicators, including system uptime, percentage of routes requiring manual correction, prediction-data latency, override frequency and time to recover from an outage. Responsibility for these indicators should be assigned before rollout. Contractual arrangements with software vendors should specify data ownership, interoperability, retraining access and the capacity to export decisions for independent auditing. A vendor-provided demonstration is not evidence of sustainable return on investment. Organizations should instead require representative pilot evidence, documented implementation costs and a credible exit strategy. These provisions make intelligent logistics investments more transparent and help avoid dependence on opaque algorithms that cannot be evaluated against the firm's changing service requirements. Moreover, leadership should treat employees as participants in research design: sales teams can explain outlet expectations, drivers can identify access constraints and warehouse supervisors can flag loading dependencies. Incorporating these perspectives early can improve problem formulation and acceptance, even before an advanced model has been trained. In that sense, organizational coordination is a necessary complement to machine-learning sophistication, not an ancillary consideration.

Table 3. Implementation roadmap and governance gates.

Phase	Deliverable	Go/no-go condition
0 — Data audit	Clean outlet, order and route records	Adequate coverage, lawful data processing
1 — Baseline	Conventional constrained routing	Feasible routes, measured existing performance
2 — Forecasting	Demand and travel-time predictors	Out-of-sample gain and drift monitoring
3 — Pilot	Supervised hybrid dispatch	No material service degradation
4 — Scale	Integration, training and governance	Replicated net benefit after full costs

13 Limitations

This article is a critical conceptual review, not a report of experiments, a systematic census of all relevant studies. The included technical literature varies considerably in problem formulation, computation budget and outcome reporting. Findings from standard benchmark instances do not directly establish real-world effect sizes in Saudi FMCG networks. Furthermore, the scope emphasizes operational distribution and may not capture all commercial incentives, formal logistics contracts or regulatory constraints that differ by product and region.

The visual architecture and numerical comparison are authored research illustrations. They are not reproductions of audited company data, calibrated forecasts or demonstrated improvements. The article does not claim that the proposed framework has been deployed. Empirical validation would

require permission to use a defined dataset, documentation of model training, registered evaluation criteria and transparent reporting of failures as well as successes.

14 Conclusion

Machine-learning-based route-to-market optimization is a promising but conditional approach to smarter Saudi retail distribution. The strongest implementation path combines demand, travel-time and service-duration prediction with constraint-aware vehicle-routing optimization, operational feedback and human governance. This hybrid formulation is more realistic than relying on neural route construction alone because it preserves hard feasibility checks and supports incremental deployment. Its likely value depends on data completeness, service-policy alignment, uncertainty management and credible performance evaluation.

The proposed research agenda aligns conceptually with Saudi Vision 2030 ambitions for logistics modernization and connectivity while retaining an appropriately narrow evidence claim. Future research should prioritize Saudi-specific benchmark data, controlled field comparisons, commercial outcome measures, climate and disruption robustness, explainable recommendations and genuine intermodal extensions. Until such evidence is generated, performance charts should remain explicitly illustrative and conclusions should be framed as testable propositions rather than verified cost savings.

References

[1] A brief survey on learning based methods for vehicle routing problems (2023). Procedia Computer Science, 221, 773–780. https://doi.org/10.1016/j.procs.2023.08.050

[2] Bengio Y, Lodi A, Prouvost A (2021). Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, 290(2), 405–421. https://doi.org/10.1016/j.ejor.2020.07.063

[3] Bogyrbayeva A, Meraliyev M, Mustakhov T, Dauletbayev B (2024). Machine learning to solve vehicle routing problems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 25(6), 4754–4772. https://doi.org/10.1109/TITS.2023.3334976

[4] Hottung A, Tierney K (2022). Neural large neighborhood search for routing problems. *Artificial Intelligence*. https://doi.org/10.1016/j.artint.2022.103786

[5] Kingdom of Saudi Arabia (n.d.). Saudi Vision 2030: Overview. https://www.vision2030.gov.sa/en/overview

[6] Kingdom of Saudi Arabia (n.d.). Vision 2030 Annual Report 2025. https://www.vision2030.gov.sa/media/ecdjfopq/vision2030_annual_report_2025_en.pdf

[7] Kool W, van Hoof H, Welling M (2019). Attention, learn to solve routing problems!. *International Conference on Learning Representations*; 2019. https://openreview.net/pdf?id=ByxBFsRqYm

[8] Kwon YD, Choo J, Kim B, Yoon I, Gwon Y, Min S (2020). POMO: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural Information Processing Systems* 33; 2020. https://proceedings.neurips.cc/paper_files/paper/2020/hash/f231f2107df69eab0a3862d50018a9b2-Abstract.html

[9] Li J (2026). Learning-based optimization for vehicle/robot routing problems: A survey. *Journal of Artificial Intelligence and Applications*; 2026. https://www.sciltp.com/journals/jaia/articles/2605004025

[10] Ministry of Transport and Logistics Services (Saudi Arabia) (n.d.). About the ministry: Objectives and roles. https://mot.gov.sa/en/web/guest/about-motls

[11] Ministry of Transport and Logistics Services (Saudi Arabia) (n.d.). National Transport and Logistics Strategy. https://mot.gov.sa/en/ntls

[12] Nazari M, Oroojlooy A, Snyder LV, Takáč M (2018). Reinforcement learning for solving the vehicle routing problem. *Advances in Neural Information Processing Systems* 31; 2018. https://papers.neurips.cc/paper_files/paper/2018/hash/9fb4651c05b2ed70fba5afe0b039a550-Abstract.html

[13] Saudi Data and Artificial Intelligence Authority (SDAIA) (n.d.). Laws and regulations: Personal Data Protection Law and data policies. https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx

[14] The vehicle routing problem in the last decade: Variants, taxonomy and metaheuristics (2023). Procedia Computer Science, 220, 398–404. https://doi.org/10.1016/j.procs.2023.03.051

[15] Zhou F, Lischka A, Kulcsár B, Wu J, Chehreghani MH, Laporte G (2025). Learning for routing: A guided review of recent developments and future directions. *Transportation Research Part E*, 202, 104278. https://doi.org/10.1016/j.tre.2025.104278