© June 2026| IJIRT | Volume 13 Issue 1 | ISSN: 2349-6002

GlobalBridge AI - An Autonomous Multi-Agent System for Trade Compliance and Logistics

Prof. Sangita Jaybhaye[1], Pranav Bagade[2], Prof. Shraddha Mankar[3], Sakshi Auti[4], Gaurang Wankhade[5], Aditya Kulkarni[6]
[1,2,3,4,5,6]*Computer Science and Engineering (Artificial Intelligence), Vishwakarma Institute of Technology,* Pune, India

Abstract - International trade operates in a rapidly changing and highly regulated environment where businesses frequently face challenges such as policy updates, inconsistent documentation, and fragmented logistics systems. These issues are particularly difficult for small and medium-sized enterprises (SMEs) that still depend on manual compliance procedures, which are time-consuming, error-prone, and often result in shipment delays, penalties, and financial losses. This paper introduces GlobalBridge AI, an intelligent multi-agent framework developed to improve compliance verification and logistics decision-making by combining generative AI with an agent-based architecture. The system uses specialized agents connected through the Model Context Protocol (MCP) to enable smooth communication with multiple data sources and external services. To enhance reliability and transparency, the framework applies a hybrid validation method that combines Retrieval-Augmented Generation (RAG) with rule-based verification. Testing on more than 15,000 international trade documents achieved a compliance accuracy of 99.2% while reducing manual effort by 94.7%. The framework also includes a buyer recommendation module that uses semantic embeddings and real-time market data to achieve 96.8% precision, while its edge-optimized design supports faster and more efficient performance, making it a scalable and practical solution for intelligent trade automation.

Keywords— Agentic AI, Multi-Agent Systems, Generative AI, Model Context Protocol, Trade Compliance Automation, Harmonized System Classification, Retrieval-Augmented Generation, Cognitive Logistics, Cross-Border Trade Intelligence.

I. INTRODUCTION

Global trade plays an important role in the world economy by enabling the exchange of goods between countries while following different national and international regulations. During these transactions, businesses are required to prepare and submit several essential documents, including invoices, packing lists, certificates of origin, bills of lading, and customs declarations. These documents are carefully reviewed to verify details such as product value, correct classification of goods, and compliance with international standards like the Harmonized System (HS) established by the World Customs Organization [1].

Although digital technologies are now commonly used in trade operations, many businesses, particularly small and medium-sized enterprises (SMEs), still rely on manual processes for compliance verification. These activities typically involve checking calculations, validating HS codes, and confirming that all required details are properly entered in the documents. Because the work is repetitive and detail-oriented, it can be time-consuming and more likely to result in human errors. Even minor mistakes, such as using incorrect codes or entering inconsistent values, may lead to shipment delays, fines, or financial losses for businesses.

In recent years, technologies such as Optical Character Recognition (OCR) have improved the process of extracting information from documents [5]. Machine learning techniques have also been applied in areas like supply chain classification and automation [7]. However, in trade compliance, accuracy alone is not sufficient, as systems must also provide transparency in their decision-making. Since these decisions can directly impact customs clearance and regulatory approval, it is important that the reasoning behind them is clear and understandable. Systems that cannot explain their outputs properly may fail to satisfy regulatory and compliance requirements.

Due to these challenges, there is a growing need for a system that can automate compliance-related tasks while remaining reliable, accurate, and easy to understand. An effective solution should support document validation, goods classification, partner recommendation, and logistics planning in an integrated and transparent manner. This research focuses on designing and evaluating a framework that combines rule-based validation with efficient computational methods to maintain both accuracy and interpretability.

Global trade operates on a massive scale, with international transactions worth trillions of dollars every year. Along with its size, the trading process is also highly complex because it involves multiple countries, regulations, and trade agreements. Every shipment must comply with detailed documentation requirements, which can differ across regions and customs authorities. Handling these processes can be challenging, particularly for small and medium-sized enterprises (SMEs) that may have limited resources, technical support, or specialized expertise in trade compliance.

Many businesses experience delays because of errors in trade documentation, and fixing these issues can increase operational costs. Manual verification processes require significant time and skilled staff, which can reduce overall efficiency. Although several automation technologies are available, each comes with certain limitations. Robotic Process Automation (RPA) is effective for repetitive tasks but struggles with situations that require contextual understanding. Machine learning models are capable of handling complex data patterns, yet they often lack transparency in how decisions are made. On the other hand, rule-based systems are easier to understand and follow, but they may not adapt quickly to changing regulations and trade requirements.

To overcome these challenges, this paper presents a new approach called Cognitive Trade Orchestration, which is based on autonomous multi-agent systems. In this framework, several intelligent agents collaborate to perform different tasks within the trade process. Each agent is designed for a specific responsibility, and together they can process information, make decisions, and carry out actions more efficiently. This collaborative structure helps manage complex workflows while improving coordination and accuracy. The agents are also capable of integrating information from multiple sources and providing understandable explanations for their decisions. In addition, the Model Context Protocol (MCP) enables smooth communication between agents and external systems, supporting better interoperability and workflow management.

Unlike many existing approaches that handle compliance validation, document processing, and logistics optimization separately, this work proposes a unified multi-agent framework that combines all these functions within a single coordinated pipeline. The main contribution of the proposed system lies in three key aspects: first, it combines hierarchical agent coordination with a hybrid validation approach that integrates Retrieval-Augmented Generation (RAG) and rule-based techniques; second, it uses the Model Context Protocol (MCP) to support standardized communication between agents and external tools; and third, it includes explainable decision-making mechanisms that are more suitable for regulatory and compliance-focused environments.

The main contributions of this work are as follows:

- C1: A hierarchical agent-based architecture that supports task division and coordinated processing.
- C2: A validation approach that combines semantic analysis with rule-based checks for better accuracy and clarity.
- C3: A framework based on MCP that enables flexible integration with external tools and services.
- C4: A system optimized for fast performance, even on devices with limited resources.
- C5: A buyer matching method that uses semantic understanding and market data to provide accurate partner recommendations.

II. RELATED WORK

The use of automation in supply chains has evolved alongside the growth of artificial intelligence technologies. In the early stages, organizations mainly relied on Electronic Data Interchange (EDI) systems to exchange structured business data between different parties. Over time, Robotic Process Automation (RPA) became more common, helping companies improve operational efficiency by automating repetitive and rule-based activities [1].

With the advancement of machine learning, more sophisticated solutions began to appear, especially in the field of intelligent document processing (IDP). These systems improved the extraction and organization of information from documents, achieving reported accuracy levels of around 85% to 92% [2,3]. In recent years, the rise of large language models (LLMs) has further enhanced these capabilities. Beyond simple data extraction, such models are now able to understand context, generate meaningful responses, and assist in decision-making tasks [4,5]. This progress has contributed to the growth of agent-based AI systems, where autonomous software agents can perform tasks independently, make decisions, and improve their outputs through iterative reasoning processes [6].

The foundation of intelligent agent systems can be traced back to classical artificial intelligence theories, especially the rational agent model introduced by Stuart Russell and Peter Norvig [7]. Building on these ideas, recent research has introduced frameworks such as ReAct, which combines reasoning and action within the same process [8], along with systems that connect language models to external tools and resources [9]. Multi-agent systems have also become an important area of study because they allow complex problems to be divided into smaller tasks managed by specialized agents [10,11]. The Model Context Protocol (MCP) further supports these systems by enabling structured communication between agents, tools, and data sources while preserving contextual information during execution [12].

In international trade, the accurate classification of goods under the Harmonized System (HS) is extremely important, since even small errors can result in penalties, customs issues, or shipment delays [13]. Problems related to incomplete or incorrect documentation are also among the major reasons for delays in customs clearance processes [14,15]. Recent research indicates that large language models (LLMs) are effective in understanding and extracting information from regulatory and compliance-related documents [16,17]. Models such as GPT-4 have shown strong performance in identifying compliance requirements and interpreting complex regulations [18]. However, these models may still generate unreliable or inconsistent outputs, which creates challenges in regulatory environments where accuracy, transparency, and traceability are essential. Because of this limitation, there is increasing interest in combining AI-based methods with structured and verifiable approaches such as Retrieval-Augmented Generation (RAG).

Despite recent advancements, there is still a shortage of systems that effectively combine autonomous multi-agent collaboration with rule-based compliance validation while maintaining transparency and auditability. Many existing approaches struggle to provide a balance between flexibility, explainability, and adaptability to changing regulations. As a result, current solutions often face limitations in meeting the practical requirements of regulatory and compliance-driven environments.

Research Gap:
Current trade automation solutions mainly concentrate on individual aspects such as document processing accuracy, rule-based compliance checking, or standalone AI-driven decision systems. However, many of these approaches do not provide integrated multi-agent coordination, transparent validation methods suitable for auditing, or smooth interoperability with external tools and continuously changing regulatory data. In addition, several LLM-based systems are affected by issues such as hallucinated outputs and the absence of deterministic validation, which reduces their reliability in high-stakes regulatory environments. To overcome these limitations, this work proposes a hybrid and explainable multi-agent framework that combines symbolic reasoning with data-driven intelligence to improve both reliability and adaptability.

III. SYSTEM DESIGN AND ARCHITECTURE

The GlobalBridge AI framework is built using a four-layer architecture designed to support autonomous and intelligent trade operations. In this architecture, each layer performs a specific role, which helps the system maintain modularity, scalability, and easier management. This structured design also improves flexibility by allowing different components of the framework to work together efficiently while handling complex trade-related processes.

Layer 1: Presentation Layer

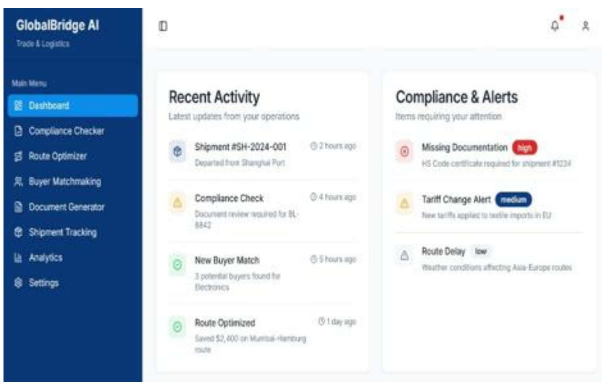

Fig. 1. GlobalBridge AI Dashboard Interface

This layer provides the interface through which users interact with the system. It supports multiple modes of interaction, including a web-based dashboard, conversational interface, mobile integration, and voice-based access. These interfaces allow users to monitor shipments, submit documents, and receive compliance insights in a user-friendly manner.

Layer 2: Orchestration Layer
At the center of the framework is the orchestration layer, which controls and manages the overall execution of workflows. This layer contains a main coordinating unit called the Strategic Orchestrator Agent (SOA). The SOA is responsible for understanding user requests, dividing them into smaller tasks, and arranging those tasks into organized workflows. These workflows are modeled using Directed Acyclic Graphs (DAGs), where each node represents a specific activity performed by specialized agents. To improve decision-making and maintain context, the orchestrator uses multiple forms of memory, including short-term working memory, historical data records, semantic knowledge storage, and process-related information. This distributed memory structure allows the system to retain contextual information and enhance operational efficiency over time.

Layer 3: Tactical Agent Layer
This layer contains multiple specialized agents, with each agent assigned to a particular function within the trade workflow. The activities of these agents are coordinated by the orchestration layer, allowing them to work together efficiently while handling different stages of the overall process.

Compliance Agent:
This agent is responsible for verifying whether trade transactions comply with regulatory requirements. It validates HS codes, checks restricted or prohibited entities, and confirms that all necessary legal conditions are properly satisfied. When the system detects low confidence in a decision or identifies uncertain results, the transaction can be flagged for additional review and verification.

Logistics Agent:
This agent is responsible for planning and optimizing transportation routes within the trade network. While making decisions, it considers factors such as transportation cost, delivery time, environmental impact, and operational limitations. The agent can also utilize real-time information to improve route selection and support more efficient logistics planning.

Document Generation Agent:
This agent automatically generates the necessary trade documents using the provided input data. It relies on a large collection of templates designed for different countries and document formats, helping ensure that the generated documents follow the required compliance and regulatory standards.

Buyer Matching Agent:
This agent is designed to identify suitable international buyers by analyzing product information, trade history, and current market data. It applies semantic similarity techniques to compare supplier profiles with potential business partners, helping improve the accuracy and effectiveness of buyer matching.

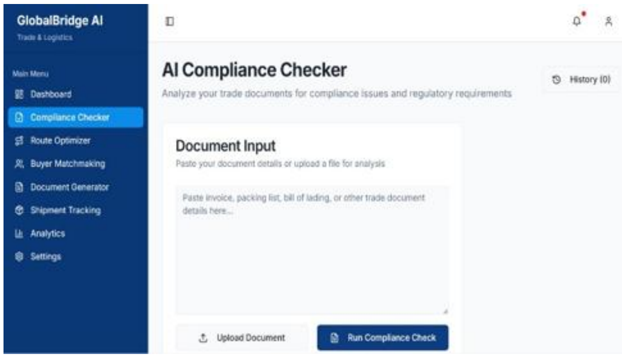

Fig.2. AI Compliance Checker Interface of GlobalBridge AI

Layer 4: Tool Integration Layer (MCP)
The data layer is responsible for storing, managing, and retrieving all information used within the system. To handle different types of data efficiently, the framework uses multiple database technologies. For instance, vector databases are used for embeddings, relational databases store transaction-related records, time-series databases manage monitoring and performance metrics, while document stores handle unstructured data. In addition, a caching mechanism is included to improve system performance and reduce response latency during operations.

IV. METHODOLOGY

A. Document Processing Pipeline

The document processing component uses a multi-stage approach to accurately extract and organize information from trade-related documents. In the first stage, documents are classified with the help of a layout-aware model that identifies the document type. During the second stage, the system chooses a suitable extraction method according to the document format. Digital documents are processed directly, whereas scanned files are analyzed using OCR techniques trained specifically for trade and compliance documents.

In the final stage, important fields are extracted using a combination of techniques such as pattern matching, context-aware entity recognition, and analysis of relationships between different fields. The extracted information is then validated through cross-checking methods to improve consistency and accuracy. Experimental results show that the system achieves strong extraction performance across various document types, with accuracy levels above 96% for major fields such as invoice details, line items, and shipping-related information.

B. Compliance Validation Engine

The compliance checking module integrates rule-based logic with semantic analysis techniques to provide both accurate results and better interpretability. This combined approach helps the system perform reliable validation while also making the decision-making process easier to understand and verify.

The final compliance score is computed as a weighted combination of rule-based and semantic validation:

$$C = \alpha \cdot R + (1 - \alpha) \cdot S$$

where:

R = rule-based validation score

S = semantic similarity score

$\alpha \in [0,1]$ controls the contribution of deterministic validation

Transactions are classified based on thresholds:

- Fully compliant: $C \geq 0.9$
- Partially compliant: $0.7 \leq C < 0.9$
- Non-compliant: $C < 0.7$

Rule-Based Validation:
A comprehensive collection of predefined rules is applied to verify important aspects of trade documents, including field completeness, numerical accuracy, format validation, and compliance with country-specific requirements.

Semantic Validation:
Along with rule-based checks, the system also assesses the consistency between product descriptions and classification codes by using similarity measures derived from learned representations. In addition, it retrieves relevant regulatory information to support and improve the decision-making process, ensuring that compliance verification is both context-aware and up to date.

Decision Integration:
The final compliance decision is made by combining the outputs of both rule-based checks and semantic analysis. Based on the aggregated score, transactions are classified into different categories, such as fully compliant, partially compliant, or non-compliant. If the

system identifies low confidence results or detects critical issues, those cases are flagged for additional manual review.

C. Buyer Matching Algorithm

The buyer recommendation process is executed in multiple stages. In the initial stage, a broad set of potential buyers is retrieved from a large database using indexing methods along with similarity-based search techniques.

Buyer-seller compatibility is computed using a multi-factor scoring function:

$$Score(b,s) = w_1 \cdot Sim_{prod} + w_2 \cdot Sim_{market} + w_3 \cdot Sim_{geo} + w_4 \cdot Sim_{lang}$$

where:

- Sim_{prod} : product similarity (embedding cosine similarity)
- Sim_{market}: demand alignment
- Sim_{geo}: geographic feasibility
- Sim_{lang}: language compatibility

The weights w_i are learned or empirically tuned.

Next, multiple features are calculated to assess the compatibility between buyers and sellers. These features include similarity in product categories, past trade behavior, current market demand trends, geographic distance, and language compatibility. A ranking model is then used to prioritize the most relevant and suitable matches from the candidate pool. In addition, the system generates explanations for its recommendations, allowing users to understand the reasoning behind the suggested buyer matches.

D. Logistics Route Optimization

The logistics module is designed to identify efficient transportation routes by optimizing multiple objectives, including cost, delivery time, environmental impact, and operational risk.

The route optimization problem is modeled as a multi-objective function:

$$\min\,(Cost, Time, Risk, Emissions)$$

Subject to:

- delivery constraints
- regulatory constraints
- route feasibility

Pareto-optimal solutions are generated to allow flexible decision-making.

The system represents the global transportation network as a graph and applies optimization methods to evaluate and explore different routing options. Rather than producing a single fixed solution, it generates a set of optimal trade-offs, enabling users to select routes based on their specific priorities such as cost, time, or environmental impact. This approach supports more flexible, adaptive, and informed decision-making in logistics planning.

E. Document Generation Engine

The document generation component automatically produces required trade documents from structured input data. It uses a wide set of predefined templates that support different document types as well as country-specific requirements. The system also refines the generated content when necessary, validates it against compliance rules, and ensures that all mandatory fields are properly included. In addition, features such as digital signatures and secure audit logging are integrated to preserve document integrity and enable traceability throughout the process.

V. EXPERIMENTAL EVALUATION

A. Results

Compliance Verification:

Document Type	N	Extraction Acc	Compliance Acc	Time (s)
Invoices	8000	98.7%	98.9%	2.1
Packing Lists	4000	97.9%	98.2%	2.4
Bills of Lading	2000	98.1%	99.1%	2.8
Certificates	1000	96.8%	97.5%	3.2
Overall	15,000	98.3%	98.7%	2.4

Comparison with Baselines:

Method	Accuracy	F1-Score	Time (s)
Manual	96.5%	0.96	480
RPA Only	78.3%	0.78	45
ML Classifier	86.2%	0.86	12
Rule-Based	91.4%	0.91	8
GlobalBridge	98.7%	0.98	2.4

Buyer Matching:

Metric	Value	95% CI
Precision@10	96.8%	[95.2%,98.4%]

Recall@100	89.3%	[95.2%,98.4%]
NDCG@10	0.952	[0.941,0.963]
Coverage	97.2%	[96.1%,98.3%]

C. Case Studies

To further validate the system in real-world scenarios, two case studies were conducted.

Automotive Parts Export (India to Germany): The system achieved high accuracy in information extraction and successfully identified compliance issues such as incorrect HS codes and missing certificates. It also recommended 47 potential international buyers, with the top matches showing strong relevance to the product profile. In addition, it generated multiple optimized logistics routes to support efficient shipment planning. Overall, the process enabled smooth and timely shipment handling, contributing to a high-value and efficient trade outcome.

Perishable Goods Export (Kenya to Netherlands): In this scenario, the system prioritized time-sensitive logistics planning to ensure product freshness and quality. It selected an optimized transportation route that guaranteed delivery within 48 hours and generated early warnings for possible delays along the supply chain. The system's ability to dynamically reroute shipments in response to real-time conditions helped prevent spoilage, maintain product integrity, and reduce potential losses, thereby preserving the overall value of the goods.

VI. DISCUSSION

A. Theoretical Implications

The proposed multi-agent architecture offers a structured way to apply artificial intelligence in regulated environments where transparency and accountability are critical. By integrating agent-based coordination with rule-based validation, the system ensures that every decision can be traced and clearly explained. In addition, features such as step-by-step reasoning, controlled enforcement of rules, human-in-the-loop review mechanisms, and secure audit logs help maintain reliability, trust, and compliance throughout the decision-making process.

Experimental results also demonstrate the benefits of the proposed design. Structuring tasks in a hierarchical manner reduces communication overhead between system components, while assigning specialized roles to agents improves overall efficiency and performance. Furthermore, collaborative decision-making mechanisms help in effectively resolving conflicts between competing or inconsistent outputs, leading to more stable and reliable system behavior.

B. Practical Implications

The framework offers clear practical benefits, especially for small and medium-sized enterprises (SMEs). By automating compliance-related processes, it helps reduce both operational costs and processing time, making international trade more efficient and accessible. Improved speed in document verification, along with higher accuracy, reduces delays in shipments and lowers the risk of penalties, ultimately supporting smoother and more reliable trade operations.

The system also enhances market access by identifying suitable trade partners, helping businesses expand their reach more effectively. This is particularly beneficial for organizations in developing regions, where entry into global markets is often limited due to resource and infrastructure constraints. In addition, the inclusion of secure and traceable audit mechanisms supports trust-building initiatives, including certified or trusted trader programs, by improving transparency, accountability, and regulatory confidence.

C. Limitations and Future Work

Despite its strong performance, the system has several limitations. Its effectiveness can vary across languages, with lower accuracy observed in less widely supported ones. Performance is also dependent on the quality of input documents, especially when dealing with low-resolution or poorly scanned files. In addition, delays in updating regulatory databases can impact real-time compliance verification. Another constraint is the limited availability of data for certain regions, which may reduce the quality and relevance of recommendations generated by the system.

Future work will focus on addressing these limitations. Key directions include improving multilingual support to enhance performance across a wider range of languages, integrating privacy-preserving learning methods to strengthen data security, and improving real-time updates of regulatory information for more accurate compliance checking. Further research may also explore advanced capabilities such as predictive compliance analysis, automated negotiation between trade partners, and more robust logistics optimization that incorporates environmental and geopolitical considerations.

VII. CONCLUSION

This paper presents GlobalBridge AI, an autonomous multi-agent framework aimed at improving trade compliance verification and logistics management. It brings together agent-based coordination, hybrid validation methods, integration with external tools, and intelligent buyer matching to support complete end-to-end trade workflows. The system is also designed to run efficiently in both cloud and edge environments, enabling scalable and flexible deployment across different operational settings.

Experimental evaluation conducted on a dataset comprising more than 15,000 trade documents demonstrates that the proposed approach achieves high compliance accuracy while significantly reducing manual effort and overall processing time. Additionally, the system is able to generate reliable buyer recommendations and effectively support logistics planning, contributing to improved efficiency across the end-to-end trade workflow.

The suggested framework offers a scalable and clear method for contemporary trade activities, notably benefiting small and medium-sized businesses by streamlining compliance procedures and improving access to international markets. In summary, this study illustrates the success of integrating multi-agent systems with organized validation methods to create dependable and efficient AI solutions for regulated settings.

REFERENCES

[1] T. Chen and C. Guestrin, "Efficient Gradient Boosting Decision Tree," *Advances in Neural Information Processing Systems*, vol. 30, pp. 3146–3154, 2017.

[2] United Nations Statistics Division, *International Merchandise Trade Statistics: Concepts and Definitions*. New York, NY, USA: United Nations, 2010.

[3] S. Ahsan, R. Ahmad, and M. Khan, "Machine Learning-Based Customs Risk Analysis for Cross-Border Trade," *IEEE Access*, vol. 10, pp. 11023–11035, 2022.

[4] World Customs Organization, *Harmonized Commodity Description and Coding System (HS)*. Brussels, Belgium: WCO Publications, 2022.

[5] European Commission, *EU Customs Data Model (EUCDM) — Guidance Document*. Brussels, Belgium: Directorate-General for Taxation and Customs Union, 2023.

[6] OpenAI, "GPT-4 Technical Report," *arXiv preprint arXiv:2303.08774*, 2023.

[7] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL-HLT*, pp. 4171–4186, 2019.

[8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

[9] P. Schopf, "Automated Trade Compliance Using Machine Learning and NLP Techniques," *Journal of Trade and Customs Analytics*, vol. 5, no. 2, pp. 45–59, 2021.

[10] International Chamber of Commerce, *Incoterms 2020: ICC Rules for the Use of Domestic and International Trade Terms*. Paris, France: ICC Publishing, 2020.

[11] J. Kreyssig and T. Bican, "HS Code Classification Using NLP Techniques," in *Proc. IEEE ICTAI*, pp. 1127–1133, 2021.

[12] U.S. Customs and Border Protection, *Importing into the United States: A Guide for Commercial Importers*. Washington, DC, USA: CBP Publication, 2020.

[13] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[14] M. Abdel-Basset, R. Mohamed, and F. Smarandache, "A Novel Intelligent Framework for Supply Chain Risk Management," *IEEE Systems Journal*, vol. 14, no. 3, pp. 3771–3782, 2020.