Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Optimized Logistics Modeling to Achieve Agile Manufacturing: An Inbound Supply Chain Framework

**G. Y. Mythili**
Vellore Institute of Technology
**Sudha C**
sudha.c@vit.ac.in

Vellore Institute of Technology

**Research Article**

**Keywords:** inbound logistics, agile manufacturing, rectilinear facility location, ALDEP, warehouse layout, packaging heuristic, reinforcement learning, NSGA-II, digital twin

**Posted Date:** November 25th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11361666/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Optimized Logistics Modeling to Achieve Agile Manufacturing: An Inbound Supply Chain Framework

G. Y. Mythili[1] and Sudha C[2*]

[1]School of Advanced Sciences, Vellore Institute of Technology, Chennai, India.

[2*]School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, India.

*Corresponding author(s). E-mail(s): sudha.c@vit.ac.in;
Contributing authors: mythili.gy@vit.ac.in;

**Abstract**

Inbound logistics in manufacturing is often fragmented across suppliers, third-party logistics (3PL) providers and intermediate warehouses, which raises handling cost and lead time, increases damage risk and limits visibility. This paper studies the inbound material flow of Precision Pressform Industries (PPF), a clutch manufacturing plant, and develops an integrated framework to support agile manufacturing. A classical baseline is built from three tools: a rectilinear minisum model to locate a consolidation hub, the Automated Layout Design Program (ALDEP) to arrange a seven-department, 25,000 sq ft facility, and a volume-based packaging heuristic to raise truck utilization. Of seven ALDEP iterations, iteration 6 gave the best layout, with a closeness score of 167 and a double-sweep score of 334. On this baseline, the paper proposes a benchmark of learning-based optimizers (genetic algorithm, simulated annealing with machine learning, reinforcement learning, deep Q-learning and graph neural networks) and a hybrid LSTM and NSGA-II framework that couples demand forecasting with multi-objective layout decisions, to be validated by simulation and statistical testing. The study addresses the static nature of classical models, limited real-time data integration and the need for explainable optimization.

**Keywords:** inbound logistics, agile manufacturing, rectilinear facility location, ALDEP, warehouse layout, packaging heuristic, reinforcement learning, NSGA-II, digital twin

# 1 Introduction

Logistics decisions such as where to locate a facility, how to lay it out and how to load a truck are usually made with analytical optimization. The field began with the truck dispatching problem Dantzig and Ramser [1959] and has since grown into a large family of routing models, summarized in Toth and Vigo [2014] and Laporte [2009]. Most of these models share an assumption: demand, capacity and travel conditions are known in advance and stay the same while the plan is carried out. That assumption fits a stable network. It does not fit a factory whose supply mix, order volumes and delivery schedules change from month to month.

Two research streams have emerged in response. The first applies learning-based methods to logistics decisions. A recent review of reinforcement learning in logistics and supply chain management finds Q-learning to be the most commonly used approach, and a rising share of work on urban logistics Yan et al. [2022]. The second connects the physical operation to a live digital model. Digital-twin research for smart manufacturing has been reviewed in terms of its reference models, application scenarios and open issues Lu et al. [2020]. Both streams are closely tied to agile manufacturing, the ability of a firm to respond quickly to disruption. Industry 4.0 technologies are widely discussed as enablers of that ability Agarwal et al. [2023].

This paper studies a case where these ideas meet a practical problem: the inbound logistics of Precision Pressform Industries (PPF), a clutch manufacturing plant. Material currently moves from suppliers to a third-party logistics (3PL) provider, then to a warehouse, and finally to the plant. This arrangement causes six recurring problems:

- Material is handled several times, which raises cost and damage risk.
- Routing through the 3PL lengthens lead time.
- Direct and warehouse-based supply are mixed without a standard rule.
- Suppliers, the 3PL and the plant have limited visibility of each other's status.
- The same stock is stored in more than one place.
- Manual verification slows receiving.

Together these amount to one core issue: the flow depends too heavily on the 3PL and is fragmented, which lowers efficiency and raises cost.

Established facility-planning tools can address much of this. A location model can place a consolidation point close to the heaviest demand. A computerized layout method such as ALDEP Seehof and Evans [1967] can turn qualitative closeness ratings between departments into a floor plan. Packing and truck-matching methods, such as the three-dimensional pallet loading approach in Chen et al. [2024], can raise the load carried per trip. These tools are transparent and easy to defend to management, but they are static. They produce one answer for one demand snapshot. Learning-based methods may adapt better, but their gain over a well-built classical baseline at a real plant has not been shown. This paper therefore asks two questions: how far can a classical, well-documented toolset improve PPF's inbound flow, and what would an adaptive, learning-based extension need in order to improve on it?

The contributions of this paper are as follows:

1. A classical baseline for PPF that combines rectilinear warehouse location, ALDEP layout planning over seven departments in 25,000 sq ft, and a volume-based packaging heuristic for truck utilization.
2. A benchmark design that compares learning-based layout optimizers (GA, SA with ML, RL, DQN and GNN) on a 12-department warehouse case, with defined evaluation criteria.
3. A proposed hybrid LSTM and NSGA-II framework for demand-aware, multi-objective warehouse decisions, together with a validation plan.

The rest of the paper is organized as follows. Section 2 reviews related work, Section 3 identifies the research gaps and states the problem, Section 4 describes the methodology, Section 5 presents the PPF implementation and results, and Section 6 concludes with limitations and future work.

# 2 Related works

## 2.1 Classical location, routing and layout models

Quantitative logistics planning began with the truck dispatching problem Dantzig and Ramser [1959], and fifty years of later work on vehicle routing are consolidated in Toth and Vigo [2014] and Laporte [2009]. Those surveys show a steady growth from the basic capacitated problem to variants with time windows and multiple depots. For large instances, heuristics and metaheuristics now do most of the work. Facility siting follows a parallel line. The rectilinear, or minisum, model places a facility so that demand-weighted travel along orthogonal aisles or roads is smallest, and it is treated in standard location texts Love et al. [1988]. Inside the building, the layout problem is handled by relationship-based methods. ALDEP builds a layout by placing departments in sweeps and scoring how well the result satisfies closeness ratings Seehof and Evans [1967], and the textbook treatment of these methods is in Tompkins et al. [2010]. The common strength of all these tools is transparency: a planner can trace why a location or layout was chosen. Their common weakness is that each answer holds for one fixed set of inputs.

## 2.2 Metaheuristics and sustainability-oriented operations research

Metaheuristics relax the need for exact solutions. Ant colony optimization is a well-known example: a population of artificial agents shares information about good tours, which helps balance exploration of new routes against exploitation of known ones Dorigo and Gambardella [1997]. Such methods need careful parameter tuning, can be computationally expensive, and are usually applied to problems whose data do not change during the search. A broader view of operations research for green logistics is given in Dekker et al. [2012], which covers transport efficiency and the handling of returned and recycled material. Together, these works show how the toolbox grew beyond cost minimization, but they say little about how a plant should adapt its flows after the plan is made.

## 2.3 Learning-based optimization

Reinforcement learning (RL) treats a logistics decision as a sequence in which an agent learns from feedback. A review of RL in logistics and supply chain management groups the literature by application and reports that Q-learning is the most widely used approach, with growing activity in urban logistics Yan et al. [2022]. Deep variants extend this to larger decision spaces. The price of this flexibility is the need for substantial training data or a faithful simulator, and the learned policies are harder to explain than the classical formulas of Section 2.1.

## 2.4 Digital twins and Industry 4.0

A digital twin keeps a virtual model synchronized with a physical system. A review of twin-driven smart manufacturing sets out a reference model, surveys application scenarios under Industry 4.0, and lists open research issues Lu et al. [2020]. In logistics, a recent study uses a shop-floor twin as the training environment for a deep RL agent that dispatches and routes automated guided vehicles, and reports lower tardiness and energy use than four comparison methods Zhang et al. [2024]. This shows how a twin can make learning-based control safer to develop. It also shows the dependency: results rest on a high-fidelity model of the site, which a small or mid-sized plant may not have. These technologies are closely tied to agility, the capacity of a manufacturer to respond to disruption, and the role of Industry 4.0 in that capacity is discussed in Agarwal et al. [2023].

## 2.5 Loading and truck utilization

Transport cost depends not only on routes but on how well each trip is filled. One study formulates truck–cargo matching for a less-than-truckload hub as a two-stage problem, packing goods onto pallets in three dimensions and then assigning loaded pallets to trucks Chen et al. [2024]; it reports pallet loading efficiency close to 90% and truck utilization of nearly 95% on a real hub. A companion study treats the online version of the problem, where orders arrive in real time, and reports average loading efficiency close to 95% Tang et al. [2024]. These methods target hubs serving many customers, whereas inbound supply to a single plant has a different structure.

## 2.6 Observations

Each reviewed study solves one sub-problem well: routing, location, layout, loading or control. None of them combines location, layout and truck utilization for the inbound flow of a single manufacturing plant, and the learning-based studies rarely compare against a documented classical baseline from the same site. These observations lead directly to the research gaps and problem statement in Section 3.

# 3 Research gaps, problem statement and objectives

## 3.1 Research Gaps

The review in Section 2 shows that individual methods are mature, but five gaps remain between that literature and the needs of a working plant.

1. **No shared benchmarks:** Public datasets for dynamic and stochastic logistics problems are scarce, so reported improvements are hard to reproduce or compare across studies.
2. **Little attention to people:** Few models account for labour constraints, driver fatigue or the social side of sustainability, although these affect whether a plan can be run day to day.
3. **Limited real-time integration:** Optimization that keeps pace with uncertain, streaming data remains difficult to scale beyond pilots.
4. **Weak explainability:** Learning-based optimizers rarely come with interpretable reasoning, which makes it hard for managers to trust or audit their decisions.
5. **Little proof from live systems:** Empirical evidence that digital twins improve operations in running logistics systems is scarce.

Section 2 also points to a sixth gap that is specific to this work: location, layout and truck loading are normally studied separately, and learning-based methods are rarely compared with a documented classical baseline measured at the same site.

## 3.2 Problem statement

PPF, a clutch manufacturing plant, receives most of its inbound material through a chain of supplier, third-party logistics (3PL) provider, warehouse and plant. This arrangement produces six recurring problems: repeated handling that raises cost and damage risk; indirect routing that lengthens lead time; a non-standard mix of direct and warehouse-routed supply; limited visibility between stakeholders; duplicate storage that raises carrying cost; and manual verification that creates bottlenecks. The core issue is an over-dependence on the 3PL and a fragmented material flow, which together reduce efficiency and raise cost.

The current state is therefore $\textit{Suppliers} \rightarrow \textit{3PL} \rightarrow \textit{Warehouse} \rightarrow \textit{Plant}$. The target state is a standardized flow from suppliers through a single consolidation hub located close to the plant, with full truck loads and fixed dispatch rhythm, replacing the 3PL chain. The problem is to design that flow, to decide where the hub should be and how its interior should be arranged, and to ask whether adaptive, learning-based optimization can improve further on the classical design.

## 3.3 Research objectives

The overall aim is to analyse the inbound logistics of the clutch manufacturing plant and to design an optimized material flow that reduces cost, improves efficiency and increases supply chain visibility. The specific objectives are:

- **O1.** Assess the current supplier → 3PL → plant flow and identify the inefficiencies in handling, routing and storage.
- **O2.** Identify the main bottlenecks in transportation, unloading and verification, and study the effect of multiple handling and warehouse dependency.
- **O3.** Locate a consolidation hub by the rectilinear method so that demand-weighted transport distance is minimized.
- **O4.** Obtain a warehouse layout with ALDEP that minimizes material-handling effort between departments.
- **O5.** Set packaging and loading rules that raise truck utilization and reduce trips.
- **O6.** Define a benchmark and an evaluation plan for learning-based optimizers, and propose a hybrid LSTM + NSGA-II framework for demand-aware, multi-objective decisions.

The expected outcome is a streamlined, cost-efficient and responsive inbound logistics system with reduced handling, faster delivery and better inventory control.

# 4 Methodology

## 4.1 Overview

The study follows four linked phases (Fig. 1). Phase I applies classical facility-planning tools to PPF and produces a measured baseline: a hub location, a department layout and a loading rule. Phase II formulates the layout as a learning problem and selects the learning-based optimizers to be tested. Phase III defines a comparative case on a 12-department warehouse. Phase IV proposes a predictive-optimization framework that couples demand forecasting with multi-objective search. Phases I and II-III provide the baseline and benchmark design. Phase IV is a design with a validation plan and is not reported as a result.

## 4.2 Supply-flow redesign

The redesign replaces the supplier → 3PL → warehouse → plant chain (Fig. 2a) with a supplier → cross-dock hub → plant flow (Fig. 2b). The hub receives material from suppliers, repacks it into full truck loads and sends it to the plant in a levelled supply. Six actions support this change:

1. **Demand and planning alignment:** The present 80:20 split between pull-based and MRP-driven supply for Chennai suppliers is re-evaluated, and the plant moves gradually toward demand-driven replenishment.
2. **Network design:** A central transit warehouse is set up near the plant as the consolidation hub, which removes the need for several intermediate storage points.
3. **Flow standardization:** Supplies are pre-picked and ready, and the ad-hoc mix of direct and warehouse supply is replaced by one flow.
4. **Load and packaging:** Material is repacked for full truck loads on 12-ton trucks, and load plans aim at maximum vehicle utilization.
5. **Transport planning:** Trucks run on fixed cycles of four to five trips per day with scheduled dispatch intervals, which improves predictability and reduces waiting.

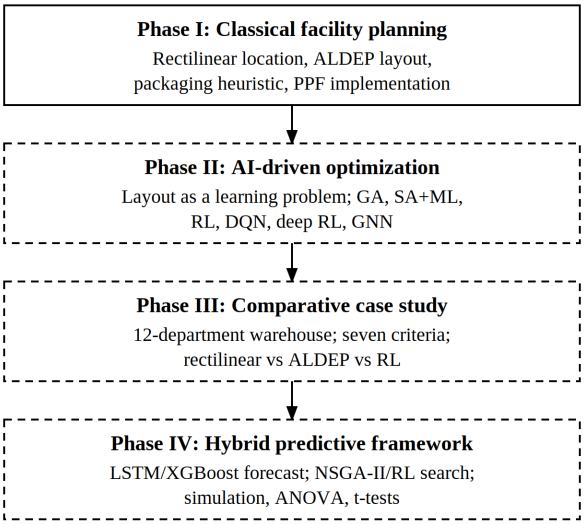


**Fig. 1** Four-phase research design.

6. **Execution and control:** Adherence to the pull system and dispatch schedule is monitored using lead time, truck utilization and cost per shipment, and the model is refined from feedback.

## 4.3 Hub location: rectilinear minisum model

Let supply and demand points $i = 1, \ldots, n$ have coordinates $(x_i, y_i)$ and load weights $w_i$, where the weight is the number of trips multiplied by the load carried. Because material moves along orthogonal roads, distance is rectilinear. The hub location $(X, Y)$ minimizes

$$Z = \sum_{i=1}^{n} w_i \left( |X - x_i| + |Y - y_i| \right) \tag{1}$$

This objective separates into one problem for each axis. Each is solved by a weighted median: $X^*$ is the point where half of the total weight lies on each side along the x-axis, and $Y^*$ is found in the same way along the y-axis. The procedure is to list all points with their coordinates and weights, sort the points on each axis and accumulate their weights, take the position at which the cumulative weight first reaches half of the total, and check the analytical optimum against feasible sites near the plant. Heavier points therefore pull the hub toward them. Gravity and minimax location models were also considered in the project plan, and the rectilinear model was used for the hub decision. Coordinates are treated as planar. The analytical result is a reference point, and the final site depends on availability.

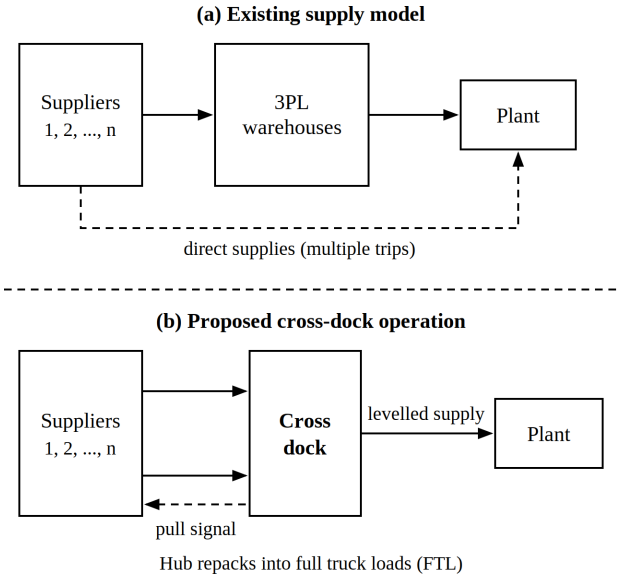


**Fig. 2** Supply flow: (a) existing model with 3PL warehouses; (b) proposed cross-dock operation.

## 4.4 Layout planning: ALDEP

ALDEP builds layouts from qualitative closeness ratings Seehof and Evans [1967], Tompkins et al. [2010]. The facility has seven departments with a total area of 25,000 sq ft, divided into 50 units of 500 sq ft (Table 1). Closeness between every pair of departments is rated on the scale in Table 2 and recorded in the relationship (REL) chart in Table 3.

**Table 1** Department data

| Code | Department | Area (sq ft) | Units |
|---|---|---|---|
| A | Receiving end | 4,000 | 8 |
| B | Shipping end | 3,000 | 6 |
| C | Storage area | 10,000 | 20 |
| D | Office | 1,000 | 2 |
| E | MHE parking | 500 | 1 |
| F | Truck waiting area | 6,000 | 12 |
| G | Common place | 500 | 1 |
| | **Total** | **25,000** | **50** |

ALDEP constructs candidate layouts by sweeping department sequences across the floor plan. Each candidate is scored by summing the closeness values of the department pairs that end up adjacent, so that pairs rated X are heavily penalized. Seven iterations with different sequences are generated, the single-sweep and double-sweep scores of each are compared, and the highest-scoring layout is selected.

**Table 2** Closeness ratings

| Rating | Meaning | Value |
|---|---|---|
| A | Absolutely necessary | 64 |
| E | Especially important | 16 |
| I | Important | 4 |
| O | Ordinary closeness | 1 |
| U | Unimportant | 0 |
| X | Not desirable | −1024 |

**Table 3** REL chart (department numbers 1–7 correspond to A–G)

| | 1 | 2 | 3 | 4 | 5 | 6 |
|---|---|---|---|---|---|---|
| 2 | U | | | | | |
| 3 | A | A | | | | |
| 4 | I | I | E | | | |
| 5 | O | O | O | U | | |
| 6 | E | E | U | O | U | |
| 7 | U | O | O | I | U | O |

## 4.5 Packaging heuristic and truck utilization

An order received at the hub is processed by a volume-based decision flow, implemented as a C program (Fig. 3). The program determines the order quantity and its volume. If the volume fits a returnable box, it computes the number of boxes. Otherwise, if the volume fits a pallet, it computes the number of pallets. Otherwise, it selects the most suitable truck, with full truck load as the target. Small orders thus go in boxes, medium orders on pallets and large orders in full trucks. The aim is to reduce empty space and partial loads. The loading logic follows the same principle as the pallet-loading and truck-matching formulations in the literature Chen et al. [2024], Tang et al. [2024], in a simpler rule-based form.

## 4.6 AI benchmark design (Phases II and III)

The layout problem is posed as a sequential decision problem, in which department placement choices are evaluated by a reward based on handling cost. Six learning-based approaches are selected for benchmarking: a genetic algorithm (GA) Goldberg [1989]; simulated annealing with a machine-learning component (SA+ML) Kirkpatrick et al. [1983]; reinforcement learning (RL) Sutton and Barto [2018]; deep Q-learning (DQN) Mnih et al. [2015]; deep RL; and a graph neural network (GNN) Wu et al. [2021], which represents departments as nodes and closeness ratings as weighted edges. They are compared on a 12-department warehouse using seven criteria: travel distance, cost, picking time, space utilization, computation time, throughput and adaptability to changing demand. The classical results (rectilinear and ALDEP) provide the reference that every learning-based method must beat.

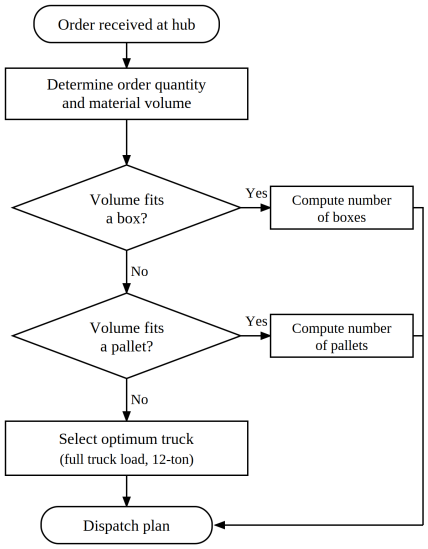


**Fig. 3** Packaging heuristic decision flow implemented in the C program.

### 4.7 Hybrid predictive framework (Phase IV)

The proposed framework has three stages. First, a demand-forecasting model, either an LSTM network Hochreiter and Schmidhuber [1997] or gradient-boosted trees Chen and Guestrin [2016], predicts future flows between departments. Second, the forecasts feed a multi-objective search, using NSGA-II Deb et al. [2002] or RL, that balances conflicting goals such as handling distance, cost and space use. Third, candidate layouts are tested in discrete-event simulation (SimPy or AnyLogic) before any change is made on the floor. Differences among methods are judged with analysis of variance and paired t-tests, with confidence intervals reported for each criterion. Because forecasts change over time, the loop can be repeated to give a layout that adapts to demand, which classical one-off methods cannot do.

## 5 Implementation and results at PPF

### 5.1 Hub location

The hub location was computed from the PPF supply and demand points, using trips multiplied by load as the weight of each point (Section 4.3). Because the heaviest flows dominate the weighted median, the hub is placed close to the highest-load node, which in practice means nearest to the PPF plant. This shortens the heaviest flows and keeps the last leg to the plant short. The exact site is then fixed from the feasible locations near this reference point.

## 5.2 Warehouse layout

ALDEP generated seven candidate layouts for the seven-department, 25,000 sq ft facility. Table 4 lists them. Iteration 6 has the highest closeness score (167) and the highest double-sweep score (334) and was selected.

**Table 4** ALDEP iterations

| Iter. | Closeness combination | Score | Double-sweep score |
|---|---|---|---|
| 1 | A+A+U+E+U+O+O+I+O+I | 155 | 310 |
| 2 | A+A+O+E+U+O+U+O+I+O | 152 | 304 |
| 3 | A+E+E+E+U+O+I+I+U+O | 122 | 244 |
| 4 | E+A+U+E+U+E+O+O+O | 115 | 230 |
| 5 | U+O+E+A+U+E+E+O | 114 | 228 |
| **6** | **E+U+A+A+E+O+I+O+U+O** | **167** | **334** |
| 7 | I+O+E+A+U+E+E+O | 118 | 236 |

Fig. 4 shows the selected layout. Receiving, storage and shipping are arranged in a straight line, so material moves in one direction from receiving to storage to shipping without backtracking. The truck waiting area is next to receiving, so trucks can unload quickly. The office is next to storage, and the common place and MHE parking sit at the dispatch end, away from the main flow.

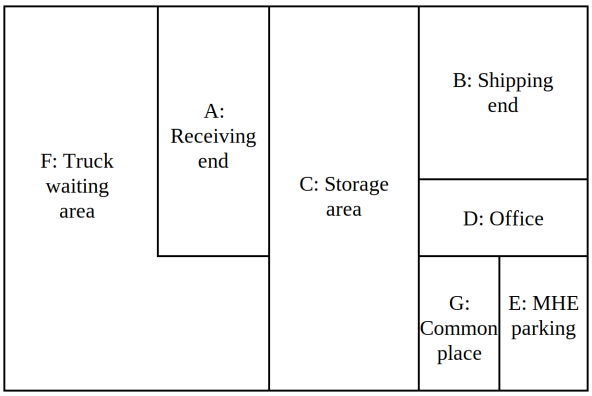


**Fig. 4** Selected warehouse layout from ALDEP iteration 6 (schematic, not to scale).

Table 5 checks the selected layout against the highest-rated relationships in Table 3. Both A-rated pairs and two of the three E-rated pairs are adjacent. The shipping–truck waiting pair is the one E relationship not satisfied, because the two areas lie at opposite ends of the building. This is a trade-off in the selected layout, and trucks waiting for outbound loads would need a short movement through the yard.

**Table 5** Adjacency of A- and E-rated pairs in the selected layout

| Pair | Rating | Adjacent in Fig. 4? |
|---|---|---|
| Receiving – Storage | A | Yes |
| Shipping – Storage | A | Yes |
| Receiving – Truck waiting | E | Yes |
| Storage – Office | E | Yes |
| Shipping – Truck waiting | E | No |

## 5.3 Truck utilization

The packaging heuristic was implemented as a C program that processes each order through the box–pallet–truck decision flow (Fig. 3). Small volumes go into returnable plastic boxes, medium volumes onto pallets and large volumes into full truck loads. In the implementation this gave better load consolidation, a truck fill rate closer to full truck load and fewer partial-load shipments.

## 5.4 Cost impact

Monthly transport cost is the product of the rate per kg, the daily volume moved and the number of working days. Because the volume handled differs between the two states, the comparison here is made on the rate per kg, which does not depend on volume. Table 6 compares the present and proposed states. The proposed rate includes the cost of the hub (godown).

**Table 6** Transport cost: present and proposed state

| | Present | Proposed | Reduction |
|---|---|---|---|
| Rate (INR /kg) | 1.60 | 0.44 (incl. godown) | 72.5% |

The proposed rate is 72.5% lower than the present rate, $(1.60 - 0.44)/1.60$. This reduction comes from consolidation into full truck loads and the removal of the 3PL step. Total monthly cost depends on the volume moved each month and is not compared here.

## 5.5 Discussion

*Operational.* The new flow removes the 3PL step, so material is handled fewer times, the damage risk falls and lead time shortens. A fixed dispatch rhythm makes arrivals predictable, which helps the plant plan production. The selected layout supports this by keeping the flow linear.

*Cost.* The rate per kg falls from INR 1.60 to INR 0.44, a reduction of 72.5%, which comes from consolidation, fuller trucks and fewer trips. The effect on total monthly cost depends on the volume moved and is not claimed here.

*Strategic.* The case shows a move from decisions made by experience to decisions made from data. Location, layout and loading each have an explicit model, so each can be re-run when volumes or suppliers change. This also sets up the AI phases: any learning-based method must beat these documented results to be worth adopting. The truck utilization reported in the literature Chen et al. [2024], Tang et al. [2024] concerns less-than-truckload hubs, so it is a different setting and these results should not be compared with it directly.

Table 7 summarizes the outcomes against the objectives of Section 3.

**Table 7** Results against objectives

| Objective | Outcome |
|---|---|
| O1, O2: assess current flow and bottlenecks | Done; six problems identified (Section 3) |
| O3: hub location | Hub placed near the dominant demand node |
| O4: ALDEP layout | Iteration 6 selected, score 167 / 334 |
| O5: packaging and truck utilization | Decision flow implemented; quantified fill rates not reported |
| O6: AI benchmark and hybrid framework | Designed; not yet validated |

# 6 Limitations, conclusion and future work

## 6.1 Limitations

Several limits apply to the results. First, the classical baseline depends on the accuracy of the demand, trip and cost data supplied for PPF, and only one plant is studied, so the findings should not be generalized without further cases. Truck fill rates and total monthly savings are not quantified here, and the cost comparison is made per kg. Second, ALDEP is a heuristic: seven iterations do not prove a global optimum, the closeness ratings are subjective judgments, and the selected layout leaves one E-rated pair (shipping and truck waiting) non-adjacent. Third, the rectilinear model treats coordinates as planar and weights as fixed, so it gives a snapshot of the flows and not a response to changing demand. Fourth, the AI layer is a design and an evaluation plan: no benchmark results or simulation outputs are reported, and no claim is made that learning-based methods outperform the classical baseline. Finally, human factors such as labour constraints are not modelled, and applying the new layout and flow in practice will need redesign effort and change management.

## 6.2 Conclusion

This paper studied the inbound logistics of a clutch manufacturing plant, where material moved through a chain of suppliers, a 3PL provider and a warehouse before reaching the plant. A classical baseline was built from a rectilinear hub-location model, ALDEP layout planning and a volume-based packaging heuristic. ALDEP iteration 6 gave the best layout for the seven-department, 25,000 sq ft facility, with a closeness

score of 167 and a double-sweep score of 334, and the layout keeps receiving, storage and shipping in one direction of flow. The proposed supplier-to-hub-to-plant flow with full truck loads is reported to lower the transport rate from INR 1.60 to INR 0.44 per kg, a reduction of 72.5%.

On this baseline, the paper defines a benchmark of learning-based optimizers on a 12-department warehouse and a hybrid LSTM and NSGA-II framework for demand-aware, multi-objective decisions. These remain proposals to be tested. The contribution at this stage is therefore a documented, transparent baseline and a clear plan against which adaptive methods can be judged.

### 6.3 Future work

Future work will run the proposed benchmark of GA, SA+ML, RL, DQN, deep RL and GNN on the 12-department case and report all seven criteria with analysis of variance and paired t-tests. The hybrid LSTM and NSGA-II framework will be implemented and tested in discrete-event simulation before any change is made on the floor. A digital twin of the PPF hub and warehouse, linked to ERP and real-time shipment data, would allow layouts and dispatch plans to adapt as demand changes. Further extensions include human-centric constraints such as labour and fatigue, explainable reporting of learned decisions, and application to more plants and suppliers.

## References


V. Agarwal, A. Z. Hameed, S. Malhotra, K. Mathiyazhagan, S. Alathur, and A. Appolloni. Role of Industry 4.0 in agile manufacturing to achieve sustainable development. *Business Strategy and the Environment*, 32(6):3671–3688, 2023. doi: 10.1002/bse.3321.

T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM, 2016.

X. Chen, W. Tang, Y. Hai, M. Lang, Y. Liu, and S. Li. Optimization of truck–cargo matching for the LTL logistics hub based on three-dimensional pallet loading. *Mathematics*, 12(21):3336, 2024. doi: 10.3390/math12213336.

G. B. Dantzig and J. H. Ramser. The truck dispatching problem. *Management Science*, 6(1):80–91, 1959.

K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6 (2):182–197, 2002.

R. Dekker, J. Bloemhof, and I. Mallidis. Operations research for green logistics – an overview of aspects, issues, contributions and challenges. *European Journal of Operational Research*, 219(3):671–679, 2012.

M. Dorigo and L. M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, 1997.

D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, USA, 1989.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9 (8):1735–1780, 1997.

S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

G. Laporte. Fifty years of vehicle routing. *Transportation Science*, 43(4):408–416, 2009.

R. F. Love, J. G. Morris, and G. O. Wesolowsky. *Facilities Location: Models and Methods*. North-Holland, New York, NY, USA, 1988.

Y. Lu, C. Liu, K. I.-K. Wang, H. Huang, and X. Xu. Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61:101837, 2020. doi: 10.1016/j.rcim.2019.101837.

V. Mnih, K. Kavukcuoglu, D. Silver, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

J. M. Seehof and W. O. Evans. Automated layout design program. *Journal of Industrial Engineering*, 18(12):690–695, 1967.

R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA, 2 edition, 2018.

W. Tang, X. Chen, M. Lang, S. Li, Y. Liu, and W. Li. Optimization of truck–cargo online matching for the less-than-truck-load logistics hub under real-time demand. *Mathematics*, 12(5):755, 2024. doi: 10.3390/math12050755.

J. A. Tompkins, J. A. White, Y. A. Bozer, and J. M. A. Tanchoco. *Facilities Planning*. Wiley, Hoboken, NJ, USA, 4 edition, 2010.

P. Toth and D. Vigo, editors. *Vehicle Routing: Problems, Methods, and Applications*. SIAM, Philadelphia, PA, USA, 2 edition, 2014.

Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021.

Y. Yan, A. H. F. Chow, C. P. Ho, Y.-H. Kuo, Q. Wu, and C. Ying. Reinforcement learning for logistics and supply chain management: Methodologies, state of the art, and future opportunities. *Transportation Research Part E: Logistics and Transportation Review*, 162:102712, 2022. doi: 10.1016/j.tre.2022.102712.

L. Zhang, C. Yang, Y. Yan, Z. Cai, and Y. Hu. Automated guided vehicle dispatching and routing integration via digital twin with deep reinforcement learning. *Journal of Manufacturing Systems*, 72:492–503, 2024. doi: 10.1016/j.jmsy.2023.12.008.