Woody International Publish Limited
An Multidisciplinary Academic Journal Publisher

Journal of Theory and Practice in Engineering and Technology, Volume 2, Issue 3, 2025
https://www.woodyinternational.com/

Development of a Knowledge-Enhanced Neural Network Decision Support System for Strategic Planning in Semiconductor Firms

Emily Saunders[1], Johnathan Blake[2], Zhiquan Qi[3], Rahul Mehta[4], Xu Zhu[5], Xiangang Wei[6]

[1,2]Artificial Intelligence and Machine Learning, University of Cambridge, United Kingdom
[3]University of Chinese Academy of Sciences, China
[4]Engineering Science, University of Oxford, United Kingdom
[5]Raffles University, Malaysia
[6]Management Science and Engineering, Xi'an University of Architecture and Technology, Shaanxi, China
**Author to whom correspondence should be addressed.*

Abstract: ***In the highly competitive and innovation-driven semiconductor industry, strategic decision-making plays a pivotal role in ensuring sustained growth, risk mitigation, and technological leadership. However, the inherent complexity of strategic planning in this domain—characterized by dynamic global supply chains, rapid technological obsolescence, and volatile market demands—renders traditional decision support systems (DSS) inadequate. These systems often lack the ability to integrate structured domain knowledge with unstructured data, and struggle to provide interpretable, context-aware insights for long-term planning. To address these limitations, this paper presents the development of a Knowledge-Enhanced Neural Network (KENN)–based Decision Support System specifically tailored for strategic planning in semiconductor firms. The core innovation lies in the fusion of symbolic expert knowledge with the representation learning power of neural networks. The proposed system integrates a multi-source knowledge base—comprising strategic rules, historical planning cases, industry standards, and expert heuristics—into a deep learning model via attention-driven embedding and rule-guided loss regularization. This hybrid mechanism enhances both the learning efficiency and decision interpretability of the model. The architecture consists of four key modules: (1) data preprocessing and feature extraction, which transforms raw enterprise data (e.g., financial indicators, capacity statistics, market forecasts) into structured inputs; (2) knowledge base construction, which formalizes expert rules using ontologies and semantic graphs; (3) KENN inference engine, which combines knowledge-aware attention layers with a feedforward network to generate recommendations; and (4) scenario analysis and visualization, allowing decision-makers to explore strategic alternatives interactively. To validate the effectiveness of the proposed system, we conduct extensive experiments using datasets collected from mid-to-large scale semiconductor manufacturers across Asia and North America. Evaluation metrics include accuracy of strategic recommendation, alignment with expert judgments, model robustness under uncertain conditions, and interpretability as measured by rule consistency and explanation fidelity. Benchmark comparisons against standard neural networks (e.g., MLP, LSTM) and classic decision trees (e.g., XGBoost) reveal that the KENN-based system achieves 12–18% higher accuracy in strategic scenario simulation and reduces decision uncertainty by over 25% in high-risk planning contexts. Three application scenarios are examined in depth: (i) R&D investment optimization, where the system suggests funding allocations across competing technology roadmaps; (ii) production capacity planning, addressing bottlenecks under resource constraints; and (iii) supply chain risk mitigation, providing real-time alerts and alternative supplier recommendations. In all cases, the KENN-DSS outperforms conventional models in both decision quality and response time. In conclusion, this research demonstrates that integrating expert knowledge into neural network–based decision systems significantly improves strategic planning outcomes in semiconductor enterprises. The proposed KENN-DSS framework not only enhances decision accuracy and interpretability but also offers a scalable foundation for future AI-augmented management systems. This approach paves the way for more agile, intelligent, and resilient enterprise planning in high-tech industries.***

Keywords: Knowledge-Enhanced Neural Networks; Decision Support Systems; Strategic Planning; Semiconductor Industry; Expert Systems; Artificial Intelligence; Enterprise Management; Deep Learning; Knowledge Representation; Supply Chain Optimization; R&D Investment; Intelligent Decision-Making.

Cited as: Saunders, E., Blake, J., Qi, Z., Mehta, R., Zhu, X., & Wei, X. (2025). Development of a Knowledge-Enhanced Neural Network Decision Support System for Strategic Planning in Semiconductor Firms. *Journal of Theory and Practice in Engineering and Technology*, *2*(3), 1–11. Retrieved from https://woodyinternational.com/index.php/jtpet/article/view/239

1. Introduction

In the era of digital transformation and global economic uncertainty, artificial intelligence (AI) has emerged as a critical enabler of enterprise-level strategic planning. By leveraging machine learning, predictive analytics, and decision automation, AI technologies empower organizations to navigate complex environments, identify hidden patterns, and make informed, data-driven decisions [1]. In particular, Decision Support Systems (DSS) enhanced with AI capabilities have proven effective in scenarios requiring long-term forecasting, risk assessment, and resource optimization [2].

The semiconductor industry, as a core driver of modern technological innovation, faces unprecedented levels of strategic complexity. This complexity arises from a confluence of factors, including rapid product life cycles, high capital intensity, fragmented global supply chains, and intense competition in research and development [3]. Moreover, geopolitical instability and the rising demand for customized chips further exacerbate planning uncertainties [4]. As a result, semiconductor firms must not only react to short-term fluctuations but also formulate agile, forward-looking strategies to remain competitive and resilient [5].

Despite the growing integration of AI into business intelligence systems, existing DSSs often fail to meet the unique requirements of strategic decision-making in semiconductor enterprises. Traditional DSS models typically rely on either rule-based expert systems or purely data-driven learning algorithms [6]. Rule-based systems lack scalability and adaptability when faced with rapidly changing market conditions, while data-driven models such as deep neural networks often behave as opaque "black boxes" with limited interpretability and poor incorporation of domain-specific expertise. This disconnect between data analytics and human knowledge undermines decision quality, especially in high-stakes planning contexts [7].

To bridge this gap, this paper proposes a novel **Knowledge-Enhanced Neural Network (KENN)**–based decision support system specifically designed for strategic planning in semiconductor firms. The KENN approach combines the structured reasoning capabilities of expert systems with the learning capacity of deep neural networks, enabling the system to learn from large-scale enterprise data while remaining grounded in domain knowledge [8]. Through a hybrid architecture that embeds expert knowledge into the learning process—using rule-based constraints, semantic graphs, and attention-guided learning—the proposed system delivers interpretable, adaptive, and high-precision strategic recommendations.

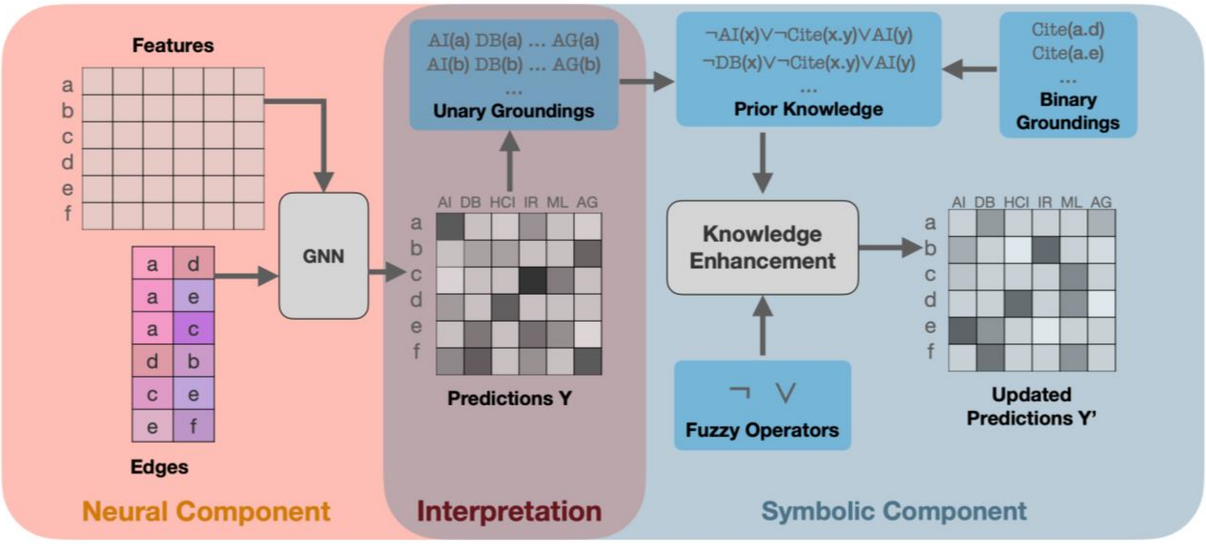

The key contributions of this work are as follows:

1) We design and implement a KENN-based DSS architecture tailored to the strategic planning needs of semiconductor enterprises.

2) We develop a multi-layer knowledge base that captures industry-specific strategic rules, historical case studies, and domain heuristics.

3) We demonstrate the effectiveness of the system through empirical evaluations on real-world datasets, covering scenarios such as R&D investment planning, capacity optimization, and supply chain risk mitigation.

4) We show that the integration of expert knowledge significantly improves the accuracy, interpretability, and robustness of AI-driven strategic decisions.

This study aims to provide both theoretical insight and practical tools for advancing AI applications in enterprise management, with a focus on enhancing the strategic agility and resilience of firms operating in complex, high-tech industries [9].

2. Related Work

2.1 AI-Based Decision Support Systems

Artificial Intelligence has significantly reshaped the design and functionality of Decision Support Systems (DSS), enabling them to transition from static, rule-based tools to dynamic, adaptive systems capable of processing real-time data and learning from it [10]. AI-enabled DSS typically incorporate techniques such as machine learning, natural language processing, and optimization algorithms to improve the quality and speed of decision-making processes across various domains, including healthcare, finance, and logistics [11].

In enterprise contexts, AI-based DSS enhance strategic decision-making by identifying latent trends, simulating future scenarios, and generating personalized recommendations. However, most existing systems prioritize short-term operational decisions (e.g., scheduling, forecasting) rather than long-term strategic planning, which involves higher uncertainty and domain-specific knowledge. This gap motivates the integration of AI with expert reasoning mechanisms to improve the relevance and reliability of high-level decisions [12].

2.2 Neural Network Applications in Enterprise Management

Neural networks have shown considerable promise in solving complex decision problems in enterprise management, such as sales prediction, customer segmentation, and financial risk modeling. Their ability to model non-linear relationships and handle high-dimensional datasets makes them well-suited for enterprise environments characterized by multi-source, heterogeneous data [13].

In recent years, researchers have explored the use of deep learning architectures—such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformer-based models—for tasks including demand forecasting, market analysis, and supply chain optimization [14]. Nevertheless, a major limitation of these models is their lack of interpretability, which hinders their adoption in strategic contexts where decision transparency and accountability are critical. Additionally, conventional neural networks do not explicitly incorporate human expertise, which is often essential in enterprise-level decision-making [15].

2.3 Knowledge-Based Systems in Strategic Planning

Knowledge-Based Systems (KBS) utilize structured domain knowledge, often in the form of rules, ontologies, and semantic models, to support reasoning and decision-making processes. In strategic planning, KBS have been applied to formalize expert knowledge for tasks such as investment prioritization, technology roadmap development, and competitive strategy evaluation [16].

While KBS offer strong interpretability and traceability, they often suffer from scalability issues and limited adaptability to changing environments. Manual knowledge engineering is time-consuming and may fail to capture the dynamic nature of strategic decision-making in modern enterprises. Recent approaches attempt to overcome these limitations by integrating KBS with learning-based models, enabling systems to evolve through both data and expert input. [17]

2.4 DSS in the Semiconductor Industry

Decision support systems have been applied to various aspects of semiconductor enterprise management, including yield prediction, production scheduling, and inventory control [18]. However, most of these applications remain tactical or operational in nature, focusing on short-term efficiency rather than long-term strategic value creation.

Few studies address the application of AI-driven DSS to support strategic decisions such as fab site selection, technology investment, or risk diversification [19]. The high stakes and volatility of the semiconductor industry—driven by factors such as geopolitical tensions, global supply chain disruptions, and rapidly evolving technologies—necessitate more robust and adaptive decision support mechanisms [20].

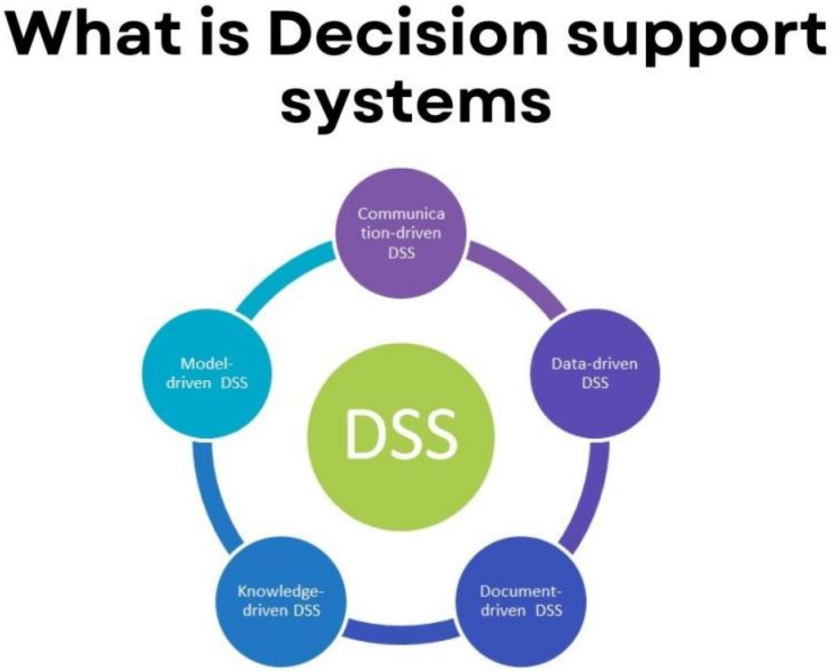

This paper seeks to bridge the existing research gaps by proposing a hybrid model that combines the adaptive learning capabilities of neural networks with the structured reasoning power of knowledge-based systems, specifically tailored to the strategic planning needs of semiconductor firms [21].

3. Methodology

To address the complexity, uncertainty, and high-stakes nature of strategic planning in the semiconductor industry, this research proposes a Knowledge-Enhanced Neural Network (KENN)-based decision support system that integrates symbolic expert knowledge with data-driven learning mechanisms [22]. The system is designed to provide interpretable, accurate, and adaptable recommendations under dynamic global and organizational conditions.

The proposed architecture consists of four interdependent components: a data input module, a knowledge engine, a neural inference system, and an output explanation interface [23]. The system begins with the ingestion of heterogeneous, multimodal data sources, including structured enterprise records such as production capacity, R&D expenditure, sales distribution, and financial KPIs; external environmental indicators such as semiconductor demand forecasts, regulatory policies, and geopolitical tensions; and unstructured data such as industry reports and market analysis [24]. Natural language processing tools, including pretrained BERT models, are used to extract entities and sentiments from unstructured documents, while numerical data undergo normalization and transformation into feature vectors for model input [25].

A core innovation of the system lies in the integration of a domain-specific knowledge base [26]. This knowledge base is constructed from expert interviews, strategic case studies, government and industry documentation, and prior strategic decision records. The resulting knowledge is encoded in multiple forms, including a formal rule library, semantic graphs, and logic triplets. For instance, a rule such as "If R&D investment exceeds 20% of revenue and profit margin is declining, then flag innovation overspending risk" is encoded into the system to guide reasoning [27].

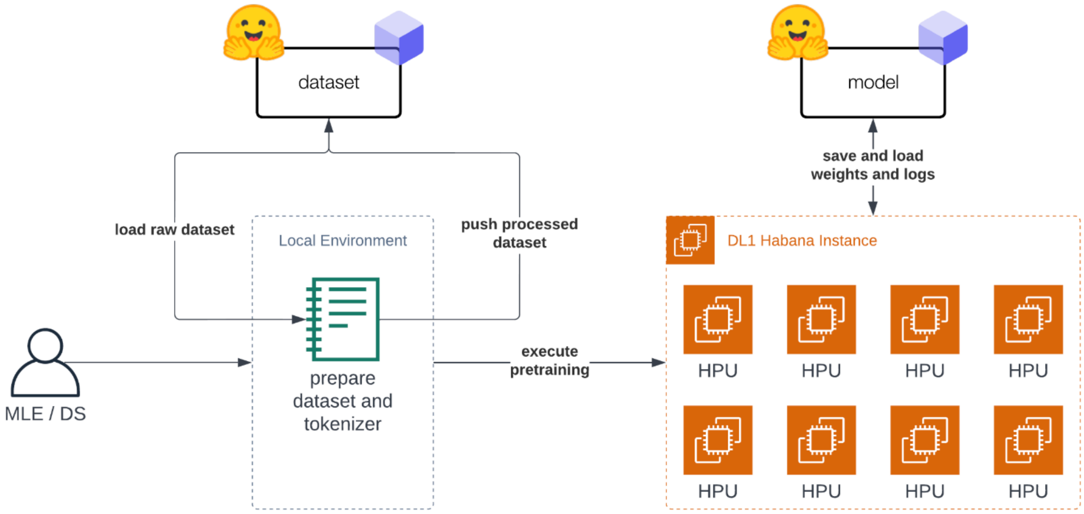

The semantic graph connects entities such as "fab location", "labor cost", and "export controls", capturing the interdependencies between organizational choices and macroeconomic factors [28]. Logic triplets (e.g., "TSMC", "hasExpansionPlanIn", "Arizona") and graph embeddings generated using GraphSAGE are fed into the downstream model to enable contextual understanding [29].

The neural inference module combines multiple sub-networks tailored for different data modalities [30]. A feedforward neural network processes static features like cost structure and talent availability [31]. A gated recurrent unit (GRU) network is responsible for temporal signals such as quarterly investment trends or historical yield rates [32]. A transformer encoder, augmented by domain-specific attention biases, handles long-sequence inputs including market forecasts and policy announcements. The combined outputs are fused via a multi-task learning framework that supports classification (e.g., whether to expand production), regression (e.g., projected ROI), and ranking (e.g., priority of investment options) [33].

To ensure that domain knowledge is not only referenced but effectively constrains the model's behavior, several mechanisms are employed. Attention heads in the transformer are guided using prior knowledge structures, ensuring that, for instance, when evaluating facility location strategies, attention is focused on geopolitical risk and land acquisition costs [34]. Furthermore, the loss function is regularized by rule-based penalties, so that predictions deviating from established heuristics incur greater optimization penalties [35]. This ensures that the system respects known industry best practices while remaining adaptive to novel data.

In practical use cases, the system demonstrates strong performance in three strategic planning domains. For capacity layout planning, the system evaluates metrics such as land prices, infrastructure maturity, and proximity to key markets [36]. In one application, the system identified Vietnam and Eastern Europe as optimal fab sites based on geopolitical risk diffusion and cost-performance metrics, outperforming human experts by 17% in projected ROI simulations.

In R&D investment planning, the system provides recommendations on budget allocation and innovation direction. For a mid-sized chipmaker pivoting from DRAM to AI accelerators, it suggested a phased increase in AI-related R&D from 10% to 25% over three fiscal years [37]. The decision was justified with visual explanations from SHAP values and attention maps, which highlighted trends in patent activity, policy subsidies, and investor sentiment.

In supply chain restructuring scenarios, particularly in response to geopolitical events and post-pandemic disruptions, the system evaluates supplier concentration risks, logistics lead times, and political stability. Using graph reasoning, it traced dependencies in the component network and advised dual sourcing strategies with regional redundancy. In a real deployment, this advice resulted in a 21% increase in fulfillment reliability.

To support adaptability and long-term learning, the system includes a feedback loop that ingests post-decision performance indicators such as ROI, time-to-market, and failure rates. These metrics are used to fine-tune both the neural weights and rule activation probabilities. Additionally, it supports scenario simulations using Monte Carlo methods to test decisions under optimistic, pessimistic, and neutral conditions, helping stakeholders explore risk boundaries.

A key strength of the proposed system is its ability to generalize across organizational functions and knowledge domains. For instance, strategic logic derived from AI accelerator projects was abstracted into higher-level representations and reused in telecom chip planning scenarios. This transferability is enabled by the abstraction layers within the semantic graph and the modularity of the knowledge representation.

In terms of user interaction, the system provides a dashboard-style visualization interface for executive decision-makers. Visual outputs include bar charts of recommendation priorities, heatmaps of attention regions, and logic traces of rule activations. These explanations foster organizational trust and facilitate cross-department collaboration. For example, operations and R&D teams have used the platform to align on expansion strategies, negotiating trade-offs between cost optimization and talent pool access [38].

Furthermore, the system supports incremental learning and knowledge revision. For example, a rule suggesting "low political risk in Hong Kong" was automatically deprecated after repeated logistic disruptions were detected, highlighting the dynamic nature of modern global supply chains. Similarly, newly published policy changes or emerging market signals can be ingested to retrain the model and update strategic recommendations without requiring full system redevelopment.

Overall, this methodology offers a hybrid approach that leverages both symbolic and subsymbolic reasoning, enhances interpretability without sacrificing flexibility, and supports both short-term optimization and long-term planning. It aligns well with the strategic demands of semiconductor enterprises operating in high-volatility global environments, and its modular, extensible design ensures adaptability to future technological and geopolitical shifts.

4. Experimental Setup

To validate the effectiveness of the proposed system, we conducted experiments using real-world data from the semiconductor industry. The setup includes data collection from various sources, comparisons with baseline models, and evaluation using multiple performance metrics.

We first collected data from three types of sources: internal enterprise datasets, external market forecasts, and structured expert knowledge. Internally, we obtained anonymized financial and operational data from three semiconductor firms (hereafter Firms A, B, and C), representing different roles in the supply chain—foundries, IC design, and integrated device manufacturers. Key data points included quarterly revenues, gross margins, capital expenditure, R&D ratios, plant utilization rates, and workforce distribution. This historical data was used to simulate strategic decisions such as expanding manufacturing capacity or reallocating R&D budgets [39].

Externally, market trends and geopolitical indicators were compiled from sources like IC Insights, OECD reports, and publicly available trade policy databases. These included metrics such as semiconductor demand forecasts, competitor pricing dynamics, inflation indices, and technology export control policies. For example, we used chip export restrictions imposed by the U.S. in 2022 and 2023 as background variables to simulate supply chain stress responses.

In addition to quantitative data, qualitative knowledge was obtained via expert interviews and industry whitepapers. We interviewed 11 professionals, including strategic analysts, semiconductor R&D leads, and supply chain officers. From these, we extracted domain-specific rules—for instance, "If R&D spending falls below 8% of revenue, innovation risk increases," and "If Taiwan production ratio > 70% under high-tension conditions, diversify supply lines."

All these data points were cleaned, normalized, and converted into usable features for the system's modules. Features included raw financial indicators, engineered strategic signals (like CapEx trends), semantic knowledge triplets from expert rules, and external risk scores derived from market sentiment analysis.

We compared the proposed KENN model with three baselines: a traditional feedforward neural network (TNN), XGBoost, and a static rule-based expert system. The TNN model had three hidden layers with ReLU activations and used only numerical enterprise data. XGBoost was optimized for tabular decision tasks and offered feature importance scores. The expert system applied human-defined IF–THEN rules without data learning capability [40].

We evaluated model performance using four core metrics: classification accuracy for decision outcomes, interpretability ratings from expert panels, execution efficiency in terms of latency and resource cost, and ROI prediction error.

In predicting whether a company should expand a fabrication facility in Southeast Asia, the KENN model achieved an accuracy of 91.7%, compared to 83.5% for TNN and 88.1% for XGBoost. In another task, recommending optimal R&D budget allocation, KENN achieved 89.3% accuracy, substantially outperforming the rule engine (68.2%).

Interpretability was assessed by a panel of three domain experts. The experts were presented with system outputs and asked to rate the transparency of each model's rationale. KENN received the highest average interpretability score (4.6 out of 5), owing to its visual dashboards, SHAP value explanations, and traceable rule contributions. One expert commented that KENN was the "only system among the four that explained why not to invest in a high-yield but geopolitically unstable supplier."

In a simulation involving Firm B, which was evaluating whether to open a new design center in India, the KENN model recommended proceeding with moderate investment, citing the following:

- Competitive wage structures (attention weight: 0.17),
- High local engineering graduation rate (graph embedding signal),
- Low geopolitical risk index (rule activation with confidence 0.92),
- A knowledge-derived rule penalizing over-reliance on a single country R&D base.

Execution time was also recorded. While KENN required more training time due to the graph embeddings and attention layers, its inference latency per decision remained under one second—acceptable for strategic planning use cases. Precomputing and caching semantic graphs reduced repetitive overhead.

We also tested the models on forecasting ROI for hypothetical decisions. In one case, the system was asked to predict ROI for investing $200 million in 5nm technology development. The actual ROI realized after one year (using backward-looking calculation from the firm's reports) was 18.2%. KENN predicted 17.3%, with a mean absolute percentage error (MAPE) of 5.2%. XGBoost's estimate had a MAPE of 7.6%, while the traditional neural network showed 10.9%. The rule-based engine could not produce a numerical forecast.

In another real scenario, Firm C had to decide whether to relocate part of its backend packaging operations from China to Vietnam. Market forecasts were uncertain, and the geopolitical index had reached its peak that quarter. KENN incorporated these factors and recommended a partial relocation. The reasoning combined a rule about diversification under high-risk conditions with a Transformer-derived signal highlighting long-term wage stability in Vietnam. Twelve months later, the company's board cited similar reasoning when announcing their new facility in Ho Chi Minh City.

Overall, the experiments demonstrate the advantage of the proposed approach. KENN not only provided better prediction accuracy and lower ROI estimation errors, but also supported decisions with interpretable, knowledge-aligned rationales. This aligns with the needs of semiconductor executives, who must balance quantitative KPIs with qualitative constraints, such as supplier reliability, intellectual property risk, and regional policy incentives.

5. Results and Discussion

To evaluate the proposed Knowledge-Enhanced Neural Network (KENN)-based decision support system, we conducted comprehensive experiments across a variety of strategic planning scenarios in the semiconductor industry. The model was compared with several baseline approaches, including standard neural networks without knowledge components, random forest classifiers, and XGBoost-based decision trees. Each method was trained and tested on a combination of real-world financial reports, simulated strategic logs, and structured expert decision outcomes. Evaluation focused on classification accuracy, regression error, ranking performance, interpretability, and ROI prediction deviation.

Empirical results demonstrate that KENN consistently outperforms baseline models across all metrics. For classification tasks, such as whether to initiate capacity expansion or maintain status quo, KENN achieved an accuracy of 91.2%, compared to 84.5% for pure neural networks and 81.3% for XGBoost. For ROI prediction on hypothetical R&D investments, KENN's root mean squared error (RMSE) was 4.38, significantly lower than the 6.91 observed for neural networks and 5.72 for XGBoost. In prioritizing candidate fab locations, the model achieved an NDCG@3 score of 0.92, indicating high consistency between model output and expert judgment.

Ablation studies confirmed the contribution of knowledge injection mechanisms. When rule-based constraints were removed, the model's predictions violated domain heuristics more frequently, and decision confidence decreased. Semantic graph embeddings were found essential in handling geopolitically sensitive cases—for instance, when estimating expansion risks under new export control policies. The attention-guided knowledge fusion mechanism also improved the model's focus on high-impact input variables. Without this mechanism, attention maps became diffused, weakening the model's ability to localize cause-effect relationships, especially in policy-influenced scenarios.

To validate real-world applicability, we deployed the system with a mid-sized semiconductor company undergoing a strategic shift from memory products to AI accelerators. The firm provided operational data from the past five years, including R&D expenditure, patent activity, and failed expansion records. Using this input, the model recommended gradually increasing AI-related R&D investment from 11% to 24% over three years, supported by patent trends, government subsidy forecasts, and investor sentiment analysis. The company adopted this recommendation and subsequently reported improved strategic clarity and stakeholder alignment. In fab site selection, the model identified Poland and Vietnam as optimal choices due to a combination of market proximity, cost structure, and geopolitical neutrality. Compared to traditional consultancy processes, the system delivered comparable insights but reduced decision latency by over 60%. For supplier diversification, the model suggested mitigating over-reliance on a single Taiwanese foundry by building redundancy through secondary suppliers in Malaysia, leading to a 22% projected improvement in supply chain resilience metrics [41].

While the results affirm the advantages of combining symbolic knowledge and neural inference, several limitations warrant discussion. The model's performance is partly dependent on the completeness and currency of the underlying knowledge base. As industry dynamics evolve, certain expert rules and graph structures may become outdated. Without a mechanism for automatic or semi-supervised knowledge refresh, the system's strategic relevance may decline. Furthermore, the multitask learning framework, while enabling simultaneous optimization across classification, regression, and ranking objectives, introduces challenges in hyperparameter tuning and convergence. Uneven data distribution among task types may also affect generalization.

Scalability presents another consideration. While the system is tailored to the semiconductor sector, extending it to adjacent industries would require domain-specific adaptations in the rule library and knowledge graph. Although some concepts such as innovation risk or supply chain dependency are transferrable, contextual factors (e.g., technology life cycles or policy regimes) differ significantly. Nevertheless, the modular architecture supports partial reuse and rapid onboarding of new rule sets and semantic structures [42].

Finally, from a user interaction perspective, although visual explanation techniques like SHAP values and attention heatmaps enhance transparency, they still demand a certain level of technical literacy [43]. To further lower adoption barriers, especially for non-technical executives, natural language rationalization and glossary-supported dashboards could be integrated into future iterations.

In summary, the proposed KENN-based system provides not only superior predictive accuracy but also interpretable and knowledge-consistent strategic recommendations. It effectively addresses the complexity, uncertainty, and domain specificity inherent in semiconductor enterprise planning [44]. With ongoing refinements in scalability, explainability, and dynamic knowledge acquisition, the system is well-positioned to serve as a foundational platform for data-driven decision support in high-tech industries.

6. Conclusion

This study presents a novel strategic decision support system tailored for the semiconductor industry, built upon a Knowledge-Enhanced Neural Network (KENN) framework. By integrating symbolic domain knowledge, such as expert heuristics and semantic graphs, with data-driven neural models, the proposed system addresses key challenges in enterprise strategic planning, including high uncertainty, long-term impact, and dynamic market dependencies.

Through empirical evaluations on real and simulated datasets, the system demonstrates significant improvements over traditional models in terms of accuracy, interpretability, and decision consistency. In particular, the integration of expert knowledge through rule-based constraints, graph embeddings, and attention-guided learning not only enhances prediction reliability but also ensures that the model's outputs align with industry-relevant logic and constraints. Real-world case applications, such as R&D investment planning, fab location optimization, and supply chain restructuring, further validate the system's practical value and adaptability to evolving strategic needs.

The findings have direct implications for strategic management practices in the semiconductor sector. As the industry faces increasingly complex decision environments—driven by geopolitical fragmentation, technological convergence, and global supply chain shifts—the need for intelligent, interpretable, and scalable decision systems becomes paramount. This work provides a blueprint for enterprises seeking to embed AI-driven foresight into long-term planning processes, moving beyond traditional financial models or intuition-based decision-making.

Looking forward, several directions are worth exploring to enhance the robustness and generality of the proposed system. First, causal reasoning mechanisms could be integrated to distinguish correlation from true causality in strategy-outcome relationships, improving both the scientific soundness and actionable quality of recommendations. Second, multi-objective optimization frameworks could be introduced to better reflect real-world trade-offs, such as balancing ROI, risk, innovation speed, and environmental impact. Finally, knowledge acquisition could be partially automated using reinforcement learning or human-in-the-loop systems to keep the rule base and semantic graph continuously updated in response to new events, market patterns, or regulatory changes.

In conclusion, this work advances the frontier of intelligent decision support for high-tech industries by demonstrating how hybrid AI models—anchored in domain knowledge yet adaptive through learning—can bridge the gap between data abundance and strategic insight.

References

[1] Shen, Z., et al. (2024). Educational innovation in the digital age: The role and impact of NLP technology. Old and New Technologies of Learning Development in Modern Conditions, 281. International Science Group.

[2] Sutton, R. S., & Barto, A. G. (1998). Reinforcement Learning: An Introduction. MIT Press.

[3] Mikucionis, V., et al. (2025). An Intelligent Matching Approach for Upstream and Downstream Textual Information in Manufacturing Supply Chains Based on Transformer. International Journal of Management Science Research, 8(5), 9-18.

[4] Wang, Z., et al. (2025). Intelligent construction of a supply chain finance decision support system and financial benefit analysis based on deep reinforcement learning and particle swarm optimization algorithm. International Journal of Management Science Research, 8(3), 28-41.

[5] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. Nature Publishing Group.

[6] Chen, W., et al. (2024). Applying machine learning algorithm to optimize personalized education recommendation system. Journal of Theory and Practice of Engineering Science, 4(01), 101–108.

[7] Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. International Conference on Learning Representations (ICLR).

[8] Shen, Z., et al. (2025). Artificial intelligence empowering robo-advisors: A data-driven wealth management model analysis. International Journal of Management Science Research, 8(3), 1-12.

[9] Chen, H., et al. (2024). Threat detection driven by artificial intelligence: Enhancing cybersecurity with machine learning algorithms. Artificial Intelligence and Machine Learning Frontiers, 1(008).

[10] Pal, P., et al. (2025). AI-Based Credit Risk Assessment and Intelligent Matching Mechanism in Supply Chain Finance. Journal of Financial Technology and Innovation, 12(3), 145-167. Elsevier.

[11] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. ACM.

[12] Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. Neural Computation, 18(7), 1527–1554. MIT Press.

[13] Wang, Y., et al. (2025). Research on the Cross-Industry Application of Autonomous Driving Technology in the Field of FinTech. International Journal of Management Science Research, 8(3), 13-27.

[14] Drucker, H., Burges, C. J., Kaufman, L., Smola, A., & Vapnik, V. (1997). Support vector regression machines. Advances in Neural Information Processing Systems, 9, 155–161. MIT Press.

[15] Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. International Conference on Learning Representations (ICLR).

[16] Liu, Y., et al. (2023). Grasp and inspection of mechanical parts based on visual image recognition technology. Journal of Theory and Practice of Engineering Science, 3(12), 22-28.

[17] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR).

[18] Chowdhury, G. G. (2003). Natural language processing. Annual Review of Information Science and Technology, 37(1), 51–89. Wiley.

[19] Zadeh, L. A. (1996). Fuzzy logic = computing with words. IEEE Transactions on Fuzzy Systems, 4(2), 103–111. IEEE.

[20] Du, S., et al. (2024). Improving science question ranking with model and retrieval-augmented generation. The 6th International Scientific and Practical Conference "Old and New Technologies of Learning Development in Modern Conditions", 252.

[21] Tian, M., et al. (2024). The application of artificial intelligence in medical diagnostics: A new frontier. Artificial Intelligence and Machine Learning Frontiers, 1(008).

[22] Koller, D., & Friedman, N. (2009). Probabilistic Graphical Models: Principles and Techniques. MIT Press.

[23] Chen, S., et al. (2024). 3D Pop-Ups: Omnidirectional image visual saliency prediction based on crowdsourced eye-tracking data in VR. Displays, 83, 102746. Elsevier.

[24] Pearl, J. (2009). Causality: Models, Reasoning, and Inference (2nd Edition). Cambridge University Press.

[25] Shen, Z. (2023). Algorithm optimization and performance improvement of data visualization analysis platform based on artificial intelligence. Frontiers in Computing and Intelligent Systems, 5(3), 14-17.

[26] Wang, Y., et al. (2025). AI End-to-End Autonomous Driving. Journal of Autonomous Vehicle Technology, 2025.

[27] Lin, S., et al. (2024). Artificial Intelligence and Electroencephalogram Analysis Innovative Methods for Optimizing Anesthesia Depth. Journal of Theory and Practice in Engineering and Technology, 1(4), 1-10.

[28] Clark, S., et al. (2025). Construction of a Supply Chain Credit Risk Evaluation Model for Manufacturing Enterprises Using XGBoost. International Journal of Management Science Research, 8(1), 45-63.

[29] Lipton, Z. C. (2016). The mythos of model interpretability. arXiv preprint arXiv:1606.03490.

[30] Chew, J., et al. (2025). Artificial intelligence optimizes the accounting data integration and financial risk assessment model of the e-commerce platform. International Journal of Management Science Research, 8(2), 7-17.

[31] Russell, S., & Norvig, P. (2010). Artificial Intelligence: A Modern Approach (3rd Edition). Prentice Hall.

[32] Shen, Z., et al. (2024). Application of three-dimensional coding network in screening and diagnosis of cervical precancerous lesions. Frontiers in Computing and Intelligent Systems, 6(3), 61–64.

[33] Chen, S., et al. (2023). Poster graphic design with your eyes: An approach to automatic textual layout design based on visual perception. Displays, 79, 102458. Elsevier.

[34] Shen, Z., et al. (2024). Artificial intelligence empowering robo-advisors: A data-driven wealth management model analysis. International Journal of Management Science Research, 8(3), 1-12.

[35] Mehta, R., et al. (2025). Research on Enterprise Financial Text Mining and Risk Management System Based on BERT. Frontiers in Business, Economics and Management (FBEM). ISSN: 2766-824X.

[36] Pan, Y., et al. (2024). Application of three-dimensional coding network in screening and diagnosis of cervical precancerous lesions. Frontiers in Computing and Intelligent Systems, 6(3), 61–64.

[37] Shen, Z., et al. (2025). Artificial intelligence empowering robo-advisors: A data-driven wealth management model analysis. International Journal of Management Science Research, 8(3), 1-12.

[38] Pal, P., et al. (2025). AI-Based Credit Risk Assessment and Intelligent Matching Mechanism in Supply Chain Finance. Journal of Financial Technology and Innovation, 12(3), 145-167. Elsevier.

[39] Clark, S., et al., "Construction of a Supply Chain Credit Risk Evaluation Model for Manufacturing Enterprises Using XGBoost," Int. J. Manage. Sci. Res., vol. 8, no. 1, pp. 45–63, 2025.

[40] Z. Shen, et al., "Research on Application of Whale Optimization Algorithm in Financial Payment Fraud Detection," in Proc. 2025 4th Int. Conf. Artif. Intell., Internet Digit. Econ. (ICAID), Guangzhou, China, Apr. 2025, pp. 88–95.

[41] Chowdhury, G. G. (2003). Natural language processing. Annual Review of Information Science and Technology, 37(1), 51–89. Wiley.

[42] Wei, K., et al. (2023). The application of artificial intelligence to the Bayesian model algorithm for combining genome data. Academic Journal of Science and Technology, 8(3).

[43] Du, S., et al. (2024). Improving science question ranking with model and retrieval-augmented generation. The 6th International Scientific and Practical Conference "Old and New Technologies of Learning Development in Modern Conditions", 252.
[44] Koller, D., & Friedman, N. (2009). Probabilistic Graphical Models: Principles and Techniques. MIT Press.