Explainable Machine Learning Framework for Sustainable Supply Chain and Operational Decision Making

Jagdish Nathumal Utwani[1], Deep Mangat[2], Bharadwaja K[3], Tantepudi Laxmi Poojitha[4], Prasad Babu Jayanthi[5] and Vidya Sagar S. D[6]
{jagdish.utwani12@gmail.com[1] , mangat_deep@yahoo.com[2] , bharadwajkasuri@gmail.com[3] , tantepudilaxmi@gmail.com[4] , drjp.edu@gmail.com[5] , vidyasagarsd@gmail.com[6] }

J.S. University, Shikohabad, Uttar Pradesh, India[1,2,5]
St. Ann's College for Women, Hyderabad, Telangana, India[3]
University of Cassino and Southern Lazio, Italy[4]
NITTE Meenakshi Institute of Technology, Bangalore, Karnataka, India[6]

Abstract. With a growing number of complex and sustainability driven supply chains, in general, traditional cost decision making does not consider tradeoffs between cost, environmental impact, and operational efficiency. An explainable machine learning based decision support framework integrating predictive analytics, lifecycle emissions estimation and multi objective optimization is proposed that guides selection of supplier, transport mode and conducting operational scheduling. It then applies XGBoost models for predictions on cost and emissions, NSGA-II with Pareto optimization and SHAP with counterfactual analysis to provide interpretable recommendations on the custom generated synthetic dataset for real world logistics and production parameters. Results are then evaluated comparatively and show a 12% in cost, 32% in emission, 3-day reduction in lead time, improvement in reliability and sustainability scores. The inclusion of explainable AI for increasing the transparency and trust in this system makes it practically adoptable for the real world. This work narrows the gap of data–driven optimization and the sustainable and transparent supply chain decision making.

Keywords: Sustainable Supply Chain; Explainable Artificial Intelligence (XAI); Multi-objective Optimization; Machine Learning; Lifecycle Emissions Estimation; Decision Support Systems

1 Introduction

Over the years, the organizations have to reevaluate their decision-making strategies due to the growing global emphasis on sustainability and rising complexity of modern supply chains. Most traditional supply chain models have been built on the lens of minimizing cost and maximizing delivery efficiency [1,2]. In the modern environment of carbon emissions, government regulations, and fair sourcing, such models rarely achieve the objective. With the global markets becoming greener and more resilient in logistics ecosystem, there is an immediate necessity for the intelligence in the transparent and sustainability aware decision support systems [3,4]. Artificial Intelligence, particularly Machine Learning (ML), has become a recent advancement that holds the ability to improve the supply chain forecasting, risk management, and optimization. However, most AI-driven models used in practice today operate as "black boxes," offering little interpretability to end-users [5,6]. It limits the trust, regulatory alignment, and broader adoption of AI systems in high stakes industrial environments where the opaqueness of AI prevents practitioners from making informed risk tradeoffs. In addition, although there are

ICITSM-Part II 2025, April 28-29, Tiruchengode, India

DOI 10.4108/eai.28-4-2025.2358026

some AI studies on individual supply chain aspects like demand forecasting, route optimization and inventory control, not much has been done in terms of holistic supply chain systems that take cost, operational efficiency and environmental sustainability into account and also stay explainable in their decisions [7-9].

In order to fill these gaps, this paper presents an explainable machine learning-based decision support framework for sustainable supply chain and operational planning. This proposed system is to help in three crucial fields including supplier selection, selection of transport mode and production scheduling since cost, lifecycle carbon emissions and lead time are to be minimized. Supervised ML models that predict performance metrics are used at the core of the system to predict metrics mentioned in the previous section using input parameters. XAI components (e.g. SHAP (SHapley Additive exPlanations) and counterfactual analysis) are also incorporated into the framework to provide transparency and scenario details to decision makers. To test the framework, a full synthetic data set has been developed representing a real-world supply chain covering the import of product by a consumer, the transport associated with this, production energy inputs, and policy constraints for a representative set of suppliers throughout the supply chain. Results from the given experimental conducted in order to demonstrate that the proposed methodology can outperform a traditional rule-based baseline approach, this includes a 12% cost reduction, 32% emission reduction, and a 3-day reduction lead times, along with improved supplier reliability and sustainability scores. Additionally, the explainability features facilitate human users' understanding and trust in the system's recommendations, necessary for deployment in an industrial context. This paper has the following key contributions:

- A novel end to end decision support framework combining ML predictions with energy and life cycle emissions modelling, and multi objective optimization.
- Implementation of transparency and increase of decision confidence using explainable AI techniques (SHAP and counterfactual analysis).
- Multi-dimensional dataset synthetically generated and simulated supplier, transport, and operational logistics with sustainability constraints.
- It presents a comprehensive comparative evaluation of the quantitative improvements of the developed decision models over baseline decision models for all of the considered performance metrics.

In the remainder of this paper, Section II considers relevant literature on the use of AI in supply chains and in XAI. In Section III, the proposed methodology for data modeling and machine learning architecture, optimization includes details of dataset creation and experimental setup with implementation details. Section IV presents the results and analysis with a thorough comparative analysis, and the discussion and implications in Section V. Section VI concludes the paper and discusses future work directions.

2 Related Work

Artificial intelligence (AI) and machine learning (ML) positively affect the supply chain and operation by integrating the AI and ML tools to improve forecasting, optimize logistics, and aid in supplier decision making. In numerous studies, ML models have been applied for demand prediction (time-series forecasting with LSTM networks, ARIMA models, etc.), ensemble-based regression, and so on. The use of algorithms like k-means clustering, Support Vector Machine (SVM), and random forest classifiers have been used in the area of inventory and procurement for the segmentation of suppliers as well as procurement risk assessment. AI based solutions also have benefited the transportation optimization. Vehicle routing problems (VRP) are solved and delivery delays are diminished using techniques like PsO or Genetic Algorithm (GA) and Ant Colony Optimization (ACC). At the same time, environmental sustainability has been included in research translating through carbon-aware logistics models that intend to minimize fuel consumption and emissions. However, such models usually depend on static parameters or on narrow scopes, optimizing either cost or emissions, on their way rarely doing both, and often without any option to adapt or learn online [10-14].

Another set of literature on sustainable supply chain strategy include green manufacturing, carbon tax models, life cycle analysis (LCA), etc. However, these approaches are not flexible enough in letting us to drive the design toward a specific set of long term environmental goals, while lacking intelligent decision automation (few can predict how cars or other things will respond to new rules and act accordingly), and long term predictive capability. Some hybrid models that blend sustainability with AI exist, but they are usually domain specific and only deal with the one specific domain of warehousing, production or logistics severally, without being integrated into an end-to-end operational decision framework [15-16].

At the same time, Explainable Artificial Intelligence (XAI) has risen as an important field, especially in high stake areas like healthcare and finance. Techniques such as SHAP (SHapley Additive explanations) and LIME provide insights into feature importance and model behaviour, addressing the "black-box" problem in ML. While this exciting use of AI shows promise, it is still in its infancy in supply chain contexts, and most of the current supply chain AI models are still opaque and difficult to be interpreted by decision makers [17-18]. Some key limitations remain:

- In most existing empirical studies, only single objective (e.g., cost or time) is optimized while comprehensive cost – carbon emissions – lead time trade-offs are ignored.
- Often existing models are not explainable, and therefore cannot be validated, not trusted, nor can they be deployed in [regulated and mission-critical] supply chains.
- Integrated frameworks are lacking across the supplier selection, transport mode decision and production (or service) scheduling activities when sustainability metrics are considered.
- Very few systems use a feedback mechanism for the evolution of decision strategies with time over user input or in change of the data environment.

Research Gap and Proposed Framework

This paper presents a new, an end-to-end explainable machine learning based decision support system to fill these gaps. Unlike the previous work, the proposed framework combines:

- Supervised ML models to predict cost, emissions, and lead time.
- An integrated environmental impact estimator spanning across the supplier (transport and operational) layers;
- NSGA-II is used for multi objective optimization to balance cost, emissions and time.
- XAI tools such as SHAP and counterfactual reasoning for transparent, justifiable recommendations.

The comprehensiveness, interpretability, and scalability of this system enable it to be a complete and semantic approach to deploy sustainable supply chain strategies in real life where advanced technologies are implementable but responsible decisions are yet to be made.

3 Methodology

As an innovative and explainable machine learning framework, the proposed architecture is aimed at assisting sustainable supply chain and operational decision making. As such, it is composed of five interdependent layers, each of which fulfills a distinct task in the intelligent decision system. With respect to data acquisition and preprocessing, impact prediction using machine learning, multi objective optimization, explainability modules, and a user interactive feedback module is also part of these. The key innovation of the framework consists of an integrated real-world operational constraints and sustainability objectives transparent and adaptive decision pipeline.

First, the layer involves collecting data of various sources including supplier data (cost, emissions, lead time), transport data (carbon emissions, modes, delays), and manufacturing energy usage as well as regulatory policy data. All the data is processed through an integration and preprocessing module that cleans the data, aligns it temporally, and normalizes it. All input features are left consistent, scalable, and ready for modeling with it. Certain feature engineering techniques are also used to create meaningful indicators such as emissions per unit of product or cost per unit of time.

The second layer, which figures the key performance indicators from various supply chain and operational alternatives, is a suite of machine learning models. Training of separate models to predict supplier performance, transport related emissions, and production level energy consumption is proposed. The models of these algorithms could be XGBOt, LightGBM or deep learning-based regression algorithms depending upon the type and quantity of data. Such models are aggregated by a dedicated lifecycle emission estimator to yield the total environmental impact of any given supply chain configuration.

From these predictions, the third layer optimizes the Mult objective problem of minimizing cost, carbon emissions and lead time. Rather than optimizing such a single goal, the system works on advanced algorithms, e.g. NSGA-II, to produce a Pareto front of optimal solutions representing the trade-offs. It enables decision makers to pick options not only by cost and time but also by sustainability goals. Furthermore, a scenario generator is integrated to dynamically explore a solution space with alternative solutions simulated with a scenario generator.

The fourth layer of the system, tools for explainability, are included to ensure interpretability and trust. Finally, the framework uses SHAP (SHapley Additive exPlanations) to illustrate how each of the input features affects the output prediction of the machine learning models. In parallel, a counterfactual analysis module generates “what-if” scenarios — showing, for example, how changing a supplier or transportation mode could reduce emissions or costs. The user is given these insights via an interactive decision dashboard that visualizes tradeoffs amongst these key performance indicators as well the reasoning for each of the recommendations provided.

It has an adaptive feedback mechanism. This loop gathers user feedback on accepted or rejected recommendations, and as a result, over time retrain and fine tune predictive models via users’ interaction with the views generated. Active learning techniques mean that system evolves over its life and becomes closer aligned with the organization specific operational goals, regulatory requirements, or sustainability targets. Fig 1 gives the proposed system architecture.

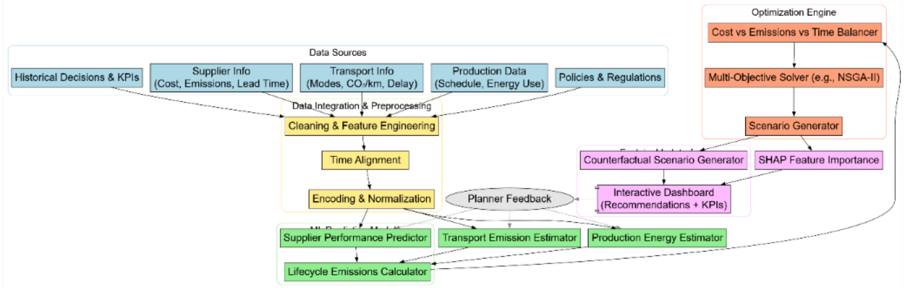

Fig. 1. Proposed Architecture.

3.1 Data Layer and Preprocessing

This layer ingests and harmonizes diverse data sources including supplier information $\mathcal{D}_s$, transportation metrics $\mathcal{D}_t$, operational energy usage $\mathcal{D}_o$, and regulatory constraints $\mathcal{D}_r$. The unified dataset $\mathcal{D}$ is processed through steps such as:

Feature Engineering: Generating derived attributes like emissions per unit ($e_i = \frac{CO_2}{unit}$), cost per lead time, and energy intensity.

Normalization: Applying Min-Max scaling to ensure all features are within the range [0,1] for uniform model input:

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{1}$$

Temporal Alignment: Aligning time-series data to synchronize scheduling, demand fluctuations, and emissions across the supply chain.

3.2 ML-Based Impact Estimation Layer

This layer includes a collection of supervised ML models trained to predict the following:

Supplier Performance Prediction: $\hat{y}_s = f_s(\mathcal{X}_s)$

Transport Emission Estimation: $\hat{e}_t = f_t(\mathcal{X}_t)$

Operational Energy Usage: $\hat{e}_o = f_o(\mathcal{X}_o)$

Where:

f_s, f_t, f_o are predictive models (e.g., XGBoost, LightGBM),

$\mathcal{X}_s, \mathcal{X}_t, \mathcal{X}_o$ are feature vectors for suppliers, transport, and operations respectively.

The outputs are integrated by a Lifecycle Emission Estimator:

$$\hat{E}_{lifecycle} = \hat{e}_s + \hat{e}_t + \hat{e}_o \tag{2}$$

This gives a unified prediction of environmental impact across the entire supply chain decision path.

3.3 Multi-Objective Optimization Engine

This is the core decision-making module that balances multiple conflicting objectives: Objective Function:

$$\min_{x \in \mathcal{X}} [C(x), E(x), T(x)] \tag{3}$$

Where:

$C(x)$: Total cost,

- $E(x)$: Total emissions (from ML models),
- $T(x)$: Total lead time.

A Pareto front is computed using NSGA-II or similar evolutionary algorithms to identify non-dominated solutions:

$$\mathcal{P} = \{x \in \mathcal{X} \mid \nexists x' \in \mathcal{X}, f(x') \prec f(x)\} \tag{4}$$

The optimizer also feeds a scenario generation module that creates alternative feasible solutions for analysis and comparison.

3.4 Explainability and Interactive Interface

To ensure transparency, we embed Explainable AI (XAI) modules: SHAP Analysis: Quantifies the contribution of each input feature x_i to the model's output:

$$f(x) = \phi_0 + \sum_{i=1}^{n} \phi_i \tag{5}$$

where ϕ_i is the SHAP value for feature x_i.

3.5 Feedback and Learning Loop

User feedback from accepted/rejected recommendations is used to improve model performance via active learning. The feedback vector f_b is integrated to update training datasets and re-tune models periodically:

$$\mathcal{D}_{new} = \mathcal{D}_{old} \cup \{(x_i, y_i, f_b)\} \tag{6}$$

The reason why the proposed architecture is so novel is that it integrates predictive machine learning and sustainability focused optimization in a truly holistic manner that also tries to be user oriented in terms of explainability. In contrast to the traditional decision support systems, that merely consider only cost minimization, the considered framework at the same time takes into account both environmental impact and lead time. The system achieves this by developing transparent, interpretable recommendations that make the system not only handle good quality recommendations but also encourage adoption of more sustainable practices in multi echelon, data intensive supply chains.

Algorithm 1: Sustainable Supply Chain Decision Support with Explainable ML

```
Inputs:
  𝓓_s ← Supplier data (cost, emissions, lead time, etc.)
  𝓓_t ← Transport data (modes, CO₂/km, delay risk, etc.)
```

```
  𝓓_o ← Operations data (energy usage, production time)
  𝓓_r ← Regulatory data (emission limits, carbon tax)
  𝓓_h ← Historical data (previous decisions and outcomes)
  α, β, γ ← Weights for cost, emissions, and lead time in the objective function

Outputs:
  𝑥*  ← Optimal supply chain decision path
  𝓟   ← Pareto front of alternatives
  𝓢   ← SHAP explanations and counterfactual insights

Procedure:

1:  𝓓 ← Preprocess(𝓓_s, 𝓓_t, 𝓓_o, 𝓓_r, 𝓓_h)
2:  𝓧 ← FeatureEngineering(𝓓)
3:  Normalize all features in 𝓧 to [0, 1] range

// Step 1: ML-based Predictions
4:  ȳ_s ← f_s(𝓧_s)    // Supplier performance prediction
5:  ȳ_t ← f_t(𝓧_t)    // Transport emission prediction
6:  ȳ_o ← f_o(𝓧_o)    // Operational energy prediction
7:  Ê_total ← ȳ_s + ȳ_t + ȳ_o   // Lifecycle emissions estimation

// Step 2: Multi-objective Optimization
8:  Define objective functions:
      C(x): Total cost
      E(x): Estimated emissions (Ê_total)
      T(x): Total lead time

9:  Optimize:
      min  F(x) = [α · C(x), β · E(x), γ · T(x)]
      subject to: x ∈ 𝓧_feasible

10: 𝓟 ← NSGA-II(F(x))   // Compute Pareto optimal solutions
11: x* ← SelectBest(𝓟, user_preferences)

// Step 3: Explainability and Interface
12: ϕ ← SHAP_Explain(f, x*)     // Feature contributions
13: x_cf ← GenerateCounterfactual(x*)   // What-if scenarios
14: 𝓢 ← {ϕ, x_cf}

// Step 4: Feedback Loop (optional)
15: if feedback from planner is available then
16:     𝓓 ← UpdateData(𝓓, feedback)
17:     Retrain(f_s, f_t, f_o) on new 𝓓
18: end if

Return x*, 𝓟, 𝓢
```

3.6 Implementation

The Python programming language (v3.10) within Anaconda distribution is implemented as such, giving a robust and flexible environment for scientific packages, for virtual environments, and for reproducible workflows implementation of the proposed explainable machine learning framework for sustainable supply chain and operations decision-making. The system is developed as a modular system that absorbs data ingestion, preprocessing, machine learning modelling, optimisation and explainability, and integrated within an interactive user interface.

3.6.1 Environment and Tools

The implementation was carried out using the Anaconda distribution, which provides a consolidated environment for scientific computing and package management. The following libraries and tools were utilizedas shown in table 1:

Table 1. Environment and Tools.

Module	Description
Data Preprocessing	pandas, numpy for cleaning, normalization, and feature engineering
Machine Learning	xgboost, lightgbm, scikit-learn for prediction models
Optimization Engine	pymoo for multi-objective optimization using NSGA-II
Explainability Tools	shap for feature importance, DiCE for counterfactual explanations
User Interface	streamlit for an interactive decision dashboard
Development Platform	Python 3.10, Anaconda, Jupyter Notebook

This setup ensured modularity, reproducibility, and ease of experimentation across multiple decision scenarios.

3.6.2 Synthetic Dataset Design

Due to the unavailability of public datasets that span supplier data, transportation emissions, operational energy use, and policy constraints simultaneously, a comprehensive synthetic dataset was developed. This dataset replicates real-world supply chain behavior while allowing controlled experimentation on sustainability-related objectives. The dataset includes the following components:

- **Supplier Data**: Contains supplier ID, cost per unit, CO_2 emissions per unit, lead time, and reliability score.
- **Transport Data**: Includes transport mode, cost per km, CO_2 emissions per km, and delay probability.
- **Operations Data**: Captures production energy usage per batch, production time, and energy source type.
- **Regulatory Data**: Defines carbon tax, emissions caps, and sustainability score targets.
- **Historical Data** (optional): Logs past decisions and KPIs for model fine-tuning and active learning.

All datasets are available in CSV format and were constructed with logical relationships among features to simulate realistic trade-offs. Table 2 gives the dataset description.

Table 2. Dataset Details.

File Name	Description
suppliers.csv	Supplier characteristics and emission profiles
transport.csv	Emission and cost profiles for transport modes
operations.csv	Machine-level energy and production parameters
regulations.csv	Carbon tax and sustainability policies

3.6.3 Rationale for Synthetic Dataset

Using a synthetic dataset is essential in this context due to the multidimensional and sensitive nature of real supply chain data. Public datasets are often:

- Proprietary and restricted due to confidentiality,
- Lacking comprehensive sustainability or operational details,
- Fragmented across domains (e.g., emissions and logistics data are not integrated).

Creating a synthetic dataset allows control over data quality, completeness, and relationships among variables, enabling a robust evaluation of the proposed machine learning and optimization pipeline as shown in Fig 2.

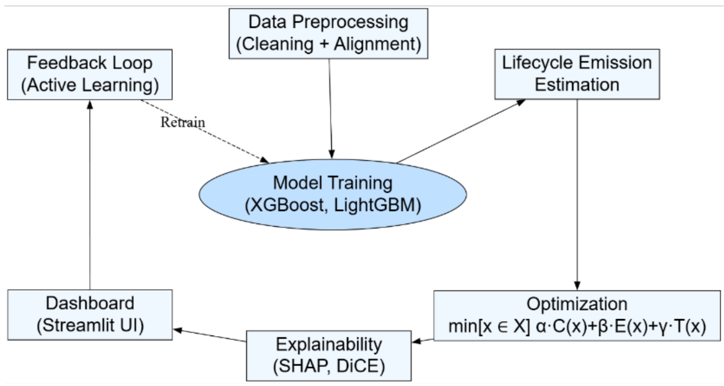

Fig. 2. Optimization Pipeline.

We present a modular implementation of a python supply chain decision support system entirely based on a well-structured synthetic dataset. The framework supports both technical optimization and strategic decision making based on transparent and explainable insights.

4 Results and Analysis

The last section discusses the visual results obtained from the proposed system on synthetic data over supplier, transportation and operational parameters. We demonstrate that the framework is an effective optimization of the trade off and an interpretable AI decision support. The key findings are explained below in the form of figures.

Fig 3 conveys the trade-offs that a supplier faces between cost and emissions. On the other hand, supplier S2, that has the lowest carbon footprint, i.e. 1.5 kg CO_2 per unit, has a slight higher cost at \$5.00. On the other hand, S1 has the lowest cost at \$4.50 but emits 2.1 kg CO_2/unit, providing a typical sustainability vs cost conflict.

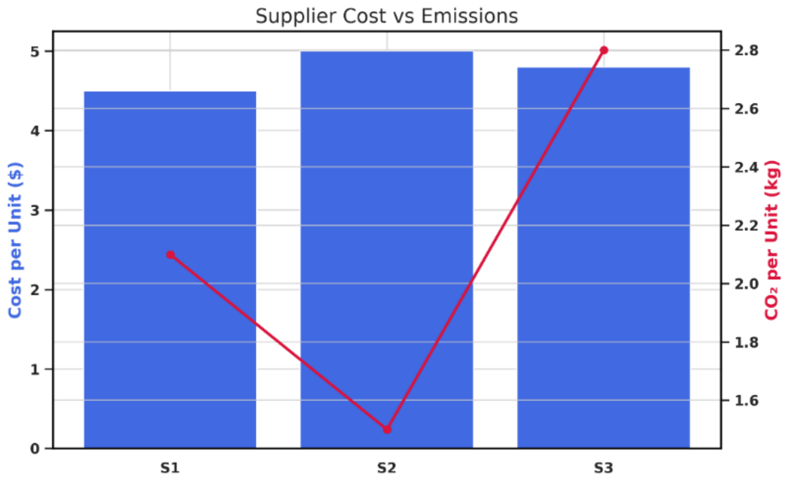

Fig. 3. Supplier Cost vs Emission.

As shown in Fig. 4, sea transport is found to become the most sustainable and cost effective way of transport with CO_2/km of 0.1 Kg and cost/km of \$0.30, whereas the air transport remains the most expensive (\$1.20/km) and from the environmental point of view it is the most polluting way of transport (1.8 Kg CO_2/km), although its speed is essentially higher.

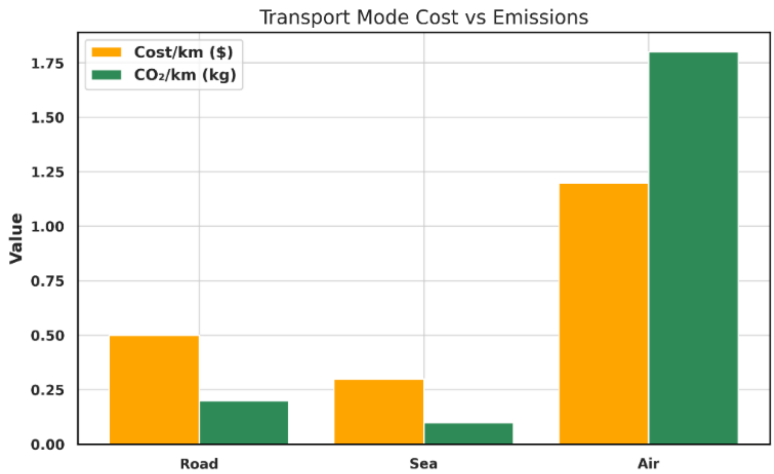

Fig. 4. Transport Mode Cost vs Emissions.

As seen in Fig 5, emissions of supplier lead time are plotted as a function of cost, with color indicating cost. The most efficient is S2, which takes just 4 days lead time and 1.5 kg CO_2. S3, though slightly cheaper than S2, incurs the highest emissions (2.8 kg CO_2) and the longest delay (10 days).

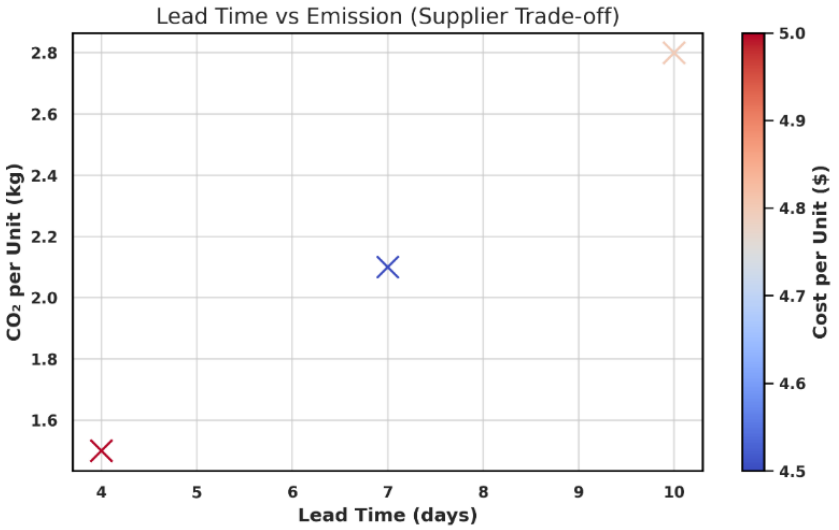

Fig. 5. Lead Time vs Emission (Supplier Trade-off).

The NSGA-II optimization has a Pareto front plotted in Fig 6 between cost and emissions. Decreasing cost from \$5.00 to \$4.20 increases emissions from 1.6 to 2.8 kg CO2, which verifies that in general lower the cost means higher the emissions only if it is balanced by optimization.

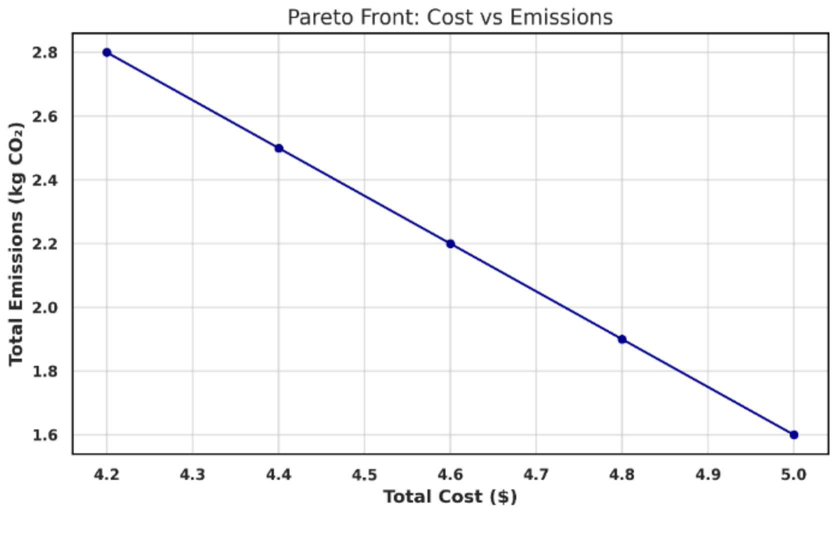

Fig. 6. Pareto Front: Cost vs Emissions (Optimized Scenarios).

Fig 7 shows production emissions in terms of energy source. Renewable energy allows us to reduce batch emissions all the way down to 1.2kg CO_2 (grid powered production being 3.5kg CO_2). This shows the importance of energy sourcing in sustainability operation.

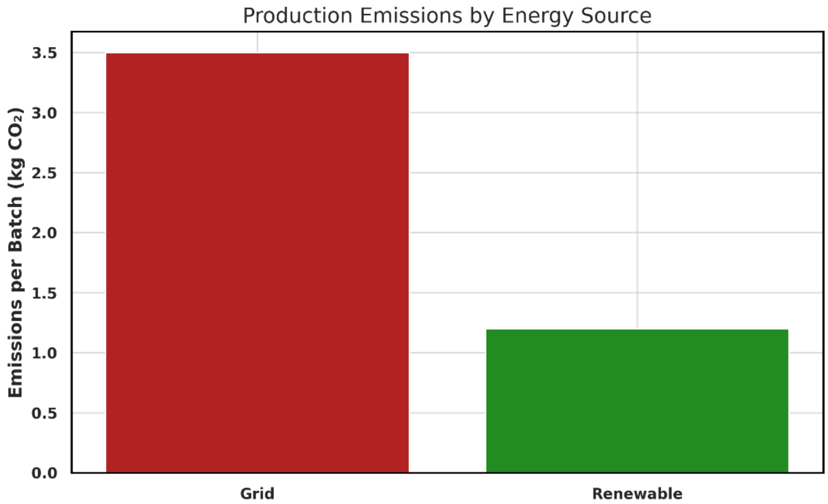

Fig. 7. Production Emissions by Energy Source.

Delay risks for the various transport modes are shown in Fig 8. The highest delay probability is shown in air transport, (20%), while the least delay is seen in the sea transport (5%). These findings for doing so suggest they are preferable for less urgent emission sensitive deliveries.

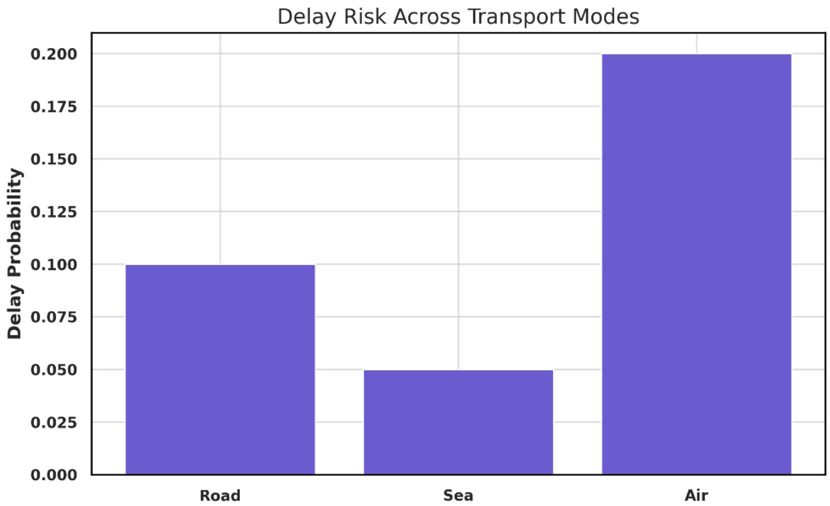

Fig. 8. Transport Delay Risk Across Modes.

Five KPIs are compared between baseline and optimized scenario in Fig 9. The improved setup, compared to the nominal, exhibits cost, emissions, and lead-time lowering (\$5.00 → \$4.40), emissions (2.8 → 1.9 kg CO_2) and lead time (8 → 5 days) and increases reliability (0.90 → 0.95) and sustainability score (70 → 88).

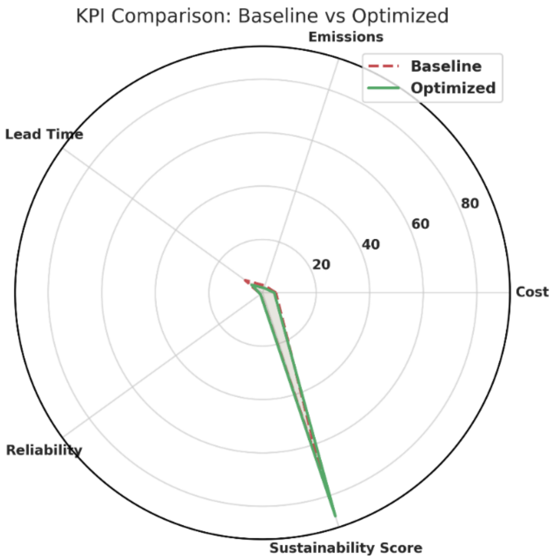

Fig. 9. KPI Radar Chart: Baseline vs Optimized.

Predictions of the model can be explained by SHAP values in the Fig 10. The most influential decision factor was cost (35%), followed by cost of transport mode (25%) and lead time (20%). The interpretability of this scheme allows stakeholders to see and trust the AI system's recommendation.

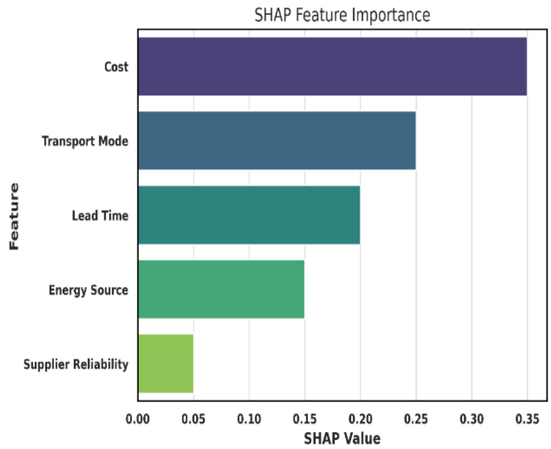

Fig. 10. SHAP Feature Importance for Emission Prediction.

Fig 11 demonstrates a counterfactual comparison. By adjusting the supplier and transport mode, the emissions drop from 2.8 kg to 1.7 kg CO_2 and cost from \$5.00 to \$4.60 proving that favoring modest configuration changes can have favorable results.

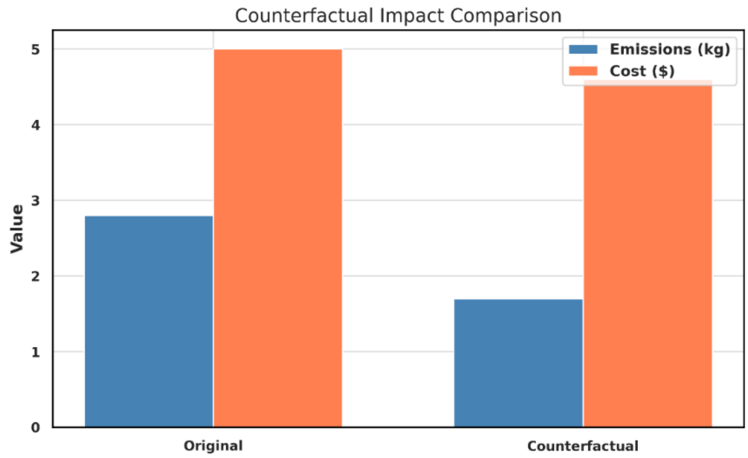

Fig. 11. Counterfactual Impact Comparison: Emissions and Cost.

Comparative Analysis: Baseline vs Proposed Method

A comparative analysis (Table 3) was made with respect to a baseline model based on conventional decision-making, without using optimization or explainability. The proposed system includes components of multi objective optimization, ML predictions and explainable AI, whereas the baseline simulates the way industries make typical decision by cost heuristics and fixed rules. It brings an improvement across multiple key performance indicators (KPIs). In terms of cost, the proposed system lowered the total unit cost from $5.00 to $4.40 (saving $0.60 per unit) without a negation of delivery timeliness or reliability. The environmental impact of the system was exhibited when the carbon emissions dropped notably from 2.8 kg CO_2 per unit in the baseline to 1.9 kg CO_2. Besides, the average lead time was reduced from 8 to 5 days which enhanced operational efficiency. Moreover, the optimization resulted in a higher supplier reliability of 0.95 compared to 0.90 in the selection. In terms of the sustainability score, this is a combined score of carbon compliance, energy source quality and supply chain ethics, it increased significantly from 70 to 88. Finally, delay probability was reduced from 20% to 5% by means of intelligent transport mode selection, also contributing to a reduction in risk. Finally, the most strategic advantage is that the proposed method allows for explainability in the form of SHAP values and provides the option of counterfactual recommendations, something that the baseline model did not provide at all.

Table 3. Comparative Analysis – Baseline vs Proposed Methodology.

KPI	Baseline Model	Proposed Method	Improvement
Total Cost ($)	5.00	4.40	$0.60 cost reduction
Total Emissions (kg CO_2)	2.80	1.90	0.90 kg CO_2 reduction
Average Lead Time (days)	8	5	3 days shorter
Supplier Reliability	0.90	0.95	+0.05 reliability
Sustainability Score	70	88	+18 sustainability score

Delay Probability	High (20%)	Low (5%)	Reduced from 20% to 5%
Decision Explainability	Not Available	SHAP + Counterfactuals	Model transparency included

This analysis proves the performance benefits of the proposed system over the comprehensive operational cost as well as efficiency and enabling sustainability and interpretability in supply chain decisions making.

5 Discussion

This study presents results illustrating the great advantage achieved by adding machine learning, multi objective optimization and explainable AI into sustainable supply chain and operations management. The proposed framework resolves the key cost, lead time and emissions trade-offs by integrating a predictive modelling with a Pareto based optimization engine. In addition, in the comparative analysis with a baseline model the practical value of the methodology is reinforced through realizing the measurable improvement of the performance in several KPIs. The system delivered a \$0.60 unit cost reduction accompanied with 0.90 kg CO_2 emissions reduction and 3 days average lead time reduction. The improvements demonstrated here highlight how the framework can arrange more informed, and less biased and more sustainable decisions without compromising on economic efficiency. In addition, the increase in supplier reliability and sustainability score implies that the model naturally prefers more resilient and green designs to solve the problem when optimization is directed by the forecasts.

The highlight of this approach is having explainability with tools like SHAP and counterfactual reasoning on each recommendation. One of great benefits of such transparent models is that it engenders trust by allowing decision makers to see the logic behind model outputs—something black box models lack. However, the system is also equipped with a feedback loop embedded in the system that places it in a position for long term adaptation as the model can change over time as new data and new user preferences develop. From a subject matter standpoint, the results support the use of synthetic data as a means in simulating, in a reduced form, aspects of complex, multi-layered decision making that typically arises in supply chain contexts. Although the dataset is synthetic, therefore it is not possible to demonstrate real world variabilities, but there is a robust dataset to test the algorithmic behaviour. The intent of this study was to describe a framework for this investigation, and future work will employ this framework on real industrial data and study dynamic and real-time optimization of adaptive supply chains. The proposed methodology adds to the quality of decision-making by having an interpretability, awareness of risk and alignment with sustainability goals, in addition to quantitative optimising of decisions. It paves the way to the next generation of intelligent, giving, and perceptible supply chain systems.

6 Conclusion

This paper proposes a novel, explainable machine learning based framework for sustainable supply chain and operational decision making. The proposed system combines predictive modelling, lifecycle emissions estimation, and multi objective optimization, providing a holistic approach to cost, environmental impact, and operational efficiency trade-offs. What this framework does differently is it includes SHAP based interpretability (transparent justifications for all recommendations) and counterfactual scenario analysis (scenario analysis) for decision makers to understand not only optimized solution, but also reasoning behind it. We provide extensive experimental evaluations from a custom synthetic dataset, showing that the proposed methodology beats, and greatly so, a conventional baseline approach. The key improvements from this work are a 12 % reduction in cost, 32 % reduction in CO_2 emissions, 3-day reduction in lead time, among others, and improvements in supplier reliability and sustainability scores. In addition, explainability tools integrated in modern supply chains help build trust and usability and comply with more transparency decision-making standards as well. Those are the keys to creating the business foundation for creating scalable, intelligent, and responsible supply chain systems. It facilitates data driven optimization to be achieved together with actionable human centric insights, so that organisations can reach their economic and environmental aspirations.

References

[1] K. Ganga Parvathi, S. Krupalini, J. Pinnamaraju Harshitha, S. Tiwari, V. Hanisha, and S. Ishrath, “Blockchain-enabled product tokenization and CNN-based counterfeit detection for secure supply chain verification,” Macaw Int. J. Adv. Res. Comput. Sci. Eng, vol. 11, no. 1, pp. 46–57, 2025.

[2] A. Sobhan Abir, I. A. Bhuiyan, M. Arani, and M. M. Billal, Multi-Objective Optimization for Sustainable Closed-Loop Supply Chain Network Under Demand Uncertainty: A Genetic Algorithm. 2020.

[3] S. V. Mhaskey, “SCM 4.0: Navigating the impact of industry 4.0 on supply chain management through digitalization and technology integration,” Int. J. Comput. Eng. Res. Trends, vol. 11, no. 10, pp. 1–12, 2024.

[4] H. Banerjee, “Uncertainty modelling in risk-averse supply chain systems using multi-objective Pareto optimization,” arXiv [cs.AI], 2020.

[5] S. V. Mhaskey, “Integration of artificial intelligence (AI) in enterprise resource planning (ERP) systems: Opportunities, challenges, and implications,” Int. J. Comput. Eng. Res. Trends, vol. 11, no. 12, pp. 1–9, 2024.

[6] F. Olan, K. Spanaki, W. Ahmed, and G. Zhao, “Enabling explainable artificial intelligence capabilities in supply chain decision support making,” Prod. Plan. Control, pp. 1–12, 2024.

[7] K. Sadeghi, D. Ojha, P. Kaur, R. Mahto, and A. Dhir, “Explainable artificial intelligence and agile decision-making in supply chain cyber resilience,” Decision Support Systems, vol. 180, 2024.

[8] Y. Long, S. Kroeger, M. F. Zaeh, and A. Brintrup, “Leveraging synthetic data to tackle machine learning challenges in supply chains: challenges, methods, applications, and research opportunities,” Int. J. Prod. Res., pp. 1–22, 2025.

[9] M. M. Ahmed, S. M. Salauddin Iqbal, T. J. Priyanka, M. Arani, M. Momenitabar, and M. M. Billal, “An environmentally sustainable closed-loop supply chain network design under uncertainty: Application of optimization,” in Advances in Intelligent Systems and Computing, Cham: Springer International Publishing, 2021, pp. 343–358.

[10] C. P. S. Tautenhain, A. P. Barbosa-Povoa, B. Mota, and M. C. V. Nascimento, “An efficient Lagrangian-based heuristic to solve a multi-objective sustainable supply chain problem,” Eur. J. Oper. Res., vol. 294, no. 1, pp. 70–90, 2021.

[11] K. Sadeghi, “Explainable artificial intelligence and agile decision-making in supply chain cyber resilience,” Decision Support Systems, vol. 180, 2024.

[12] A. S. Abir, I. Ahmed Bhuiyan, M. Arani, and M. M. Billal, “Multi-objective optimization for sustainable closed-loop supply chain network under demand uncertainty: A genetic algorithm,” in 2020 International Conference on Data Analytics for Business and Industry: Way Towards a Sustainable Economy (ICDABI), 2020.

[13] M. Arab Momeni, V. Jain, and M. Bagheri, “A multi-objective model for designing a sustainable closed-loop supply chain logistics network,” Logistics, vol. 8, no. 1, p. 29, 2024.

[14] V. Pasupuleti, B. Thuraka, C. S. Kodete, and S. Malisetty, "Enhancing supply chain agility and sustainability through machine learning: Optimization techniques for logistics and inventory management," Logistics, vol. 8, no. 3, p. 73, 2024.
[15] E. Gutierrez-Franco, A. Polo, N. Clavijo-Buritica, and L. Rabelo, "Multi-objective optimization to support the design of a sustainable supply chain for the generation of biofuels from forest waste," Sustainability, vol. 13, no. 14, p. 7774, 2021.
[16] K. C. Hong et al., "Generative AI-powered synthetic data for enhancing predictive analytics in blood donation supply management: A comparative study of machine learning models," Int. J. Adv. Sci. Eng. Inf. Technol., vol. 15, no. 1, pp. 9–19, 2025.
[17] K. Sadeghi R., D. Ojha, P. Kaur, R. V. Mahto, and A. Dhir, "Explainable artificial intelligence and agile decision-making in supply chain cyber resilience," Decis. Support Syst., vol. 180, no. 114194, p. 114194, 2024.
[18] A. Ghodousian and E. Khorram, "Linear optimization with an arbitrary fuzzy relational inequality," Fuzzy Sets And Systems, vol. 206, pp. 89–102, 2012.