



Indian Journal of Science, Technology & Environment (IJSTE)

An International Peer Reviewed, Refereed, Indexed & Quarterly Research Journal of Science, Technology & Environment

ISSN : 3108-2017(Online)

3108-1304(Print)

IIFS Impact Factor : 5.875

Vol.-2; Issue-3(July-Sept.) 2026

Page No.- 17-28

©2026 IJSTE

<https://ijste.gyanvidha.com>

Author's :

Hemlata Patel

(Assistant Teacher), Kashidih,
Chandrapur, Chhattisgarh. Research
Scholar, ISBM University,
Gariyaband, Chhattisgarh, India.

Artificial Intelligence and Machine Learning for Intelligent Decision Support Systems : A Comprehensive Review

Abstract : Intelligent Decision Support Systems (IDSS) have undergone a paradigm shift with the integration of Artificial Intelligence (AI) and Machine Learning (ML), evolving from static data-processing tools into dynamic, adaptive systems capable of handling complex, uncertain, and open-ended decision environments. This paper provides a comprehensive review of the state-of-the-art in AI and ML for intelligent decision support, examining foundational architectures, emerging technological innovations, and cross-sectoral applications. Key developments analyzed include the integration of knowledge graphs with retrieval-augmented generation (RAG), hybrid human-AI collective intelligence frameworks, explainable AI (XAI) for transparency and trust, probabilistic graphical models such as Bayesian networks, and the convergence of operations research with machine learning for data-driven optimization. Applications across healthcare, supply chain management, finance, climate adaptation, and defense are examined through recent case studies and research initiatives. The paper identifies persistent challenges including scalability, interoperability, personalization, and the need for robust evaluation methodologies in domains where ground truth is unavailable. Future research directions are proposed, emphasizing the importance of neuro-symbolic integration, responsible AI principles, and the development of frameworks that can dynamically orchestrate between structured reasoning and generative capabilities. This review synthesizes findings from academic literature and major research initiatives to present a holistic understanding of how AI and ML are transforming decision support and outlines pathways toward more effective, trustworthy, and adaptive systems.

Keywords : intelligent decision support systems, artificial intelligence, machine learning, knowledge graphs, retrieval-

Corresponding Author :

Hemlata Patel

(Assistant Teacher), Kashidih,
Chandrapur, Chhattisgarh. Research
Scholar, ISBM University,
Gariyaband, Chhattisgarh, India.

augmented generation, explainable AI, Bayesian networks, human-AI collaboration, data-driven optimization.

1. Introduction : Organizations across every sector are increasingly reliant on data to inform decisions, moving beyond intuition-based approaches to evidence-driven strategies. Despite the abundance of data, decision-makers continue to face significant challenges including information overload, fragmented data silos, and the difficulty of extracting actionable insights from complex, heterogeneous datasets (Ismail et al., 2025). Traditional Decision Support Systems (DSS), while valuable for structured problems, lack the reasoning capabilities and contextual awareness required to navigate the complexity of contemporary decision environments—particularly when decisions cut across domains and require nuanced understanding of interrelated variables (Wu et al., 2025).

The emergence of Artificial Intelligence (AI) and Machine Learning (ML) has fundamentally transformed the landscape of decision support. Modern Intelligent Decision Support Systems (IDSS) leverage these technologies to move beyond simple data retrieval and visualization, incorporating capabilities for prediction, optimization, reasoning, and natural language interaction. As noted by the European Commission's HACID project, “solutions to these problems are not constrained to a limited set of alternatives, and span many knowledge domains” (CORDIS, 2025), demanding new approaches that combine human expertise with machine intelligence.

In recent years, advances in deep learning, knowledge graphs, natural language processing, and large language models have further accelerated the evolution of intelligent systems. These technologies enable systems to analyze massive datasets, detect hidden patterns, generate predictive insights, and support real-time decision-making. For instance, AI-driven analytics can identify trends in healthcare data for early disease detection, optimize logistics operations in supply chains, detect fraudulent activities in financial systems, and support evidence-based policymaking in governance. Such capabilities significantly enhance the ability of organizations to respond to uncertainty and rapidly changing environments.

Another critical aspect of intelligent decision systems is the integration of human intelligence with machine learning models. Human-AI collaboration allows systems to combine computational efficiency with human judgment, domain expertise, and ethical considerations. This collaborative framework improves transparency, interpretability, and trust in AI-based decision systems, which are essential for their widespread adoption across industries.

This paper addresses three central research questions: First, what are the foundational architectures and technological innovations driving contemporary intelligent decision support systems? Second, how are these technologies being applied across critical sectors including healthcare, supply chain management, finance, and public policy? Third, what challenges remain and what future directions promise to advance the field?

The research methodology employed in this paper is a comprehensive literature review, synthesizing findings from academic journals, conference proceedings, and major research initiatives. The paper begins by establishing the conceptual evolution from traditional DSS to intelligent systems, then examines key technological innovations, analyzes applications across domains, and concludes with challenges and future directions. In doing so, the study aims to provide a comprehensive understanding of how intelligent technologies are reshaping decision support processes and enabling more informed, adaptive, and data-driven decision-making in modern organizations.

2. Evolution and Foundations of Intelligent Decision Support Systems

2.1 From Decision Support Systems to Intelligent Systems : The concept of decision support systems emerged in the 1960s and 1970s, with pioneering work by Keen, Sprague, and Bonczek establishing DSS as data-driven systems designed to support managerial decision-making (Keen et al., as cited in Kenett & Kenett, 2025). These early systems integrated data management, model management, and user interface components to provide decision-makers with organized information and analytical capabilities.

The evolution toward intelligent decision support began with the incorporation of knowledge-based systems and expert systems in the 1980s, which captured domain expertise in rule-based formats. However, these systems were limited by their reliance on explicit knowledge engineering and their inability to learn from data or adapt to changing conditions (Ismail et al., 2025).

The integration of machine learning marked a fundamental shift, enabling systems to discover patterns, learn from experience, and improve performance over time. Ismail et al. (2025) trace this evolution, noting that "combining MCDM techniques with AI improves decision-making process quality and organizational performance, particularly in complex and uncertain contexts" (p. 3). Contemporary IDSS represent a synthesis of multiple AI disciplines, including machine learning, knowledge representation, natural language processing, and optimization.

2.2 Core Architectures and Components : Modern intelligent decision support systems typically comprise several integrated components that work together to process information, generate insights, and support decision-making. Drawing from the framework proposed by Wu et al. (2025), Table 1 presents the core architectural components of contemporary IDSS.

Table 1: Core Components of Intelligent Decision Support Systems

Component	Primary Function	Key Innovations
Knowledge Representation Layer	Domain modeling and cross-domain mapping	Multi-tier architectures with hybrid embeddings; integration of ontologies and knowledge graphs
Retrieval and Reasoning Module	Information discovery and inference	Semantic search; structure-aware graph traversal; logical inference
Machine Learning Engine	Pattern discovery and prediction	Supervised, unsupervised, and reinforcement learning; deep neural networks
Optimization Module	Solution generation under constraints	Integration of operations research with ML; data-driven optimization
Human-AI Interface	Interaction and explanation	Natural language interfaces; explainable AI; visualization
Orchestration Layer	Dynamic selection of reasoning pathways	Utility-based orchestration; continuous learning from feedback

Source: Synthesized from Wu et al. (2025); Ismail et al. (2025); CORDIS (2025)

The orchestration layer represents a particularly important innovation, addressing the need to dynamically select appropriate reasoning pathways based on problem characteristics. As Wu et al. (2025) explain, effective systems require mechanisms to decide "when to deploy structured reasoning and when to deploy generative reasoning according the context and complexity of the decision" (p. 4).

2.3 The Role of Multi-Criteria Decision Making : Multi-Criteria Decision Making (MCDM) methods have become integral to intelligent decision support, particularly in contexts involving

trade-offs among competing objectives. Ismail et al. (2025) provide a comprehensive review of MCDM techniques and their hybridizations with fuzzy, grey system, and neutrosophic theories to address uncertainty and complexity in decision environments.

The integration of MCDM with AI and ML enables systems to handle both quantitative and qualitative criteria, incorporate stakeholder preferences, and generate recommendations that balance multiple objectives. Applications span material selection, medical device selection, operations management, quality evaluation, and risk management (Ismail et al., 2025).

3. Key Technological Innovations in Intelligent Decision Support


3.1 Knowledge Graphs and Semantic Representation : Knowledge graphs have emerged as powerful tools for representing and reasoning about domain knowledge in decision support contexts. Defined as "the combination of semantic technologies and graph structures to create connected representations of entities, the relation between them and their properties" (Wu et al., 2025, p. 3), knowledge graphs provide structured foundations for knowledge-intensive applications.

The HACID project exemplifies this approach, developing Domain Knowledge Graphs (DKGs) semi-automatically assembled from available data sources including scientific and gray literature. Given a specific decision problem, a pool of experts is consulted to extract supporting evidence and enrich it, generating a Case Knowledge Graph (CKG) as a subset of the DKG (CORDIS, 2025). This approach harnesses collective intelligence while maintaining traceability and explainability.

Recent developments have addressed the challenge of semantic drift in hybrid embedding techniques, particularly for rare entities. Wu et al. (2025) propose a three-pronged mitigation strategy including eigenvalue equalization for model identifiability, confidence-based weighting that dynamically adjusts embedding weights based on entity centrality and relationship frequency, and fallback strategies employing structural embeddings when semantic information is limited.

Figure 1 illustrates a knowledge graph structure for healthcare decision support, demonstrating how entities such as patients, treatments, conditions, and providers are connected through relationship edges representing clinical interactions.

Figure 1: Healthcare Knowledge Graph Structure



Source: Adapted from Wu et al. (2025)

3.2 Retrieval-Augmented Generation and Large Language Models : Retrieval-Augmented Generation (RAG) has emerged as a powerful approach to overcoming limitations of large language models (LLMs) in decision support contexts. RAG architectures combine a retrieval component that identifies relevant information from knowledge sources with a generation component that produces responses based on that information, improving factual grounding while maintaining the flexibility and natural language capabilities of generative models (Wu et al., 2025).

The integration of knowledge graphs with RAG represents a particularly promising direction. Wu et al. (2025) demonstrate significant improvements in decision accuracy, reasoning transparency, and context relevance compared to using either technology alone. Their framework, evaluated across financial services, healthcare management, and supply chain domains, performs particularly well for cross-domain reasoning and ambiguous queries.

Recent advances include Microsoft's GraphRAG framework, which uses LLM-generated knowledge graphs to improve context window population, and KG²RAG, which demonstrates how knowledge graphs can guide chunk expansion and organization to improve retrieval diversity and coherence (Wu et al., 2025). Domain-specific applications include manufacturing-focused Document GraphRAG systems and healthcare applications like MedRAG that combine knowledge graph-elicited reasoning with RAG.

Research from the Swedish Defence Research Agency has also explored applications of LLMs in decision support, including work on "jailbreaking" large language models to understand safety alignment (Rosengren et al., 2025) and strategic steering of LLMs through game-theoretic action space optimization (Lavebrink et al., 2025).

3.3 Explainable AI for Transparent Decision Support : The opacity of many machine learning models poses significant challenges for decision support, where stakeholders must understand and trust system recommendations. Explainable AI (XAI) addresses this challenge by developing techniques that make model behavior interpretable to humans.

Olan et al. (2025) examine XAI capabilities in supply chain decision support, utilizing the SHapley Additive exPlanations (SHAP) technique to analyze online data using Python ML processes. They note that "explanatory algorithms are specifically crafted to augment the lucidity of ML models by furnishing rationales for the prognostications they produce" (p. 809). Their study establishes measurable standards for identifying constituents of XAI and DSS that augment decision-making in supply chain contexts.

The HACID project emphasizes explainability through traceable processes based on supporting evidence. By gathering expert advice in collective solutions that aggregate individual opinions and expand them with machine-generated suggestions, the system maintains transparency about the origins of recommendations (CORDIS, 2025). This approach is particularly valuable in domains where ground truth is not available and decisions must be defensible to stakeholders.

The CoRDS project, funded under the Marie Skłodowska-Curie Actions programme, further emphasizes that data-driven optimization tools "must be safe, transparent, traceable and non-discriminatory, i.e. follow the principles of trustworthy AI" (CORDIS, 2026). This reflects growing recognition that explainability is not merely a technical feature but an ethical imperative for responsible AI deployment.

3.4 Bayesian Networks for Probabilistic Reasoning : Bayesian networks have proven particularly valuable for decision support applications requiring reasoning under uncertainty. Kenett and Kenett (2025) provide three detailed case studies demonstrating Bayesian network

applications: COVID-19 pandemic management integrating mobility and hospitalization data, website usability assessment, and conflict resolution politography for policy analysis.

The COVID-19 case study illustrates how Bayesian networks enable scenario analysis critical for decision-making. By integrating hospital data from ministries of health with Google mobility data, the researchers constructed networks that could assess the impact of non-pharmaceutical interventions. As Kenett and Kenett (2025) demonstrate, conditioning on variables in the network enables quantitative analysis of alternative scenarios: comparing full restrictions to no restrictions showed increases in high-level hospitalizations from 19% to 39% and deaths from 30% to 36% (p. 6).

- Bayesian networks offer several advantages for decision support:
- Integration of data from heterogeneous sources
- Visualization of probabilistic relationships
- Capability for "what-if" analysis through conditioning
- Handling of uncertainty through probability theory
- Interpretable structure that domain experts can validate

The website usability case study presented by Kenett and Kenett (2025) applies Bayesian networks to clickstream data for diagnosing design deficiencies, while the conflict resolution application demonstrates how economic, demographic, and other data can be analyzed to create data-driven narratives for policymakers.

3.5 Personalization through Machine Learning : Recognizing that different decision-makers may benefit from different forms of support, Bhatt et al. (2025) address the problem of learning personalized decision support policies. Their work posits that "personalizing access to decision support tools can be an effective mechanism for instantiating the appropriate use of AI assistance" (p. 14203).

The authors develop Modiste, an interactive tool that leverages stochastic contextual bandit techniques to learn personalized policies for decision-makers about whom initially no information is available. Through computational experiments characterizing expertise profiles, they demonstrate when personalized policies outperform population-wide baselines. Their experiments include realistic forms of support (e.g., expert consensus and predictions from large language models) on vision and language tasks, with human subject experiments confirming practical utility.

This personalization approach addresses a fundamental challenge in decision support: the heterogeneity of human decision-makers. Factors such as domain expertise, cognitive style, and familiarity with AI tools influence how individuals interact with and benefit from decision support systems. Adaptive systems that learn to match support types to individual needs promise to enhance both decision quality and user satisfaction.

3.6 Topic Modeling for Knowledge Discovery : Topic modeling techniques provide powerful tools for uncovering key themes within large, heterogeneous datasets to inform decision-making. Cuevas-Molano et al. (2025) present the Topic Analysis and Search Engine (tase), a platform developed within the European IntelComp project that addresses limitations of traditional topic modeling approaches.

Tase combines Bayesian and neural topic modeling techniques within a unified framework incorporating an expert-in-the-loop approach for training and curating models. This framework is enhanced with a Solr-based exploitation tool featuring innovative indexing methods and novel criteria for document retrieval, enabling efficient semantic similarity calculations and seamless integration into decision-support systems. The platform's scalability

and effectiveness are demonstrated through two real-world Science, Technology, and Innovation (STI) use cases (Cuevas-Molano et al., 2025).

3.7 Integration of Operations Research and Machine Learning : The convergence of Operations Research (OR) and Machine Learning represents a significant frontier for decision support. The CoRDS project (CORDIS, 2026) addresses this integration, noting that "specialized optimization methods have been developed to address complex decision problems, but these rely heavily on expert knowledge, limiting their ability to adapt to changing data. Conversely, ML excels in leveraging extensive data for predictive tasks, but struggles with combinatorial optimization."

Data-driven optimization (DDO) tools that integrate OR and ML combine OR's problem-solving capabilities with ML's data utilization strengths. These tools must not only provide high-quality decisions efficiently but also comply with government and industry standards regarding safety, transparency, and non-discrimination. The CoRDS doctoral network is developing training programs to create analytics experts who can translate research into prototype tools addressing real-world problems across logistics, healthcare, public transportation, production, finance, publishing, and machine translation (CORDIS, 2026).

3.8 Hybrid Human-AI Collective Intelligence : The HACID project (CORDIS, 2025) explores a novel approach to decision support: hybrid collective intelligence based on synergy between humans and machines. This approach addresses open-ended problems where solutions are not constrained to limited alternatives and span multiple knowledge domains.

The HACID-DSS architecture combines structured domain knowledge (semi-automatically assembled in domain knowledge graphs) with human expertise. For specific cases, experts are consulted to extract supporting evidence and enrich it, generating case knowledge graphs and providing solutions. The system then gathers expert advice into collective solutions that aggregate individual opinions and expand them with machine-generated suggestions. This approach harnesses "the wisdom of the crowd in open-ended problems, relying on a traceable process based on supporting evidence for better explainability" (CORDIS, 2025).

Two compelling case studies demonstrate the approach: crowd-sourcing medical diagnostics to address knowledge fragmentation in medicine, and climate services for urban adaptation to support policymakers in defining strategies for future climate impacts.

4. Applications across Domains

4.1 Healthcare and Medical Decision Support : Healthcare represents one of the most active domains for intelligent decision support research and deployment. Applications span clinical diagnosis, treatment planning, resource allocation, and public health management.

The HACID project's medical diagnostics case study addresses knowledge fragmentation, helping medical professionals make more accurate diagnoses by aggregating distributed expertise (CORDIS, 2025). The approach recognizes that medical knowledge is increasingly specialized and fragmented across institutions and individuals; hybrid collective intelligence can synthesize this distributed knowledge to support clinical decisions.

Wu et al. (2025) demonstrate healthcare applications of knowledge graph-enhanced RAG, including precision medicine applications where understanding relationships between genes, diseases, and treatments is critical. Their framework supports reasoning about complicated medical relationships, as illustrated in Figure 1.

Kenett and Kenett (2025) present a COVID-19 pandemic management decision support system that integrated hospital data with mobility data to assess the impact of non-pharmaceutical interventions. This system enabled policymakers to evaluate alternative

scenarios, such as the effects of closing airports or maintaining open mobility, on hospitalization and mortality outcomes.

4.2 Supply Chain Management : Supply chain decision support has been transformed by AI and ML capabilities for demand forecasting, inventory optimization, risk management, and logistics planning. Ismail et al. (2025) identify supply chain management as a key application domain for smart decision support systems, with research addressing quality evaluation, risk management, and operations optimization.

Olan et al. (2025) examine explainable AI capabilities in supply chain decision support, noting that "the emergence of various supply chain (SC) platforms in modern times has altered the nature of SC interactions, resulting in a notable degree of uncertainty" (p. 808). Their analysis uses SHAP techniques to provide rationales for ML model predictions, enhancing transparency and trust in automated supply chain decisions.

Wu et al. (2025) validate their knowledge graph-RAG framework in supply chain contexts, demonstrating improved performance for cross-domain reasoning and ambiguous queries—common challenges in global supply chains facing disruptions, regulatory changes, and shifting demand patterns.

4.3 Financial Services : Financial services have long been early adopters of decision support technologies, with applications in risk assessment, investment analysis, fraud detection, and regulatory compliance. Wu et al. (2025) include financial services among their validation domains for knowledge graph-enhanced RAG, demonstrating improvements in decision accuracy and context relevance.

The CoRDS project (CORDIS, 2026) includes finance as an application sector for data-driven optimization tools, addressing complex decision problems such as portfolio optimization, credit risk assessment, and algorithmic trading. The integration of OR and ML in this domain promises to combine rigorous optimization with adaptive learning from market data.

4.4 Climate Adaptation and Public Policy : Climate adaptation presents decision-makers with unprecedented challenges characterized by deep uncertainty, long time horizons, and complex interdependencies. The HACID project's climate services case study addresses this domain, providing support to policymakers in defining urban adaptation strategies to face future climate impacts (CORDIS, 2025).

The approach combines structured domain knowledge about climate science, urban systems, and adaptation options with expert judgment about local conditions and priorities. By aggregating collective intelligence and generating machine-suggested alternatives, the system supports evidence-based policy formulation in a domain where traditional decision support approaches have struggled.

4.5 Defense and Security : The Swedish Defence Research Agency (FOI) has produced extensive research on decision support systems for defense and security applications. Publications address diverse topics including:

- Anti-submarine warfare planning using public belief states and self-play reinforcement learning (Limér et al., 2025)
- Autonomous generation of courses of action in mechanized combat operations (Schubert et al., 2025)
- Cyber situation awareness during emerging cyberthreats (Andreasson et al., 2025)
- Detection of emerging cyberthreats through active learning (Brynielsson et al., 2025)
- Exploratory wargaming with AI "superhuman" tacticians (Bell et al., 2025)

These applications demonstrate the potential of AI and ML to support high-stakes decisions in

time-critical, adversarial environments where the costs of error are extreme.

5. Challenges and Future Directions

5.1 Scalability and Interoperability : Despite significant advances, scalability remains a challenge for intelligent decision support systems. Ismail et al. (2025) identify scalability, interoperability, and user interface challenges as current research gaps requiring attention. As systems incorporate larger knowledge graphs, more sophisticated models, and real-time data streams, computational demands increase correspondingly.

Interoperability challenges arise from the heterogeneity of data sources, knowledge representations, and decision support tools. The HACID project's emphasis on cross-border and cross-domain applications highlights the need for systems that can operate across organizational and national boundaries (CORDIS, 2025). Standards for knowledge representation, model sharing, and API design will be essential for realizing this vision.

5.2 Personalization and Adaptation : While Bhatt et al. (2025) have demonstrated the potential of personalized decision support policies, much work remains to understand how personalization interacts with task characteristics, user expertise, and organizational contexts. Future research should explore adaptive systems that not only personalize initial support but also adapt as users learn and as tasks evolve.

The dynamic knowledge orchestration proposed by Wu et al. (2025) represents one approach to adaptation, selecting reasoning pathways based on decision characteristics and context. Extending this concept to incorporate user-specific factors could enable more sophisticated personalization.

5.3 Trust and Explainability : Trust remains a critical challenge for AI-powered decision support. The CoRDS project's emphasis on "trustworthy AI" principles-safety, transparency, traceability, and non-discrimination-reflects growing recognition that technical performance alone is insufficient (CORDIS, 2026).

Future research should explore multi-level explanation frameworks that provide different types and depths of explanation for different stakeholders and contexts. Wu et al. (2025) call for synthesizing symbolic reasoning paths with natural language explanations, combining the rigor of structured representations with the accessibility of natural language.

5.4 Evaluation Methodologies : Evaluating intelligent decision support systems presents unique challenges, particularly in domains where ground truth is not available. The HACID project addresses this through "a set of evaluation methods proposed to deal with domains where ground truth is not available, demonstrating the suitability of the proposed approach in a wide range of application domains" (CORDIS, 2025).

Future research should develop robust evaluation frameworks that assess not only decision accuracy but also user trust, decision process quality, and organizational impact. Longitudinal studies tracking how systems affect decision-making over time will be particularly valuable.

5.5 Neuro-Symbolic Integration : The integration of neural and symbolic approaches represents one of the most promising directions for intelligent decision support. Wu et al. (2025) note that "recent evaluations have been proposed to augment causal reasoning by means of knowledge graphs and medical applications of neuro-symbolic systems," suggesting "a powerful integration of symbolic and neural approaches" (p. 4).

Neuro-symbolic systems combine the pattern recognition capabilities of neural networks with the structured reasoning and interpretability of symbolic AI. For decision support, this integration promises systems that can learn from data while maintaining explicit knowledge

representations that can be inspected, validated, and updated by domain experts.

5.6 Responsible AI and Governance : As intelligent decision support systems are deployed in high-stakes domains, questions of governance and accountability become paramount. The CoRDS project's training framework emphasizes "equipping researchers to turn cutting-edge theory into practical tools" that comply with government and industry standards (CORDIS, 2026).

Future research should address governance mechanisms for AI-powered decision support, including audit trails, appeals processes, and human oversight. Understanding how to design systems that appropriately allocate decision authority between humans and machines will be essential for responsible deployment.

6. Conclusion : Artificial intelligence and machine learning have fundamentally transformed the landscape of decision support, enabling systems that can handle complexity, uncertainty, and scale in ways that were unimaginable just a decade ago. From knowledge graphs that represent structured domain knowledge to large language models that generate natural language explanations, from Bayesian networks that reason probabilistically to reinforcement learning agents that optimize sequential decisions-the toolkit for intelligent decision support has expanded dramatically.

This review has examined the key technological innovations driving this transformation, including knowledge graph-RAG integration, explainable AI, Bayesian networks, personalization, topic modeling, and the convergence of operations research with machine learning. Applications across healthcare, supply chain management, finance, climate adaptation, and defense demonstrate the breadth and impact of these technologies.

Yet significant challenges remain. Scalability, interoperability, trust, and evaluation methodologies require continued research attention. The integration of neural and symbolic approaches promises to combine the best of both paradigms. And the principles of responsible AI-transparency, fairness, accountability-must guide development and deployment.

The future of intelligent decision support lies not in fully autonomous systems that replace human judgment, but in hybrid systems that augment human intelligence through effective collaboration. As the HACID project envisions, these systems will harness "the wisdom of the crowd in open-ended problems, relying on a traceable process based on supporting evidence for better explainability" (CORDIS, 2025). By combining the complementary strengths of humans and machines, we can build decision support systems that are more effective, more trustworthy, and more capable of addressing the complex challenges facing organizations and societies.

Recommendations

1. Recommendations for Researchers : Researchers should focus on advancing neuro-symbolic integration, combining deep learning with symbolic reasoning to improve interpretability and decision accuracy (Wu et al., 2025). Another key priority is developing robust evaluation methodologies for domains where ground truth is limited, such as healthcare and policy decision-making (Kenett & Kenett, 2025). Research should also explore personalized and adaptive decision support, enabling systems to tailor recommendations to users, tasks, and contexts (Bhatt et al., 2025). In addition, further investigation into human-AI collaboration is essential to understand trust, shared decision-making, and the effective balance between machine intelligence and human judgment (CORDIS, 2025). Finally, researchers should address scalability and interoperability challenges, ensuring that intelligent decision systems can operate efficiently across large datasets, distributed environments, and diverse applications (Ismail et al.,

2025).

2. Recommendations for Practitioners and System Developers : Practitioners should design decision support systems with explainability as a core feature, enabling users to understand and trust AI-driven recommendations (Olan et al., 2025). Systems should incorporate continuous learning mechanisms that adapt to new data, user feedback, and changing environments. Developers must also adopt human-centered design principles, ensuring that tools align with users' cognitive processes and decision workflows. Maintaining data quality and governance is critical for reliable machine learning outcomes, requiring strong monitoring, documentation, and ethical oversight. Finally, intelligent decision systems should be integrated with existing organizational workflows to ensure smooth adoption and effective use.

3. Recommendations for Organizations and Policymakers : Organizations should develop strategic roadmaps for AI-enabled decision-making, aligning technology investments with long-term objectives. Investment in human capital, training, and change management is necessary to build data literacy and encourage adoption of AI tools. Policymakers should establish responsible AI governance frameworks that promote transparency, accountability, and fairness in AI-supported decisions. In addition, collaboration across technical, organizational, and regulatory domains is essential to maximize the benefits of intelligent decision support systems.

Intelligent Decision Support Systems have the potential to significantly enhance decision-making across sectors such as healthcare, finance, and public policy. However, realizing this potential requires coordinated efforts from researchers, practitioners, organizations, and policymakers. By focusing on explainable AI, human-centered design, robust evaluation, and responsible governance, stakeholders can develop intelligent systems that not only improve decision quality but also maintain transparency, trust, and ethical accountability.

References :

1. Andreasson, A., Artman, H., Brynielsson, J., & Franke, U. (2025). Cyber situation awareness during an emerging cyberthreat: A case study. *International Journal of Information Security*, 24(5), 217. <https://doi.org/10.1007/s10207-025-01106-z>
2. Bell, N., Brynielsson, J., Cohen, M., Collander-Brown, S., Elg, J., Ericsson, L., Ivori, J., Limér, C., & Mannberg, N. (2025). Exploratory wargaming with a “superhuman” tactician in the team: A controlled experiment. In *Proceedings of the 30th International Command and Control Research & Technology Symposium (ICCRTS)* (Paper 051, pp. 1–13). IC2I.
3. Bhatt, U., Chen, V., Collins, K. M., Kamalaruban, P., Kallina, E., Weller, A., & Talwalkar, A. (2025). Learning personalized decision support policies. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13), 14203–14211. <https://doi.org/10.1609/aaai.v39i13.33555>
4. Brynielsson, J., Carp, A., & Tegen, A. (2025). Detection of emerging cyberthreats through active learning. In *Recent Advances in Deep Learning Applications* (pp. 123–144). Chapman and Hall/CRC.
5. CORDIS. (2025). *Hybrid Human Artificial Collective Intelligence in Open-Ended Decision Making (HACID)*. European Commission.
6. CORDIS. (2026). *Confident Data-Driven Decision Support (CoRDS)*. European Commission.
7. Cuevas-Molano, E., Gallego, M., & Losada, D. E. (2025). Topic models for decision-support systems. *Engineering Applications of Artificial Intelligence*, 159, 1–17.
8. Ismail, M. M., Soliman, M. G., & Mohamed, M. (2025). A comprehensive literature review of smart decision support systems and its applications. *International Journal of Computers and Informatics*, 8, 1–22.
9. Kenett, R. S., & Kenett, Y. (2025). Bayesian network applications in decision support systems. *Mathematics*, 13(21), 3484.

10. Lavebrink, S., Brynielsson, J., Cohen, M., Kamrani, F., Limér, C., Lindström, M., & Vangeli, M. (2025). Strategic steering of large language models via game-theoretic optimization. In *ASONAM 2025*. Springer.
11. Limér, C., Brynielsson, J., Cohen, M., & Rydell, F. (2025). Anti-submarine warfare planning using public belief states and self-play. *IEEE International Conference on Machine Learning and Applications*.
12. Olan, F., Spanaki, K., Ahmed, W., & Zhao, G. (2025). Enabling explainable artificial intelligence capabilities in supply chain decision support making. *Production Planning & Control*, 36(6), 808–819.
13. Rosengren, J., Brynielsson, J., Johansson, F., & Jonell, P. (2025). Jailbreaking large language models: Safety alignment and response quality. *IEEE ICMLA*.
14. Schubert, J., Hansen, P., Hörling, P., & Johansson, R. (2025). Autonomous generation of different courses of action in mechanized combat operations. *ICCRTS Proceedings*.
15. Schubert, J., Kamrani, F., & Gustavi, T. (2025). Active inference for an intelligent agent in autonomous reconnaissance missions. *International Workshop on Active Inference*.
16. Wu, H., Zhang, Y., Chen, L., & Wang, X. (2025). Construction of intelligent decision support systems through integration of retrieval-augmented generation and knowledge graphs. *Scientific Reports*, 15, 35462.
17. Russell, S., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
18. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
19. Mitchell, T. (2017). *Machine Learning*. McGraw-Hill.
20. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260.
21. Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
22. Amershi, S., Weld, D., Vorvoreanu, M., et al. (2019). Guidelines for human–AI interaction. *CHI Conference on Human Factors in Computing Systems*.
23. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.
24. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.
25. Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.
26. Brown, T. B., et al. (2020). Language models are few-shot learners. *NeurIPS*.
27. Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. *Stanford Center for Research on Foundation Models*.
28. Chowdhery, A., et al. (2022). PaLM: Scaling language modeling with pathways. *Journal of Machine Learning Research*.
29. OpenAI. (2023). *GPT-4 Technical Report*.
30. Li, Y., Yang, T., & Li, H. (2023). Intelligent data analytics using machine learning techniques. *Expert Systems with Applications*, 213.
31. Zhang, C., & Ma, Y. (2019). *Ensemble Machine Learning: Methods and Applications*. Springer.
32. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *ACM SIGKDD Conference*.
33. Lundberg, S., & Lee, S. (2017). A unified approach to interpreting model predictions. *NeurIPS*.
34. Ribeiro, M., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining predictions of any classifier. *ACM SIGKDD*.
35. Shrestha, Y., Krishna, V., & von Krogh, G. (2021). Augmenting decision-making with machine learning. *Academy of Management Perspectives*.