Journal of Organizational and End User Computing
Volume 38 • Issue 1 • January-December 2026

Agentic Artificial Intelligence for Information Fusion:

Architectures, Coordination, and Decision Intelligence

Dawa Chyophel Lepcha
http://orcid.org/0000-0002-1424-0449
Graduate Institute of Biomedical Informatics and International Center for Health Information Technology, College of Medical Science and Technology, Taipei Medical University, Taiwan & Biomedical Sensors and Systems Lab, University of Memphis, USA

Aaliya Ali
Centre of Research Impact and Outcome, Chitkara University, Rajpura, India & Biomedical Sensors and Systems Lab, University of Memphis, USA

Bhawna Goyal
Marwadi University Research Center, Department of Engineering, Marwadi University, India & Faculty of Engineering, Sohar University, Oman

Shabbir Syed Abdul
http://orcid.org/0000-0002-0412-767X
Graduate Institute of Biomedical Informatics and International Centre for Health Information Technology, College of Medical Science and Technology, School of Gerontology and Long-Term Care, College of Nursing, Taipei Medical University, Taiwan

Loveleen Gaur
http://orcid.org/0000-0002-0885-1550
Alliance University, India & The University of the South Pacific, Fiji

Ilan Alon
http://orcid.org/0000-0002-6927-593X
Department of Economics and Business Administration, Ariel University, Israel & Department of Strategic Management, Institute of Strategy and Management, Corvinus University of Budapest, Hungary

Justin Z. Zhang
http://orcid.org/0000-0002-4074-9505
Department of Management, Coggin College of Business, University of North Florida, USA

Received: April 8th, 2026 | **Accepted:** July 31st, 2026

ABSTRACT

Organizations across healthcare, enterprise analytics, supply chains, and cybersecurity struggle to integrate heterogeneous data sources for timely, reliable decision-making. Traditional information fusion relies on static, pipeline-oriented architectures inadequate for dynamic, distributed, and real-time environments. This study presents a PRISMA-guided systematic review integrating agentic decision theory with organizational information systems perspectives, including the Technology Acceptance Model, Task-Technology Fit, and Sociotechnical Systems Theory. A search of four databases yielded 1,205 records, from which 380 unique publications were identified. A six-dimension taxonomy covering architectures, learning paradigms, coordination mechanisms, decision coupling, trust modeling, and uncertainty representation is developed. Hybrid cloud-edge and federated architectures with multi-agent reinforcement learning offer favorable scalability and robustness trade-offs. Bias-aware design, human-in-the-loop accountability, and explainable outputs are necessary for responsible deployment.

DOI: 10.4018/JOEUC.418738

KEYWORDS

Agentic Artificial Intelligence, Agentic Decision Support Systems, Information Fusion, Multi-Agent Systems, Decision Intelligence, Reinforcement Learning (RL), Distributed Fusion, Uncertainty Modeling

1. INTRODUCTION

Modern organizations operate in environments characterized by heterogeneous, distributed, and high-volume data sources, including enterprise databases, Internet of Things (IoT) sensors, healthcare monitoring systems, and cyber-physical infrastructures. Despite the availability of such rich data, organizations consistently face significant challenges due to data silos, in which information remains fragmented across platforms, limiting the effectiveness of agentic decision support systems and system-wide intelligence. Enterprise analytics platforms, business intelligence systems, supply chain monitoring infrastructures, and cybersecurity decision systems often operate across fragmented data silos, leading to delayed insights, inconsistent interpretations, and suboptimal decision outcomes (LI et al., 2024; Tsanousa et al., 2022). Decision makers, analysts, managers, and operators across these domains require timely, reliable, and interpretable fusion of multi-source information to support strategic and operational decisions. These challenges highlight the need for adaptive, intelligent, and decision-centric fusion frameworks that can dynamically integrate multi-source information while supporting human-centric decision processes.

Despite these advances, a critical gap persists in the literature: no existing review systematically integrates the technical dimensions of agentic AI, including architectures, learning paradigms, coordination mechanisms, trust modeling, and uncertainty representation, with the organizational and end-user information systems perspectives required for responsible, human-centric deployment. This review addresses that gap.

This review is positioned within the Journal of Organizational and End User Computing for three reasons. First, agentic fusion systems are, at their point of use, organizational decision support tools whose adoption depends on end-user trust, interpretability, and workflow alignment. Second, the organizational embedding of agentic AI raises fundamental IS questions about how information systems theory must evolve to accommodate autonomous, learning-enabled decision agents. Third, this review provides IS researchers and practitioners with a structured taxonomy and design guidelines enabling theoretically grounded evaluation and deployment of agentic decision support systems across enterprise, healthcare, and cybersecurity contexts.

There are three types of information fusion levels based on the amount of semantics behind them: data-level fusion, feature-level fusion, and decision-level fusion (Senel et al., 2023). These levels can be applied to different applications within various domains, such as medical imaging, traffic monitoring, remote sensing, and manufacturing (G. Chen et al., 2021). These advances in fusion capabilities have been possible due to advancements in machine learning and deep learning, allowing for better integration of multimodal data and representation learning across multiple sources of complex and heterogeneous data (J. Li et al., 2022). As detailed in Section 3.4, current fusion methodologies exhibit fundamental limitations in adaptability, autonomy, and decision-making integration, creating significant constraints for dynamic organizational environments (LI et al., 2024). Multisensor fusion has been shown to help improve perception accuracy; however, it has also created issues such as delays in the availability of fused information, the presence of inconsistent data, and the increased computational complexity associated with fusing data from many different sensor (Wang et al., 2025). As a result, there is increasing demand for a decision-centric approach to fusing multiple sensor systems that enables intelligent, adaptive, and context-sensitive decision-making.

Figure 1. Structured Taxonomy of Agentic Artificial Intelligence for Information Fusion

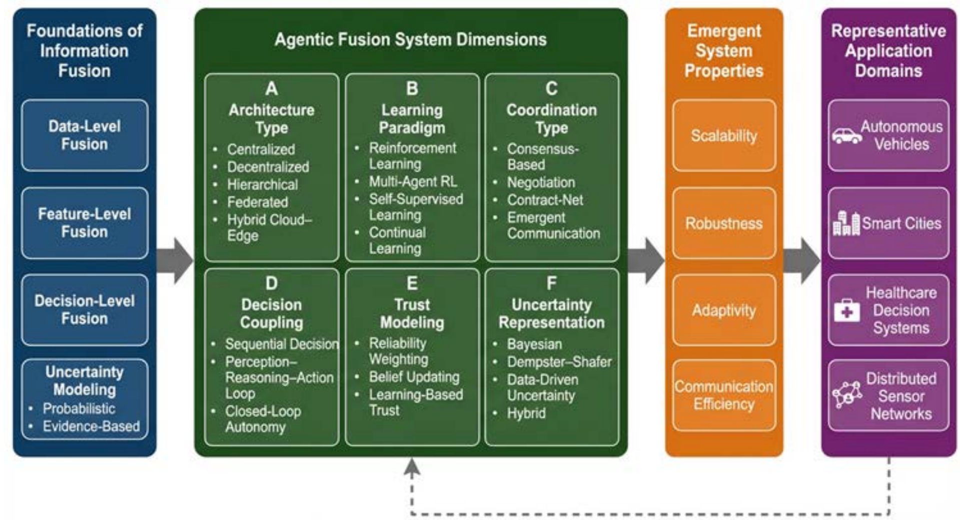

Figure created by the authors for this study.

Agentic fusion systems are recognized through existing theoretical perspectives from Information Systems. The technology acceptance model (TAM) defines the end user's acceptability of AI-based decision-making systems as a function of the end user's perceived usefulness and perceived ease of use (Qian et al., 2025; Wang et al., 2025). Task-Technology Fit (TTF) describes the degree to which a system's characteristics align with the requirements of a specific decision-making task, ensuring that both the system and organizational workflows align properly. Sociotechnical Systems Theory highlights the interdependencies among technology, intelligence, humans, and their organization, underscoring the need for collaboration between end users and intelligent agents in environments where both parties make decisions together (LI et al., 2024). Therefore, these different perspectives position Agentic Fusion as a human-centered agentic decision support system that bridges data-based intelligence and organizational action.

Agentic AI is a new way of thinking about intelligent technology based on its ability to learn and act autonomously on behalf of users while being able to perceive, process, and respond to changes within its environment in order to carry out its programmed objectives (Buşoniu et al., 2006; Sycara, 1998). In contrast to traditional fusion pipelines, agentic fusion applications utilize continuous feedback loops from their environment so they can continuously learn how best to use their knowledge to make decisions under constantly changing circumstances (C. Li et al., 2025). The authors organized the literature reviewed on this subject so that researchers could easily compare data and establish a basis for further discussion of developing a unified taxonomy, a set of guidelines for future research, and an understanding of the differences among the various types of agentic AI and information fusion research. Based on bibliometric analyses of the available literature, both agentic AI and information fusion research are rapidly gaining popularity as fields of study. The primary focus areas in these fields are multi-agent systems, RL, and decision intelligence, indicating a trend among researchers away from static data integration toward more dynamic, agent-based methods to enhance their ability to make effective decisions based on the fusion of real-time data from multiple sources. The potential types of agentic fusion systems and their relative arrangement in a conceptual framework (shown in Figures 1 and 2) illustrate how agentic fusion combines perception, reasoning, and coordination through the closed-loop framework detailed in Section 4.1 to transform data from multiple sources into usable decision-making intelligence. Key terminology used throughout this review is summarized in Appendix A (Table A1).

Figure 2. Conceptual Framework of Agentic Decision Support Systems

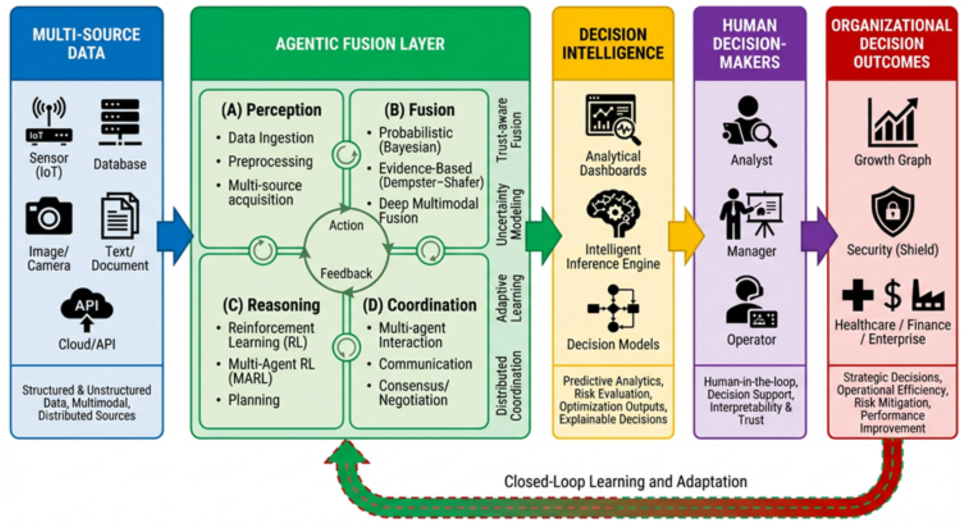

Figure created by the authors for this study.

The main contributions of this paper are summarized as follows:

- This paper presents a systematic review of recent advances in agentic AI for information fusion, highlighting the transition from static fusion pipelines to adaptive, decision-centric paradigms suitable for organizational deployment as agentic decision support systems.
- A unified taxonomy is developed that captures key dimensions including architecture, learning paradigms, coordination mechanisms, trust modeling, and uncertainty representation, providing a structured framework for analysis and design.
- The study examines centralized, distributed, hierarchical, and hybrid cloud–edge fusion architectures, along with a structured quantitative synthesis evaluating trade-offs across accuracy, latency, robustness, scalability, and communication efficiency.
- The paper investigates trust-aware fusion, uncertainty modeling, and human-in-the-loop decision intelligence, and critically evaluates applications across healthcare, autonomous systems, smart cities, and enterprise domains, while outlining ethical, managerial, and future research challenges.

The remainder of this paper is organized as follows. Section 2 presents the methodology. Section 3 presents the foundational concepts and mathematical formulations of information fusion. Section 4 introduces agentic AI and develops a taxonomy of agentic fusion systems. Section 5 discusses architectural paradigms for agentic fusion, followed by coordination and learning mechanisms in Section 6. Section 7 examines trust, uncertainty, and decision intelligence. Section 8 provides application-oriented analysis, while Section 9 presents a quantitative synthesis of agentic fusion systems. Section 10 discusses ethical and governance considerations, and Section 11 outlines challenges and future research directions. Finally, Section 12 concludes the paper.

2. METHODOLOGY

This study adopts a structured, reproducible literature search protocol guided by the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework to ensure methodological transparency, comprehensive coverage, and minimize selection bias. Four major bibliographic databases were systematically queried: Scopus, Web of Science Core Collection, IEEE Xplore Digital Library, and ScienceDirect. These repositories collectively index leading journals and conferences in information fusion, artificial intelligence, pattern recognition, robotics, and distributed intelligent systems, covering venues such as Information Fusion, IEEE Transactions series, ACM journals, and premier AI conferences including NeurIPS, ICML, ICLR, AAAI, and ICRA.

A framework for searching based on Boolean principles was used to identify common ground between agentic intelligence and the integration of different types of information, that is, information fusion. The primary search strings utilized in this process were combinations of words and phrases related to the concept of an agent (e.g., "agentic AI," "autonomous agents," "multi-agent systems," "multi-agent reinforcement learning," and "decision intelligence") and the concept of information fusion, that is, the integration of heterogeneous data from multiple sources (e.g., "information fusion," "multi-sensor fusion," "distributed fusion," and "evidence theory"). There were also some additional phrases related to the concepts being searched that were particular to certain areas, namely "trust-aware fusion," "federated fusion," "POMDP-based fusion," and "foundation models for multimodal fusion."

A total of 1,205 records were identified across four databases (Scopus: n=301; Web of Science: n=40; IEEE Xplore: n=256; ScienceDirect: n=608). After removing 258 duplicates, 947 unique publications remained for the title and abstract screening phase. Publications were excluded from this phase if they focused solely on traditional fusion without autonomous agent components, were unrelated to autonomous agent systems, or were editorial, commentaries, or published in a language other than English. This phase excluded 512 publications, leaving 435 publications for eligibility assessment. The eligibility criteria for full-text review were much stricter than those for the title and abstract screening phase, requiring that fusion, as defined in this study, be integrated with autonomous agent reasoning, adaptive learning, coordination mechanisms, and decision-making architectures. Publications that treated fusion as a separate component rather than an integrated process, did not provide sufficient methodological transparency to allow others to adequately replicate the research, or did not present empirical findings but instead offered theoretical discussions of fusion, were also excluded. For the purposes of full-text eligibility, integrated fusion was operationally defined as requiring that the fusion mechanism (a) receive inputs from the agent's perception component, (b) produce outputs that directly inform the agent's reasoning or decision component, and (c) be dynamically modifiable by the agent's learning process. Studies applying fusion solely as a preprocessing or post-processing step were classified as treating fusion as a separate component and excluded. Following cross-database deduplication of included records using DOI and title matching, 55 redundant entries were identified and removed, yielding a final synthesis corpus of 380 unique publications for the qualitative synthesis. Additionally, foundational theoretical works are retained as background references grounding the taxonomy but are not part of the 380-study synthesis corpus. As illustrated in Figure 3, the complete PRISMA flow from identification through inclusion reflects a systematic and transparent screening process across all four databases.

Each of the studies that were included has been extracted using a standardized template to capture bibliographic information, the application domain of the study, architectural configuration, the level of fusion used, the type of learning paradigm employed, the coordination mechanism, the uncertainty modeling framework, and the components of memory or knowledge representation. Using data obtained by coding each study across multiple analytical dimensions, a taxonomy was developed based on the degree of autonomy, decision tightness, scalability, and robustness mechanisms.

Inter-rater reliability was assessed across a stratified random sample of 45 publications selected to represent each application domain proportionally. Three expert raters, all holding doctoral-level expertise in AI systems or information systems research, independently coded each sampled study across the six taxonomy dimensions. Raters were blinded to one another's codings during the independent phase. Cohen's kappa was computed separately per dimension, yielding the following scores: Architecture (k=0.84), Learning Paradigm (k=0.82), Coordination (k=0.86), Trust Modeling (k=0.79), Uncertainty Representation (k=0.78), and Decision Coupling (k=0.81), with an overall composite k ranging from 0.78 to 0.86, indicating substantial to near-perfect agreement (Landis & Koch, 1977). Disagreements were resolved through structured discussion; where discussion did not yield consensus, which occurred in fewer than 5% of cases, a fourth senior reviewer provided a tiebreaking assessment. The standardized data extraction template, including all coded fields and response categories, is provided in Appendix B (Table B1) to support replication to support replication and transparency consistent with PRISMA 2020 standards. The full list of 380 included studies is publicly available at https://doi.org/10.17605/OSF.IO/SQ87G. The full coded corpus is available from the corresponding author upon reasonable request.

Figure 3. PRISMA 2020 Flow Diagram Illustrating the Systematic Identification, Screening, Eligibility Assessment, and Final Inclusion of Studies on Agentic Artificial Intelligence for Information Fusion. A total of 1,205 records were identified across four bibliographic databases (Scopus n=301; Web of Science n=40; IEEE Xplore n=256; ScienceDirect n=608). Following removal of 258 duplicates, 947 records were screened at the title and abstract stage. After applying inclusion and exclusion criteria, 435 records were retained, from which 55 cross-database duplicates were removed, yielding a final synthesis corpus of 380 unique publications.

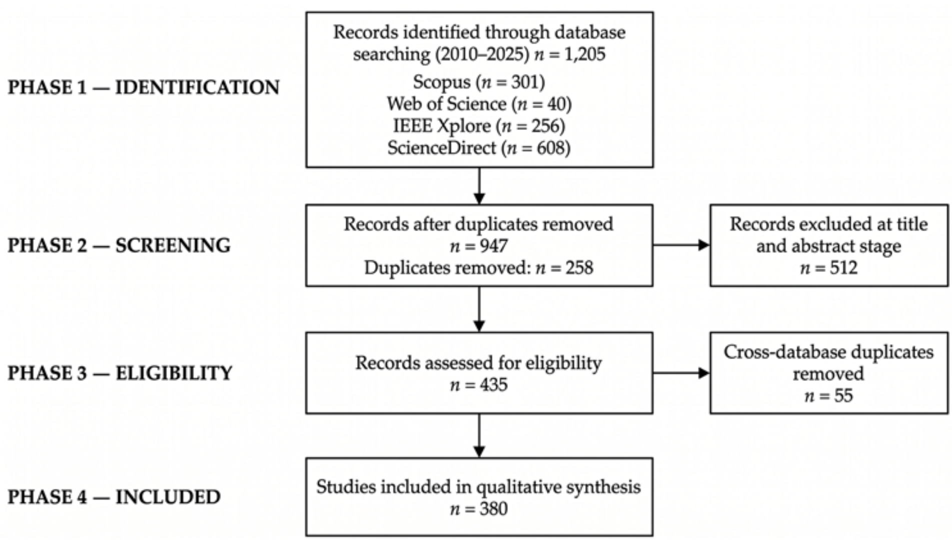

Figure created by the authors for this study.

3. FOUNDATIONS OF INFORMATION FUSION

3.1 Classical Fusion Levels

Information fusion has traditionally been categorized into hierarchical levels based on the stage at which data integration occurs, namely data-level, feature-level, and decision-level fusion. Data-level fusion combines raw signals from multiple sources to improve signal quality and reduce noise, making it suitable for low-level sensing applications such as radar, medical image fusion, and IoT-based monitoring systems (LI et al., 2024). Feature-based fusion provides a way to build richer semantics among heterogeneous modalities, including images, text, and sensor data, through the combination of extracted representations, and has proven particularly valuable in clinical decision support, where imaging and electronic health records must be jointly interpreted (Senel et al., 2023). Decision-level fusion aggregates many model or agent outputs into a final inference, leading to greater robustness in uncertain environments and forming the basis for enterprise-level agentic decision support systems where multiple analytical subsystems must converge on a common organizational recommendation (Senel et al., 2023). Although these fusion levels provide a useful conceptual framework, they tend to be implemented in static, sequential ways rather than adapting to dynamic environments. Due to the need for real-time, context-aware fusion solutions in modern applications, there is now increased interest in moving from static to more intelligent and autonomously-based fusion paradigms (Qian et al., 2025).

3.2 Mathematical Foundations

Mathematical modeling provides the theoretical basis for handling uncertainty, conflict, and decision-making in information fusion systems.

3.2.1 Bayesian Fusion

Bayesian inference is a foundational probabilistic framework for updating beliefs based on observed evidence. Given a hypothesis Hand evidence E, the posterior probability is computed as:

$$P(HE) = \frac{P(EH)P(H)}{P(E)} \tag{1}$$

For multi-source fusion with independent observations $E_1, E_2, \ldots, E_n$, the posterior becomes:

$$P\left(H \mid E_1, \ldots, E_n\right) \propto P\left(H\right)\prod_{i=1}^{n} P(E_i \mid H) \tag{2}$$

Bayesian fusion is widely used in tracking, localization, and probabilistic reasoning due to its strong theoretical foundation and interpretability (Kendall & Gal, 2017).

3.2.2 Dempster–Shafer Evidence Theory

Dempster–Shafer (DS) theory extends probabilistic reasoning by explicitly modeling uncertainty and ignorance. A basic belief assignment (BBA) $m(A)$ distributes belief mass over subsets of hypotheses. The combination of two independent evidence sources $m_1$and m_2 is defined as:

$$m_{12}(A) = \frac{1}{1 - K}\sum_{B \cap C = A} m_1(B)m_2(C) \tag{3}$$

where the conflict coefficient Kis:

$$K = \sum_{B \cap C = \varnothing} m_1(B)m_2(C) \tag{4}$$

This framework is particularly effective in high-conflict and incomplete information scenarios, where probabilistic models may struggle to represent uncertainty(Lin & Xie, 2021).

3.2.3 Markov Decision Processes and Dec-POMDPs

To incorporate sequential decision-making, fusion systems can be modeled as a Markov Decision Process (MDP):

$$M = (S, A, P, R, \gamma) \tag{5}$$

The state space is indicated by S, the action space with A, the transition function with P, the reward function with R and γ as the discount factor. Using fusion scenarios to classify belief combinations into different states and selecting a sensor from which to send information is called the selection or fusion strategy (C. Li et al., 2025). In decentralized or distributed environments, the decision process would then be represented as a decentralized partial observable Markov decision process (Dec-POMDP):

$$\mathscr{M} = \left(I, S, \{A_i\}, P, R, \{\Omega_i\}, O, \gamma\right) \tag{6}$$

where multiple agents *I* operate under partial observability. Dec-POMDPs support multisystemic agent collaboration via decentralized decision-making and fusing knowledge about elements of mutual interest (Oliehoek & Amato, 2015). The theoretical rigor is intended to provide a practical foundation for an adaptive decision-support system that could be used in organizational systems outside the laboratory or under real-life conditions. In a practical organizational deployment, the belief state b(s) is initialized from available historical data at system startup. For example, prior transaction records establish probability distributions over fraud states in an enterprise monitoring system. During operation, belief states are updated at each time step using Bayes rule applied to new observations arriving from sensors or data streams. Agents exchange belief updates through the coordination layer, enabling coherent global inference without requiring centralized state aggregation. Li et al. (2025) demonstrate a concrete operationalization in which the belief state is maintained as a learned latent representation updated continuously through multimodal sensor observations in a reinforcement learning-based fusion system.

3.3 Deep Learning-Based Fusion

Recent advances in deep learning have significantly enhanced information fusion by enabling end-to-end multimodal representation learning. Deep neural networks, including convolutional neural networks (CNNs), transformers, and attention-based models, allow the integration of heterogeneous data sources such as medical images, clinical text, and sensor signals within a unified framework (Gao et al., 2020). The ability of multimodal transformers to utilize imaging, laboratory, and patient record data has proven highly effective in providing healthcare decision support. Additionally, multimodal conditional cross-modal fusions of transactional, behavioural and external data sources have enabled real-time business intelligence using enterprise analytics. Transformer-based architectures have demonstrated significant capabilities by enabling enhanced perceptual understanding and prediction across multiple modalities in autonomous vehicle operations, medical image fusion, and other organizational analytics applications.

However, deep learning-based multimodal fusion methods also face several challenges, including interpretability, which is highly dependent on how they produce data and generalize across domains. Most deep fusion systems are static post-training and, as a result, maladaptive to the continuously-changing and dynamic nature of organizations (LI et al., 2024).

Table 1 depicts the progression of the various fusion paradigms (i.e., Bayesian, Dempster-Shafer, deep multimodal, and agentic) in their evolving capabilities across uncertainty, adaptability, scalability, and decision integration, transitioning from static, data-driven fusion methods to adaptive, multi-agent-based decision-support systems.

Table 1. Comparative Analysis of Probabilistic, Evidence-Based, Deep Multimodal, and Agentic Fusion Paradigms

Criterion	Bayesian Fusion	Dempster–Shafer Fusion	Deep Multimodal Fusion	Agentic Fusion
Theoretical Foundation	Bayesian probability theory; prior–posterior inference (Kendall & Gal, 2017)	Evidence theory; belief mass functions (Lin & Xie, 2021)	Neural representation learning (CNN, Transformer) (Gao et al., 2020)	Agent-based decision theory + RL + probabilistic reasoning (Buşoniu et al., 2006)(C. Li et al., 2025)
Uncertainty Representation	Precise probabilities	Belief intervals; explicit ignorance modeling	Implicit (attention weights, softmax)	Explicit + learned (probabilistic + trust-aware + adaptive belief update)
Conflict Handling	Likelihood-based; sensitive to prior mismatch	Dempster's rule; unstable under high conflict	Learned weighting; no explicit conflict modeling	Negotiation, trust-weighting, consensus, adaptive source reweighting
Adaptivity	Static unless re-estimated	Static rule-based	Moderate (retraining required)	High (online RL, continual learning, dynamic policy adaptation)
Scalability	Moderate; computational cost grows exponentially	Moderate; combinatorial explosion with hypotheses	High (GPU scaling, distributed training)	Very high (distributed multi-agent, federated, hierarchical scaling)
Multi-Agent Coordination	Not inherent	Not inherent	Limited (attention-based weighting)	Native support (MARL, consensus, task allocation, negotiation)
Decision-Centric Integration	Post-fusion decision layer	Post-fusion belief decision	Classifier-driven decision	Integrated perception–reasoning–action loop
Interpretability	High (mathematical transparency)	Moderate–High (belief structure)	Low (black-box deep networks)	Moderate–High (reasoning + trust modeling + explainable decision path)
Trust Modeling	Implicit via priors	Partial via belief mass	Not explicit	Explicit trust-aware weighting and source reliability modeling
Architecture Dependency	Centralized/hierarchical	Centralized / distributed	Centralized / distributed	Centralized, decentralized, hierarchical, hybrid cloud–edge

Note. Ratings reflect consensus characterizations across reviewed studies, relative to the other paradigms in this comparative framework. Contested entries include inline citations. The Moderate-High interpretability rating for Agentic Fusion reflects the availability of reasoning traces and explainable decision paths (Schick et al., 2023; Park et al., 2023). The Moderate-High rating for Dempster-Shafer Theory reflects the partial transparency of belief mass structures (Lin and Xie, 2021).

3.4 Limitations and Transition Toward Agentic Fusion

Although classical and modern fusion techniques provide powerful tools for integrating multi-source data, they exhibit several limitations that restrict their applicability in complex real-world organizational systems. Traditional architectures are largely static and pipeline-based, lacking the ability to adapt to changing environments or evolving data distributions (Ding et al., 2019). Data fusion is mainly about aggregating data, not about making decisions; therefore, there is very little integration of decision-making and action (Ding et al., 2019). Coordination, trust, and communication among the multiple organizations that comprise a system development organization present a unique challenge that current models of coordination do not support (Ding et al., 2019). In addition, while both deep learning and probabilistic methods (like Bayesian networks) are excellent at modeling uncertainty and learners' representations of data, they are generally not integrated into adaptive decision-making or multi-agent systems for coordination; therefore, there is a lack of integration of real-time, collaborative, and context-sensitive decision-making within organizations. This has led to a new area of research related to agentic AI for information fusion, where autonomous agents operate as an integrated system that incorporates perception, reasoning, and action and provides a framework for creating next-generation agentic decision support systems.

4. AGENTIC ARTIFICIAL INTELLIGENCE FOR INFORMATION FUSION

4.1 Definition and Core Components

Agentic AI represents a paradigm shift from static computational pipelines to autonomous, goal-driven systems capable of perceiving, reasoning, and acting in dynamic environments. In the context of information fusion, agentic AI integrates data processing with adaptive decision-making, enabling systems to continuously refine their behavior based on environmental feedback and evolving organizational objectives (LI et al., 2024; Park et al., 2023). It is important to distinguish agentic AI from classical multi-agent systems. While multi-agent systems research focuses on interaction protocols and coordination structures among multiple autonomous entities, agentic AI emphasizes the individual agent's capacity for goal-directed autonomy, continuous environmental perception, closed-loop learning, and adaptive decision-making. Agentic AI agents maintain persistent objectives, update their behavior through experience, and integrate reasoning and action within a unified framework (Park et al., 2023). Classical multi-agent systems research provides a necessary foundation, but agentic AI extends it toward self-directed, organization-embedded intelligence. Given the static pipeline limitations outlined in Section 3.4, existing fusion approaches cannot adequately support enterprise analytics, fraud detection, healthcare decision support, and supply chain management environments where data distributions and operational constraints evolve continuously.

It is equally important to distinguish the agentic decision support system construct used in this paper from classical DSS frameworks. Traditional decision support systems(Power, 2002; Turban & Aronson, 1998), were conceived as model-driven or data-driven tools providing decision-relevant analysis to human decision-makers while leaving the decision itself to the human. Agentic DSS extends this paradigm along three dimensions: first, autonomy, meaning the system can formulate and execute decision policies without continuous human direction; second, closed-loop adaptation, meaning the system continuously learns from decision outcomes to refine its fusion and reasoning strategies; and third, organizational integration, meaning that perception, reasoning, and action are embedded within organizational workflows rather than operating as standalone analytical tools. Human oversight remains a design requirement in agentic DSS, particularly in high-stakes domains, but the nature of that oversight shifts from operational control to supervisory validation.

The interconnected perception-reasoning-action (PRA) loop is at the core of agentic fusion behavior. During Perception, the agent takes heterogeneous data from sensors, databases, electronic health records, transactional systems, and external APIs, combining heterogeneous data from sensors, databases, electronic health records, transactional systems, and external APIs through classical fusion processes using a probabilistic or learning-based approach to deal with uncertainty and noise (Qian et al., 2025). During the Reasoning phase, the fused data is converted into a structured knowledge base and into actionable insights through probabilistic reasoning, RL, and neural networks. This enables the agent to produce decision policies that support organizational objectives (Bengio et al., 2003). After a set of decisions are made, the Action phase executes the decisions based on the data collected during both the Perception and the Reasoning phases, closing the loop by obtaining continual feedback from the environment. By being closed-loop in its design, agentic fusion differs from traditional systems by enabling continuous improvement and contextually aware decision-making embedded directly into an organization's workflow (Park et al., 2023).

From a Technology Acceptance Model perspective, the perceived usefulness of agentic fusion outputs depends critically on interpretability and trust transparency, while perceived ease of use is influenced by the degree to which the system's decision rationale is surfaced in a form accessible to analysts and managers. The Architecture and Trust Modeling dimensions of the proposed taxonomy have the strongest implications for end-user adoption: centralized architectures with transparent trust outputs reduce perceived complexity, while interpretable uncertainty representations enhance perceived usefulness by giving decision-makers confidence in system outputs.

4.2 Agent Types

Agentic fusion systems encompass diverse agent types suited to different organizational roles. Reactive agents respond directly to environmental stimuli without maintaining an explicit internal model, making them suitable for real-time fusion tasks with low latency requirements, such as cybersecurity threat alerting and IoT anomaly detection (Wooldridge, 2009).

Deliberative agents use internal representations to think through various situations; this allows them to make decisions that are neither more nor less sophisticated based on what is known about the world around them. They use this ability to work towards healthcare diagnostics and to mitigate enterprise risk through methods such as using an expert system (Jennings, 2002). On the other hand, hybrid agents combine reactive and deliberative capabilities, providing an advantageous trade-off between reactivity and decisiveness for autonomous systems or intelligent business dashboards. Recently, learning agents utilizing RL or self-supervised learning methods, along with collaborative agents that coordinate their actions with other agents towards achieving an organizational objective, have allowed agentic systems to be more versatile than ever before across distributed enterprise environments (Dautov & Distefano, 2017).

4.3 Relationship with Reinforcement Learning, Multi-Agent Systems, and Foundation Models

Agentic AI for information fusion is closely related to RL, multi-agent systems (MAS), and foundation models. RL provides a principled framework for sequential decision-making under uncertainty, enabling fusion systems to adaptively select sensor combinations, fusion strategies, and decision policies that optimize organizational performance over time (C. Li et al., 2025).

Multi-agent reinforcement learning (MARL) can extend the capabilities of distributed environments and improve coordination among agents in the enterprise ecosystem. Multi-Agent systems constitute the framework for agentic fusion, allowing agents to communicate with each other, negotiate with each other, and cooperate with each other across distributed data sources, while at the same time, enabling decentralized processing and horizontal scalability for the processing of large volumes of heterogeneous organizational data (Buşoniu et al., 2006) (Schick et al., 2023).

4.4 Taxonomy of Agentic Fusion Systems

The complexity and diversity of agentic fusion systems necessitate a structured taxonomy that captures their key design dimensions, providing a systematic framework for analyzing, comparing, and designing agentic decision support systems across organizational contexts. It is important to distinguish between the conceptual framework presented in Figure 2 and the taxonomic classification system in Figure 1 and Table 2. The conceptual framework describes the functional data flow through agentic fusion systems, from multi-source inputs through perception, fusion, reasoning, and coordination to decision intelligence outputs. The taxonomy classifies the design space of such systems along six dimensions, enabling comparative analysis and design specification. Taxonomy dimensions were derived iteratively: initial dimensions emerged from the first 50 coded studies; additional dimensions were added when consistent patterns appeared in more than 10% of subsequently coded studies; and all six dimensions stabilized at the 150-study mark, after which no new dimensions emerged.

4.4.1 Dimension Design

The proposed taxonomy is organized along six primary dimensions. The Architecture dimension categorizes systems based on structural organization, including centralized, distributed, hierarchical, federated, and hybrid cloud–edge configurations, each with distinct implications for scalability, communication patterns, and computational efficiency (Noh, 2020).

The learning paradigm dimension models how agents learn and adapt through a combination of RL, self-supervised learning, or continual learning, which allows for multiple fusion strategies and adaptive decision policies that are applicable in organizations that are evolving and have changing environments (Buşoniu et al., 2006). The coordination dimension describes how agents coordinate with one another through a variety of channels, such as consensus, negotiation, or communication protocols, which are key to creating efficient distributed decision support systems for agents (Dorri et al., 2018). The Trust Modeling dimension focuses on the credibility or reliability of sources of data and agents as well as the ability of the system to have trust-aware fusion methods that can automatically update their weights based on the source of trust for each piece of evidence, thus providing robustness to both uncertainty and adversarial conditions. The full theoretical treatment of trust calibration mechanisms, including Bayesian updating and reputation-based models, is provided in Section 7.1; the Uncertainty Representation dimension is fully operationalized in Section 7.2.

The Decision Coupling dimension characterizes the degree to which fusion and decision-making are integrated within a single closed-loop system, ranging from sequential post-fusion decision layers to tightly coupled perception-reasoning-action loops and fully closed-loop autonomous systems (Oliehoek and Amato, 2015). Full operationalization of this dimension is provided in Section 7.3. The Uncertainty Representation dimension classifies how systems model and communicates epistemic and aleatoric uncertainty, including explicit probabilistic methods such as Bayesian and Dempster-Shafer approaches, implicit learned representations such as attention weights and softmax outputs, and hybrid combinations. Full operationalization is provided in Section 7.2.

4.4.2 Validation and Consistency

To ensure the robustness and reliability of the taxonomy, a structured validation process was employed. A representative set of reviewed studies was mapped to the defined dimensions to verify completeness and consistency. Inter-rater reliability details, including per-dimension kappa scores, rater qualifications, sample size, and disagreement resolution protocol, are reported in full in Section 2. The taxonomy reflects validated theories from information fusion and MAS architecture and has been extended to include agent-centered and decision-centered views, providing a solid foundation for analyzing and developing future generations of agentic decision support systems. Table 2 illustrates how key taxonomy dimensions translate into representative organizational applications, bridging theoretical framework development and practical deployment across enterprise, healthcare, and cybersecurity domains.

Table 2. Mapping of Individual Agentic Fusion Taxonomy Dimension Values to Representative Organizational Applications. Note: Each row illustrates one configuration option within a single dimension. Real organizational deployments combine specific values across all six dimensions simultaneously. For illustrative purposes, a federated healthcare decision support system would combine: Federated Architecture, Reinforcement Learning paradigm, Consensus Coordination, Dynamic Trust Modeling, Bayesian Uncertainty Representation, and Closed-Loop Decision Coupling.

Taxonomy Dimension Value	Description in Agentic Fusion Context	Representative Organizational Application Example
Centralized Architecture	Fusion and decision-making coordinated by a central agent or system (G. Chen et al., 2021) (Gao et al., 2020)	Enterprise analytics and corporate decision intelligence platforms (Tao & Chao, 2023)
Distributed Architecture	Decentralized agents perform local fusion with peer-to-peer coordination (Jennings, 2002) (Dautov & Distefano, 2017)	IoT-enabled supply chain monitoring and logistics optimization (P. Sun et al., 2025) (Sharma et al., 2022)
Reinforcement Learning	Adaptive policy learning for dynamic fusion and decision optimization (Buşoniu et al., 2006) (C. Li et al., 2025)	Dynamic pricing and demand forecasting systems (Tao & Chao, 2023)
Multi-Agent Coordination	Cooperative or competitive interaction among multiple agents (Dautov & Distefano, 2017) (Dorri et al., 2018)	Cybersecurity threat detection and response systems (H. Chen & Wu, 2025) (Möller, 2023)
Trust Modeling	Reliability-aware weighting and credibility assessment of information sources (Mendelson & Zwillinger, 2024) (X. Sun & Zeng, 2025)	Fraud detection and risk assessment systems with confidence scoring (Halbouni et al., 2022)
Decision Coupling	Integration of fusion with real-time decision-making loops (Oliehoek & Amato, 2015) (Kendall & Gal, 2017)	Real-time business intelligence dashboards and operational analytics (Tao & Chao, 2023)

5. ARCHITECTURES OF AGENTIC FUSION SYSTEMS

Agentic information fusion systems rely on architectural design to determine how perception, reasoning, coordination, and decision-making processes are distributed across computational entities. Compared to traditional fusion architectures, agentic systems offer advantages such as decentralization, fluidity, and distributed operation, and can function effectively in large, rapidly changing, organized environments by leveraging distributed intelligence and scalable coordination. As depicted in Figure 4, agentic fusion architectures provide an evolution from static types of fusion (at the data, feature, and decision levels) through to deep learning-based types of fusion, to then agentic fusion architectures that unify perception, reasoning, action, and learning into a cyclical, closed-loop, autonomous decision-making framework. As illustrated in Figure 4, this evolution progresses from data-level, feature-level, and decision-level static fusion approaches through deep learning-based representation learning using CNN and transformer architectures, ultimately arriving at agentic fusion where increasing autonomy and contextual reasoning enable closed-loop decision intelligence across organizational environments. As such, agentic fusion architectures provide higher levels of autonomy, distributed coordination, and decision-centric integration. In this document, we will explore four of the primary architectural paradigms, such as centralized, distributed, hierarchical, and hybrid, as they relate to their properties, trade-offs, and support for various organizations.

5.1 Centralized Architectures

Centralized agentic fusion architectures are characterized by a single coordinating entity that aggregates data, performs fusion, and generates decisions. All relevant data from sensors, databases, or distributed organizational sources is transmitted to a central node, where perception, reasoning, and action components are integrated within a unified framework, enabling global optimization and consistent decision-making as the central agent has access to complete system information (LI et al., 2024). Centralized architectures that are best suited to be used in enterprise analytics platforms, centralized decision support systems for the healthcare industry, or business intelligence dashboards for organizations where there are reliable communications infrastructure, provide global context and utilize advanced fusion techniques (e.g., probabilistic inference, deep learning, optimizing based on decisions) through a combination of technologies (Gao et al., 2020). A representative example is enterprise clinical decision support platforms that aggregate electronic health records, medical imaging, and laboratory data at a central inference node, enabling global diagnostic optimization across a hospital system while providing analysts and clinicians with consistent, centrally coordinated outputs.The Technology Acceptance Model view offers insight into why centralized architectures have a positive impact on decision-makers (i.e., individuals who make decisions) and analysts and managers by allowing them access to a simplified (i.e., presented to them) interface that provides them with the ability to gain consistent outputs for their decision-making process from a central coordinating intelligence.

Figure 4. Evolution of information fusion paradigms from data-centric static approaches toward agentic fusion, where perception, reasoning, action, and learning are integrated into a closed-loop autonomous decision-making framework

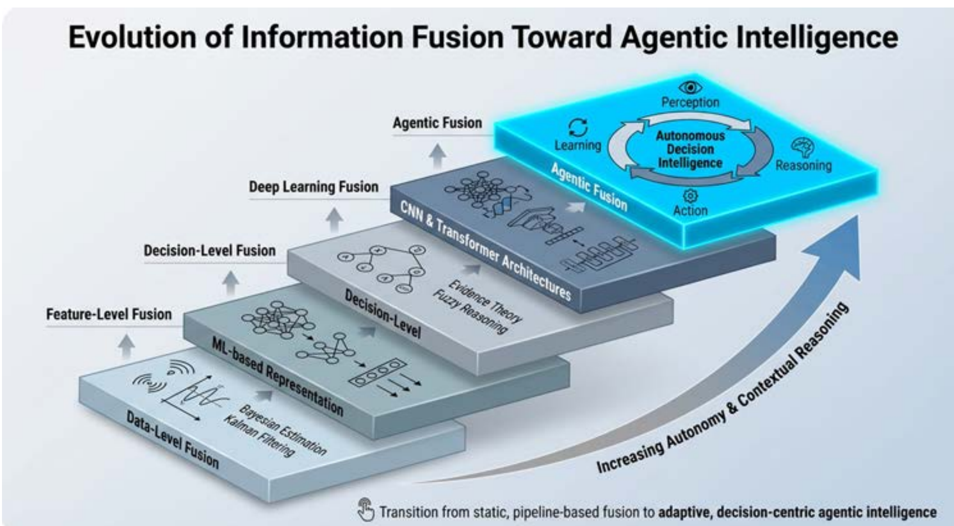

Figure created by the authors for this study.

5.2 Distributed Architectures

Distributed architectures represent a fundamental shift toward decentralized information processing, where multiple autonomous agents perform local perception, fusion, and decision-making. Agents act on locally available data and collaborate with other agents to create a coherent organizational understanding, thereby providing a scalable, resilient, and adaptive fusion system consistent with agentic AI principles (Dautov & Distefano, 2017). Consensus algorithms, belief propagation, and cooperative RL are used to share information, make decisions based on collective input, and aggregate information from a pool of distributed data sources. The advantages of these solutions have been realized in numerous applications involving geographically distributed data sources, such as IoT-enabled supply chain monitoring, autonomous vehicle networks, and smart grid management (Buşoniu et al., 2006). Warehouse robotics networks exemplify this architecture: each robotic agent performs local environmental perception and collaborates via consensus protocols for coordinated path planning without requiring a central coordinator, demonstrating the fault tolerance and scalability advantages of distributed design. These distributed computing architectures distribute the computational load across multiple agents in a system, reducing system latency and improving robustness to failure, making them suitable for enterprise environments requiring a continuous stream of real-time operational intelligence.

5.3 Hierarchical Architectures

Hierarchical architectures combine elements of centralized and distributed paradigms by organizing agents into multi-level structures, where local agents perform initial fusion and higher-level agents aggregate and refine results. The three-tier architecture is made up of three separate tiers of agents: lower-tier (data acquisition) agents perform collection, middle-tier (aggregating/abstracting) analysts aggregate and synthesize results, and upper-tier (supervisory) agents integrate lower-tier results into a single coherent global output for decision-making purposes, thereby minimizing the amount of communication overhead by transmitting only aggregate information up to upper-tier agents (Senel et al., 2023). The three-tier architecture facilitates efficient processing of large-scale organizational data while also providing whole-of-organization coordinated operation, which is typical in large-scale smart city management, industrial monitoring operations, surveillance systems, and healthcare networks since the same data often exists at multiple levels of detail (Ignatious et al., 2023). Urban traffic management systems illustrate this structure: street-level edge sensors perform raw data collection, intersection-level aggregators compute traffic flow estimates, and a city-level cloud agent optimizes arterial signal timing across the entire network (H. Chen and Wu, 2025). Hierarchical architectural designs facilitate modularity, allowing each organizational level to be optimized independently of the others with respect to its particular performance objective.

5.4 Hybrid Cloud–Edge Architectures

Hybrid cloud–edge architectures represent an emerging paradigm that integrates the strengths of centralized and distributed approaches by combining edge computing with cloud-based intelligence. Edge agents perform real-time perception and local fusion close to organizational data sources, reducing latency and bandwidth requirements, while cloud agents provide global reasoning, model training, and long-term optimization using large-scale computing resources (P. Sun et al., 2025). The separation of roles in these hybrid architectures is valuable for time-sensitive applications within an organization, such as autonomous vehicles, real-time patient monitoring, automated industrial systems, and the prevention of financial fraud, because they require both immediate local response and strategic information on a larger scale. Because the models built in the cloud can be sent to edge devices for use and constantly updated with new data as it becomes available to an organization, these hybrid architectures support adaptive learning and evolution of the system, while also providing integration capabilities with other federated learning models to support privacy and securement of the data (Pinto Neto et al., 2023). Autonomous vehicle systems exemplify hybrid architecture: onboard edge processors perform real-time object detection at low latency while cloud systems aggregate fleet-wide observations to continuously update the global perception model (Wang et al., 2025).

5.5 Task-Technology Fit Analysis of Architectural Paradigms

Synthesizing across architectural paradigms through a Task-Technology Fit lens, Table 3 can be interpreted normatively. Task-Technology Fit theory predicts that adoption and decision performance will be highest when architectural properties match the task's complexity, geographic distribution, and latency requirements. Distributed architectures best fit geographically distributed, latency-sensitive tasks such as IoT monitoring and autonomous navigation. Centralized architectures best fit analytically complex, globally coordinated tasks requiring complete situational awareness, such as enterprise command systems and centralized healthcare diagnostics. Hierarchical architectures fit large-scale tasks with both local and global coordination requirements, including smart city management. Hybrid cloud-edge architectures optimize fit for tasks requiring both immediate local response and strategic global reasoning, such as autonomous vehicle navigation and real-time financial fraud detection.

Table 3. Architectural Trade-Off Analysis in Agentic Fusion Systems: Latency–Scalability–Robustness

Architecture Type	Latency	Scalability	Robustness	Communication Overhead	Suitable Applications
Centralized Agentic Fusion	Moderate–High	Limited	Moderate	High (data aggregation)	Command/control systems, healthcare DSS (Kumar, 2025)
Decentralized Agentic Fusion	Low	High	High (fault tolerance)	Moderate	Sensor networks, IoT systems (Ding et al., 2019)
Hierarchical multi-Agent Fusion	Moderate	High	High	Structured (layered)	Smart cities, surveillance (Sharma et al., 2022)
Hybrid Agent–Cloud–Edge Fusion	Adaptive (edge low, cloud higher)	Very High	Very High	Optimized via tiering	Cyber-physical systems, autonomous vehicles (Qian et al., 2025; Wang et al., 2025)
Federated Agentic Fusion	Moderate	Very High	High (privacy-aware)	Model-update only	Healthcare, distributed sensing (Hussain et al., 2024)

Note. Latency, Scalability, and Robustness ratings are consensus characterizations from the reviewed literature. Communication overhead ratings are relative assessments. See Section 9 for supporting citations per architectural category.

6. COORDINATION AND LEARNING IN AGENTIC FUSION

Agentic information fusion systems fundamentally rely on coordination and learning mechanisms to enable distributed agents to interact, share information, and collectively achieve decision objectives under uncertainty. Unlike traditional fusion systems with defined and static forms of coordination, agentic fusion enables adaptive communication, dynamic coordination, and experience-driven decision-making that supports effective operation within complex and changing organizational environments (Dautov & Distefano, 2017; Noh, 2020).

6.1 Communication Models

Communication forms the backbone of coordination in agentic fusion systems, determining how agents exchange information and maintain a consistent understanding of the organizational environment. Agents in existing agent-based systems use event-driven or probabilistic methods (like publish-subscribe or gossip protocols), which provide agents with the ability to communicate with one another asynchronously and selectively, thereby reducing bandwidth overhead and ensuring that the overall behavior of the system remains consistent (Amin & Hossain, 2021). In multi-tiered IoT networks and enterprise monitoring environments, agents typically send updates only when there is a significant change in the environment they are monitoring, thereby enabling more efficient communication within distributed enterprise data infrastructures.

Innovations in machine learning have led to the emergence of learning-based communication methods, in which agents adaptively decide what information to share, how to share it, and with whom. This has been demonstrated in differentiable communication architectures such as CommNet (Sukhbaatar et al., 2016), DIAL (Foerster et al., 2016), and TarMAC (Das et al., 2019), in which communication channels are included within the end-to-end differentiable MARL framework, enabling agents to jointly optimize their communication policy alongside their action policy with respect to shared organizational task objectives. The result is a framework that facilitates agents to optimize their communication policy with the same objectives as their decision-making (Schick et al., 2023). Although these techniques will improve performance and scalability for agentic decision support systems in organizations, care must be taken to ensure they are designed to maintain the ongoing stability, robustness, and interpretability of those systems during use by the intended end-users, including analysts and managers.

6.2 Coordination and Negotiation

Coordination in agentic fusion systems involves managing dependencies among agents' actions and information states to achieve coherent organizational behavior. In distributed environments, agents often operate with partial knowledge, requiring mechanisms to align local decisions with global organizational objectives (Wooldridge, 2009). Through the use of market-based and auction-based coordination methods, agents can bid on work that needs to be done or resources they want to use to complete the work, and thus enhance the efficiency of the assignment of resources in situations where there are limited resources available to the organizations using them (Buşoniu et al., 2006). These types of market- and auction-based systems are also modeled using game theory, which provides a rigorous method for analyzing the interactions of agents and establishing the optimal strategies for accomplishing individual and collective goals associated with the organization (Mendelson & Zwillinger, 2024).

6.3 Consensus and Conflict Resolution

In multi-agent fusion systems, conflicting information is inevitable due to sensor noise, incomplete observations, or heterogeneous organizational data sources. Consensus-based approaches enable agents to iteratively update their beliefs through local interactions until a stable agreement is reached, using distributed averaging, consensus filtering, and belief propagation techniques well-suited for environments with reliable communication (Amin & Hossain, 2021)(Noh, 2020). Conflict resolution extends consensus processes to resolve discrepancies between information sources.

The Dempster-Shafer method for analyzing evidence helps combine conflicting information while accounting for uncertainty and missing information. This method is especially important when detecting fraud or identifying cybersecurity attacks within organizations with many different data sources that may not provide reliable information or may be aimed at harming your organization (Lin & Xie, 2021). In addition to the Dempster-Shafer method, trust-aware fusion methods assign "trust" values to both the agent (the person or thing performing the data fusion) and the data itself (the source of the information), which allows the final product to have some level of credibility based on information from unreliable sources.

6.4 Learning Paradigms: Reinforcement Learning and Multi-Agent Reinforcement Learning

Learning is a fundamental capability of agentic fusion systems, enabling adaptation to changing organizational environments and continuous improvement in decision performance. Among various learning paradigms, RL has emerged as a key approach for decision-centric fusion, where agents learn optimal policies by interacting with the environment and maximizing cumulative rewards, making it well-suited for sequential organizational decision-making problems (Buşoniu et al., 2006). In the context of information fusion, systems can use RL to adopt and change their sensor fusion combinations, strategies, and decision-making processes based on changes in the organization, which helps reduce computing and communication costs while at the same time increasing the quality of the decisions that are made (C. Li et al., 2025). One way RL can be used in enterprise agentic decision support systems is to enable these systems to dynamically prioritize data sources with the greatest potential to inform them, so they can accomplish their purpose (e.g., monitoring financial risk or detecting fraud).

RL-based frameworks learn from historical data within the organization to improve the effectiveness of data-driven decisions.

MARL is an extension of RL that allows multiple agents to learn individually while coordinating their behavior. Agents can use MARL to create joint policies to collectively obtain the best solution to a problem that they must solve together, despite their partial observability and lack of control (Buşoniu et al., 2006; Dorri et al., 2018). MARL's relevance to agent-based fusion systems stems from the need to coordinate agents to achieve the overall organizational goal. This includes coordinating situational awareness coherently across distributed nodes of the enterprise. It can include coordinating in order to measure anomalies in a synchronized manner across cybersecurity monitoring agents (Buşoniu et al., 2006).

7. TRUST, UNCERTAINTY, AND DECISION INTELLIGENCE

7.1 Trust Modeling

Trust modeling is essential in agentic fusion systems to evaluate the reliability and credibility of data sources, sensors, and agents operating within distributed organizational environments. In enterprise analytics, fraud detection, and healthcare decision support, information originates from heterogeneous sources with varying levels of accuracy, necessitating dynamic weights based on trustworthiness rather than relying on static weighting schemes that cannot reflect the changing reliability of organizational data streams. Dynamic trust estimation algorithms are incorporated into newly developed trust-based fusion frameworks to update trust levels based on historical data and the context of relationships between the source(s) providing the data and the fusion output(Granatyr et al., 2015; Jøsang et al., 2007). Probabilistic models such as Bayesian inference provide a principled method for updating beliefs about the reliability of the source(s) as new information is generated, whereas evidence-based models improve upon this by explicitly modeling the uncertainty and conflict between sources in order to support a better assessment of the source(s)' trustworthiness where there is uncertainty (X. Sun & Zeng, 2025). When viewed through the Technology Acceptance Model, a transparent trust model provides analysts, managers, and operators with greater confidence in agentic decision support systems by clearly communicating to the end user the reliability of the data source(s) that support the fusion output.

In addition to providing an evaluation scheme for determining how trustworthy an agent may be based on its past interactions with other agents, trust modeling in multi-agent organizations can also evaluate all types of agent behavior (e.g., cooperative and competitive) in multi-agent environments. Reputation-based systems indicate that agents can achieve decisions based upon their reliability (i.e., degree of trustworthiness), regardless of the presence of unreliable and/or adversarial agents. This occurs through the aggregation of many different responses from many different previous interactions an agent has had with other agents, thereby providing a more robust way of making decisions regarding the reliability of agents when the decisions are made within real-time operations (e.g., Cybersecurity monitoring of data networks and Enterprise supply chain networks (Granatyr et al., 2015). In addition, more recent research investigates models of trust based on learning rather than past interactions. In such trust models, agents can update their assessments of peers' reliability using RL, thereby enabling more robust systems that adapt to organizational changes over time (Buşoniu et al., 2006).

7.2 Uncertainty Representation

Uncertainty is inherent in information fusion due to sensor noise, incomplete observations, model limitations, and the heterogeneous nature of organizational data ecosystems. Accurate representation and management of uncertainty are therefore critical for reliable decision-making in agentic systems deployed across healthcare, enterprise, and cybersecurity domains. Probabilistic methods are commonly used for understanding and quantifying uncertainty in many areas, including localization, risk assessments (clinical), and predictive modeling of organizations(Kendall & Gal, 2017). The unknowns that arise from insufficient information or conflicting information can be difficult to represent within a traditional probability framework. Dempster-Shafer Theory addresses this limitation by providing a new way to quantify uncertainty using belief masses over sets of hypotheses. This theory is a very useful method for performing multi-source data fusion due to the possibility that the data being used may be incomplete or contradictory (X. Sun & Zeng, 2025).

Recent developments in uncertainty-aware neural technologies have resulted in methods based on deep learning to determine both an estimate of predictive uncertainty as well as the level of confidence associated with the predictions made by a model using the new techniques of Bayesian neural networks and ensembles (Kumar, 2025). In addition to improving the reliability of agentic fusion systems, developing these types of models also enables end users (e.g., analysts and operators) making high-stakes organizational decisions to have greater confidence in the predictions made by agentic fusion systems under conditions of uncertainty through the use of models such as medical image fusion evaluation; financial risk assessment; and real-time responses to cyber threats.

7.3 Decision Coupling and Decision Intelligence

Decision coupling is a defining characteristic of agentic fusion systems, distinguishing them from traditional fusion approaches, in which fusion and decision-making are treated as separate stages. In contrast, agentic fusion tightly integrates perception, reasoning, and action, enabling continuous interaction between data integration and organizational decision execution (Park et al., 2023). In Agentic Systems, the two decision-making frameworks used to formalize agentic decision-making are Decision Theory (MDP) and Decentralized-POMDP. Fusion is the common belief state of every actor involved in the decision-making process, and Action is the agreed-upon plan of action that forms the basis for all of the organizational agents' actions (Oliehoek & Amato, 2015). Dynamically adjusting the Fusion component of the agentic decision-making process enables agentic systems to adapt their Fusion Strategy based on predicted outcomes of their decisions, thereby improving performance in three key operational areas of the organization, such as Financial Risk Monitoring, Clinical Decision Support, and Enterprise Operational Optimization.

For non-technical organizational stakeholders, including senior managers, compliance officers, and clinical administrators, the interpretability of agentic fusion outputs must be operationalized through user-facing explainability interfaces rather than algorithmic transparency alone. Practically, this means providing confidence scores expressed as plain-language probability ranges, natural-language rationales generated by language model components, and visual decision trails tracing which data sources most influenced a given recommendation. Such interfaces align with the Technology Acceptance Model construct of perceived usefulness: if a manager cannot interpret a system output, its usefulness and consequently its adoption are severely constrained regardless of underlying technical accuracy.

By promoting decision coupling, agentic systems can, in the future, infuse uncertainty and trust directly into the decision policy. Therefore, agents will be able to assess how confident they are in the accuracy of their sources of information before making an organizational commitment. For example, agents may choose to defer making a decision or gather more data if they do not have high confidence in their data source. They may also take decisive action if they are confident in their data source. Therefore, as agents continue to operate adaptively, they will improve organizational robustness and operational efficiency (Park et al., 2023). Another benefit of agentic fusion systems is that they enable human-in-the-loop (HITL) decision intelligence. Thus, human operators, including clinicians, analysts, and managers, can validate and interpret the system's decision-making before executing it. The ability for humans to interact and evaluate agentic decisions is critical in organizational domains where a regulatory requirement exists and where the consequences for the decisions are very high (Topol, 2019).

8. APPLICATIONS OF AGENTIC FUSION SYSTEMS

Agentic information fusion systems are increasingly deployed across diverse application domains where heterogeneous data integration, adaptive decision-making, and distributed intelligence are critical. Unlike traditional fusion approaches that primarily focus on data aggregation, agentic fusion systems operate as agentic decision support systems, integrating perception, reasoning, and action to enable real-time, context-aware decisions for analysts, managers, operators, and other organizational end users. Table 4 provides a domain-wise system-level comparison of architectures, data modalities, fusion methods, outcomes, and limitations across representative deployment contexts.

8.1 Autonomous Systems

Autonomous systems, including self-driving vehicles, robotics, and unmanned aerial systems, are among the most prominent application areas of agentic fusion. These systems rely on multi-sensor perception, integrating data from LiDAR, radar, cameras, GPS, and inertial sensors to construct a comprehensive understanding of the environment (Qian et al., 2025). The designs of today's autonomous platforms incorporate an agentic fusion structure that consists of distinct modules for perception, planning, and control operating together as cooperating agents to implement this cooperation, we utilize a probabilistic method combined with a deep-learning architecture within a multi-modal transformer fusion processing pipeline collected from multiple sensors over space and time (Hussain et al., 2024). The agentic fusion of LiDAR and camera-based object detection will improve the overall accuracy of object location analysis in difficult situations, thereby illustrating how agentic fusion transforms raw sensor data into useful navigation information (Wang et al., 2025).

8.2 Healthcare Decision Support Systems

Healthcare is a critical organizational domain in which agentic fusion systems support clinical decision-making by integrating multimodal patient data. These systems combine information from electronic health records (EHRs), medical imaging, laboratory tests, wearable sensors, and genomic data to provide comprehensive insights into patient health (X. Chen et al., 2024). Most agentic fusion architectures in the healthcare domain will use centralized or federated architecture, where multiple sources of data can be processed (using deep learning and probabilistic models) in a multimodal diagnostic system (using both fused clinical and imaging data) in order to improve disease diagnosis/prognosis (Sharma et al., 2022). Bayesian and evidence-based fusion techniques help address uncertainty and variability associated with medical data and provide consistent and interpretable results that meet the needs of clinical decision making for physicians, nurses, and healthcare administrators (Lin & Xie, 2021).

8.3 Smart Cities and IoT Systems

Smart cities and IoT ecosystems are large-scale organizational environments in which agentic fusion systems enable real-time monitoring, control, and optimization of urban infrastructure. These systems integrate data from sensors, cameras, environmental monitors, and communication networks to support traffic management, energy optimization, public safety, and urban service delivery (Zaman et al., 2024). Distributed and hierarchical architectures are often used within these environments, where edge-based devices execute local fusion and cloud-based systems facilitate the overall organizational decision-making as a system to enable the performance of transportation management systems by collecting many forms of input data (e.g., road sensors, GPS, and video) to determine traffic flow and to minimize congestion in real time (H. Chen & Wu, 2025). Distributed agent-based fusion systems in IoT environments utilize event-driven methods of communication and consensus mechanisms between agents to reduce the overall overhead of communication while allowing agents to share relevant information (Amin & Hossain, 2021).

8.4 Cybersecurity and Enterprise Decision Systems

Cybersecurity and enterprise environments represent increasingly important organizational application domains for agentic fusion systems, where decision intelligence is critical for detecting threats, managing financial risks, and optimizing operations. The systems implement integrated information from: network logs, user activity, system operations, transaction history, behavior analytics, and external threat intelligence sources, allowing users to detect anomalies and to assist in making decisions by the organization (Möller, 2023). Cybersecurity systems that employ distributed multi-agent architectures allow for collaboration of multiple agents, where each agent has its own area of responsibility in monitoring aspects of the network and detecting suspicious activity through the combination of statistical analysis, machine learning and rule-based fusion processes, with multi-agent co-ordination allowing for the creation of alerts and for rapid response actions to include isolation of infected network nodes and updates to the existing security policy (Halbouni et al., 2022). Enterprise decision support systems employ agentic fusion to create an integrated financial system, supply chain, and customer interaction records, while fraud detection systems use transaction records, behavior analytics, and external indicators to achieve fraud detection rates that exceed 99% accuracy (Tao & Chao, 2023).

Table 4. Domain-Wise System-Level Analysis of Agentic Information Fusion Architectures

Domain	System	Architecture	Data Modalities	Fusion Method	Outcome	Limitation
Autonomous Systems	Multi-sensor autonomous driving perception system	Distributed / Hybrid	LiDAR, radar, camera, GPS, IMU (Qian et al., 2025; Wang et al., 2025)	Deep multimodal fusion + probabilistic fusion + MARL (Buşoniu et al., 2006)	Improved perception accuracy, robust environment understanding	High computational cost, latency constraints, communication overhead
Healthcare DSS	Multimodal clinical decision support system	Centralized / Federated	EHR, medical imaging, lab data, wearables, genomics (Hussain et al., 2024)	Bayesian + evidence-based + deep learning fusion (Lin & Xie, 2021)(X. Chen et al., 2024)	Enhanced diagnosis, personalized treatment, decision support	Privacy concerns, interpretability issues, regulatory constraints
Smart Cities / IoT	Intelligent traffic & urban monitoring system	Distributed / Hierarchical	Sensors, cameras, GPS, environmental data (Sharma et al., 2022) (Zaman et al., 2024)	Consensus-based + event-driven + adaptive fusion (Pinto Neto et al., 2023)	Real-time traffic optimization, energy efficiency, urban intelligence	Scalability challenges, heterogeneous data reliability, security risks
Cybersecurity Systems	Intrusion detection and threat intelligence system	Distributed multi-agent	Network logs, user behavior, system events, threat data (H. Chen & Wu, 2025)	ML/DL-based fusion + rule-based + anomaly detection (Möller, 2023)	Improved threat detection, rapid response, anomaly identification	False positives, adversarial attacks, model generalization issues
Enterprise Decision Systems	Fraud detection & business intelligence system	Hybrid / Centralized multi-agent	Transactional data, behavioral analytics, external signals (Halbouni et al., 2022)	Statistical + ML + trust-aware fusion (X. Sun & Zeng, 2025) (Tao & Chao, 2023)	Accurate fraud detection, risk assessment, operational optimization	Data bias, scalability issues, explainability limitations

9. COMPARATIVE SYNTHESIS OF AGENTIC FUSION SYSTEMS

Comparative evaluation is essential for understanding the performance characteristics and trade-offs of agentic information fusion systems across different architectural and operational settings. This section presents a structured qualitative comparative synthesis drawing on performance characterizations reported across the reviewed literature. A formal numerical meta-analysis is precluded by the heterogeneity of reported metrics, evaluation contexts, and experimental designs across included studies, a known limitation reflecting the early maturity of this research area. The synthesis identifies consistent patterns in accuracy, latency, scalability, and robustness trade-offs across centralized, decentralized, hierarchical, and hybrid agentic fusion architectures. Qualitative labels such as High and Moderate reflect consensus characterizations derived from the frequency and consistency of reporting across included studies, with supporting citations provided. This synthesis enables organizational designers and system architects to make informed architecture selection decisions aligned with specific operational requirements.

9.1 Comparative Metrics Analysis

Performance evaluation of agentic fusion systems is typically based on three core metrics: accuracy, latency, and robustness. Deep fusion combined with RL improves the ability to gather, integrate, and assess multimodal organizational data by increasing the effectiveness of information sources and associated decision-making strategies (Gao et al., 2020). Multimodal fusion has improved both detection and prediction accuracy compared to unimodal baselines in a number of areas, including autonomous systems, health care decision support, medical image fusion, and enterprise analytics. As a result, analysts, clinicians, and risk managers will have access to higher-quality output. Latency, which is the time it takes to process data and execute decisions based on that data, has significant implications for a variety of applications, including autonomous driving, industrial control, and financial fraud detection. Centralized systems provide high accuracy but suffer from bottlenecks in communication and processing (G. Chen et al., 2021), while distributed and edge-based systems have lower latency for local fusion, allowing organizations to respond more quickly. Low-latency systems may have less complete representations of the organization's state globally. This has the potential to affect the quality of decisions made by managers who require a comprehensive situational awareness (Dautov & Distefano, 2017).

9.2 Scalability Trends

Scalability is a critical factor for agentic fusion systems deployed in large-scale organizational environments such as IoT networks, smart cities, and enterprise systems. Centralized architectures exhibit limited scalability as the number of data sources increases, with the central node becoming a performance constraint, leading to degraded organizational efficiency (G. Chen et al., 2021). Distributed architectures offer the advantage of horizontal scaling, allowing agents to be added as needed while maintaining performance. Large-scale deployments of the Internet of Things (IoT) show that distributed fusion can significantly reduce the amount of data that must be transmitted and provide a means for parallel processing across an organization's nodes (Kulkarni et al., 2024). Hierarchical Architectures use local processing with accumulation at higher levels to provide vertically scaled systems with minimal communication overheads and greater global coordination (Tsanousa et al., 2022). Hybrid cloud/edge architectures facilitate large-scale Agile processes by providing real-time edge processing with cloud-based model training and strategic reasoning within the enterprise business functions (P. Sun et al., 2025).

9.3 Communication–Performance Trade-off

Communication efficiency is a key factor influencing the performance of agentic fusion systems, particularly in distributed and multi-agent organizational environments. In centralized frameworks, communication is primarily unidirectional toward a single processing node, simplifying coordination but creating overhead and bottleneck risks that limit organizational scalability (Noh, 2020). Distributed architectures reduce reliance on a single node through a peer-to-peer model through a peer-to-peer model; however, it is now challenging to coordinate data between organizational nodes due to consistency issues with synchronized data (Dautov & Distefano, 2017). Researchers have developed communication methods that allow you to send messages using event-driven systems, with compression and the sharing of only essential information to limit the passing along of unnecessary messages, while retaining the information needed to support quality organizational decisions (Berna-Koes et al., 2004). Learning-based communication methods enable agents to configure their communication policies based on performance and efficiency trade-offs and to optimize the balance between information richness and bandwidth use (Buşoniu et al., 2006). Table 5 provides a cross-architectural comparison of representative agentic fusion approaches.

Table 5. Comparative Performance Analysis of Representative Agentic Fusion Studies Across Architectures

Study	Architecture	Agents	Accuracy	Latency	Robustness
Multi-sensor autonomous driving fusion (Qian et al., 2025 ; Wang et al., 2025)	Distributed / Hybrid	Multi-agent (vehicle + sensors)	High (enhanced perception via multimodal fusion)	Low–Moderate (edge processing)	High (sensor redundancy + adaptive fusion)
RL-based multimodal fusion decision system (C. Li et al., 2025)	Hybrid / Distributed	Multi-agent	High (adaptive policy learning improves decision accuracy)	Moderate (learning overhead)	High (adaptive strategy under dynamic conditions)
Deep multimodal fusion frameworks (Gao et al., 2020)	Centralized / Distributed	Single / multi-agent	Very High (representation learning across modalities)	Moderate–High (computational cost)	Moderate (limited adaptability without retraining)
Distributed IoT data fusion systems (Ding et al., 2019) (Jennings, 2002)	Distributed	Large-scale multi-agent	Moderate–High	Low (local processing)	High (fault tolerance, decentralized processing)
Federated healthcare fusion systems (Kulkarni et al., 2024) (Hussain et al., 2024)	Federated / Hybrid	Multi-agent (institution-level)	High (privacy-preserving multimodal learning)	Moderate	High (privacy + distributed robustness)
Multi-agent coordination & consensus fusion (Dautov & Distefano, 2017) (Dorri et al., 2018)	Distributed	Multi-agent	Moderate–High	Low–Moderate	High (consensus + cooperative robustness)

continued on following page

Table 5. Continued

Study	Architecture	Agents	Accuracy	Latency	Robustness
Trust-aware fusion systems (Mendelson & Zwillinger, 2024) (X. Sun & Zeng, 2025)	Distributed / Hybrid	Multi-agent	High (reliability-weighted decisions)	Moderate	Very High (handles unreliable/adversarial data)
Centralized enterprise decision systems (Tao & Chao, 2023) (G. Chen et al., 2021)	Centralized	Single / limited agents	High	High (communication bottleneck)	Moderate (single point of failure)

Note. Performance labels are qualitative consensus characterizations derived from frequency and consistency of reporting across reviewed studies. No numerical aggregation is implied. Supporting citations are provided in the accompanying text.

10. ETHICAL, TRUST, AND GOVERNANCE CONSIDERATIONS

Agentic information fusion systems operate in high-stakes, data-intensive organizational environments where decisions directly influence operational outcomes, safety, and societal impact. As these systems transition from static pipelines into fully autonomous, decision-centric infrastructures deployed across healthcare, enterprise analytics, supply chain monitoring, and cybersecurity domains, ethical, trust, and governance aspects become integral components of system design rather than supplementary concerns. The European Union AI Act (2024) provides a risk-based framework that classifies all agentic fusion systems (e.g., in healthcare, autonomous systems, and critical infrastructure) as high-risk. Therefore, these technologies must have transparency, traceability, ongoing human oversight, and continuous monitoring by human operators (van Kolfschooten & van Oirschot, 2024).

Beyond these general frameworks, agentic fusion systems deployed in specific sectors face additional compliance layers. In healthcare, the Health Insurance Portability and Accountability Act imposes strict constraints on protected health information use and disclosure, directly shaping the design of federated fusion architectures where patient data must remain within institutional boundaries. FDA guidance on Software as a Medical Device further requires analytical and clinical validation before deployment of clinical decision support systems. In the financial sector, the Digital Operational Resilience Act, enacted across EU financial institutions, mandates ICT risk management frameworks encompassing AI-based decision systems, requiring documented incident response procedures and third-party risk assessments for agentic components. These sector-specific compliance layers are complementary to, and in some cases more stringent than, the EU AI Act and NIST Risk Management Framework.

In addition, compliance will directly influence the design of agentic fusion architectures by requiring audibility of the decision-making process and explainability of the outputs from fusion, as well as quantifying uncertainty that is present within fusion by quantifying how likely it is that each fusion output is accurate, and providing decision-makers, analysts, and managers access to this measurable uncertainty. The NIST AI Risk Management Framework also provides an organized way to identify, assess, and mitigate the risks associated with AI during the lifecycle of the AI system (*AI Risk Management Framework | NIST*, n.d.). This aligns with the principles of agentic fusion design as it focuses on governance through mapping, measurement, and management of organizational risks.

The issue of bias in multi-source fusion systems poses a major ethical dilemma. If data from multiple organizational sources are used to make decisions, then each of those datasets will likely carry some form of bias stemming from factors such as non-representative sampling, historic inequities and inconsistencies between different measurement methodologies. Furthermore, fusion (or combination) of the data through learning will often exacerbate rather than mitigate any underlying biases (LI et al., 2024). In sectors such as healthcare and corporate analytics, this is an especially serious problem, as biased data distributions can lead to unfair or inaccurate organizational decisions that affect patients, customers, or employees. As such, bias-aware fusion methods that utilize fairness-constrained optimization, subgroup performance assessment, and continuous monitoring of organizational decision outcomes are critically important (Barocas et al., 2023).

From a Sociotechnical Systems Theory perspective, successful deployment of agentic fusion systems requires co-optimization of the technical subsystem, including agent architectures, learning algorithms, and coordination protocols, and the social subsystem, including human roles, organizational workflows, and accountability structures. A purely technical optimization that maximizes accuracy while ignoring decision-maker cognitive load, trust calibration needs, or regulatory accountability requirements will predictably fail in organizational adoption even when technically superior. This implies that agentic DSS designers must engage information systems scholars and organizational stakeholders as co-designers from the earliest stages of system specification, ensuring that sociotechnical joint optimization governs design tradeoff decisions.

Human-in-the-loop accountability for agentic decision-making support systems is an important factor in their responsible deployment. Organizations use these systems for high-stakes applications, such as in healthcare, finance, cybersecurity, and critical infrastructure. Human intervention is necessary to validate or interpret system-generated outputs, or to override the system's decisions when necessary (Topol, 2019). When agentic systems offer interpretable outputs, confidence metrics, and justifications for decision-making, clinicians, analysts, risk managers, and compliance officers will have increased confidence in utilizing these systems.

11. CHALLENGES AND FUTURE DIRECTIONS

Despite significant advances, agentic information fusion systems face several open challenges that must be addressed to realize scalable, trustworthy, and general-purpose agentic decision support systems. These challenges span architectural scalability, trust calibration, explainability, integration with foundation models, and the development of unified fusion intelligence frameworks applicable across diverse organizational contexts.

Scalability presents a key architectural challenge for agentic fusion systems in large-scale organizational environments. The specific research question to be addressed is: how can multi-agent coordination protocols maintain coherent global decision intelligence as agent counts scale beyond thousands in enterprise IoT and smart city deployments? The recommended methodological approach involves simulation-based benchmarking using established large-scale MARL testbeds, comparing communication-efficient protocols such as sparse attention and event-triggered communication against accuracy-latency trade-off curves at increasing agent counts (Dautov & Distefano, 2017). The expected contribution is a set of evidence-based design guidelines for scalable agentic fusion, deployable in national-scale smart city and enterprise infrastructure.

Trust calibration is a critical challenge in dynamic, heterogeneous organizational environments. The specific research question is: how can trust models accurately estimate and dynamically update source reliability weights under adversarial conditions and unpredictable data quality fluctuations (Granatyr et al., 2015; Jøsang et al., 2007)? The recommended approach involves longitudinal field studies combined with reinforcement learning-based trust calibration experiments in simulated adversarial conditions (Buşoniu et al., 2006). The expected contribution is an adaptive, context-aware trust calibration framework validated under realistic organizational data quality conditions.

Explainability is a foundational deployment requirement across healthcare, financial, and cybersecurity organizational domains. The specific research question is: how can causal inference and symbolic reasoning be integrated into agentic fusion systems to provide decision-makers with intelligible, quantifiably confident justifications (van Kolfschooten & van Oirschot, 2024)? The recommended methodological approach combines neuro-symbolic architecture experiments with user studies measuring decision quality and trust among organizational stakeholders with varying technical backgrounds (Schick et al., 2023). The expected contribution is an explainability framework bridging technical interpretability and managerial usability requirements, validated across healthcare and financial deployment contexts.

Foundation model integration presents both opportunities and challenges for agentic fusion. The specific research question is: how can large-scale multimodal foundation models be efficiently integrated into real-time agentic fusion pipelines without prohibitive computational cost or unacceptable latency (Schick et al., 2023)? The recommended approach involves model compression experiments, edge deployment studies, and domain-specific fine-tuning validation (Murphy, 2022). The expected contribution is a deployment framework enabling foundation model-powered agentic fusion on resource-constrained edge devices, validated in autonomous vehicle and industrial IoT contexts.

12. CONCLUSION

This study presented a comprehensive review of agentic artificial intelligence for information fusion, highlighting its evolution from traditional static fusion pipelines to adaptive agentic decision support systems serving diverse organizational needs across healthcare decision support, medical image fusion, enterprise analytics, cybersecurity, smart cities, and autonomous systems. A systematic search of four major bibliographic databases, namely Scopus, Web of Science Core Collection, IEEE Xplore, and ScienceDirect, yielded 1,205 records, from which a final synthesis corpus of 380 unique publications was identified through rigorous PRISMA-guided screening and cross-database deduplication using Rayyan systematic review software. Classical fusion approaches exhibit the fundamental static pipeline limitations outlined in Section 3.4, including lack of autonomy, scalability, and real-time decision integration. In contrast, agentic fusion systems integrate perception, reasoning, and action within the unified closed-loop framework described in Section 4.1, enabling continuous learning, distributed coordination, and context-aware decision-making aligned with organizational objectives and end-user requirements. A key contribution is the development of a multi-dimensional taxonomy that captures architecture, learning paradigms, coordination mechanisms, trust modeling, and uncertainty representation, validated through an inter-rater reliability assessment with an overall composite Cohen's kappa ranging from 0.78 to 0.86 across six taxonomy dimensions. In healthcare, agentic fusion integrates electronic health records, medical imaging, wearable sensors, and genomic data within federated and hybrid architectures, enabling personalized clinical decision support, adaptive treatment optimization, and privacy-preserving distributed learning across institutional boundaries.

The analysis demonstrates that modern agentic decision support systems increasingly favor decentralized and adaptive designs to address organizational challenges of scalability, robustness, and real-time responsiveness. Based on the Technology Acceptance Model, Task-Technology Fit, and Sociotechnical Systems Theory, this review reinforces the need for agentic decision support systems to align with end-user workflows, organizational decision contexts, and human oversight requirements to achieve meaningful adoption. Ethical, regulatory, and managerial dimensions, including bias mitigation, privacy preservation, and explainable decision outputs, are foundational requirements for responsible organizational deployment. Managers must ensure trust through transparency, maintain human oversight in high-stakes decisions, and align fusion capabilities with organizational workflows. In the end, data fusion is no longer a standalone computational function but a central mechanism of adaptive, autonomous, and human-centered organizational decision-making, positioning agentic AI as a defining paradigm for next-generation intelligent organizational systems.

DATA AVAILABILITY STATEMENT

The full list of 380 included studies is publicly available at https://doi.org/10.17605/OSF.IO/SQ87G. Foundational background references cited in the manuscript are retained as theoretical grounding and are not part of the synthesis corpus, consistent with Section 2. The full coded corpus is available from the corresponding author upon reasonable request.

COMPETING INTERESTS STATEMENT

The authors of this publication declare there are no competing interests.

FUNDING STATEMENT

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Funding for this research was covered by the authors of the article.

CORRESPONDING AUTHOR

Correspondence should be addressed to Shabbir Syed-Abdul: drshabbir@tmu.edu.tw

REFERENCES

AI Risk Management Framework | NIST. (n.d.). Retrieved March 3, 2026, from https://www.nist.gov/itl/ai-risk-management-framework

Amin, S. U., & Hossain, M. S. (2021). Edge Intelligence and Internet of Things in Healthcare: A Survey. *IEEE Access : Practical Innovations, Open Solutions*, *9*, 45–59. DOI: 10.1109/ACCESS.2020.3045115

Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning : limitations and opportunities*. MIT Press. https://fairmlbook.org/

Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A Neural Probabilistic Language Model. *Journal of Machine Learning Research*, *3*, 1137–1155. DOI: 10.5555/944919.944966

Berna-Koes, M., Nourbakhsh, I., & Sycara, K. (2004). Communication efficiency in multi-agent systems. *Proceedings - IEEE International Conference on Robotics and Automation, 2004*(3), 2129–2134. DOI: 10.1109/ROBOT.2004.1307377

Buşoniu, L., Babuška, R., & De Schutter, B. (2006). Multi-agent reinforcement learning: A survey. *9th International Conference on Control, Automation, Robotics and Vision, 2006, ICARCV '06*. DOI: 10.1109/ICARCV.2006.345353

Chen, G., Liu, Z., Yu, G., & Liang, J. (2021). A New View of Multisensor Data Fusion: Research on Generalized Fusion. *Mathematical Problems in Engineering*, *2021*(1), 5471242. DOI: 10.1155/2021/5471242

Chen, H., & Wu, Y. (2025). Intelligent traffic signal optimization algorithm based on multi-source data fusion. *Intelligent Decision Technologies*, *19*(6), 4346–4355. DOI: 10.1177/18724981251369386

Chen, X., Xie, H., Tao, X., Wang, F. L., Leng, M., & Lei, B. (2024). Artificial intelligence and multimodal data fusion for smart healthcare: topic modeling and bibliometrics. *Artificial Intelligence Review, 57*(4), 91-. DOI: 10.1007/s10462-024-10712-7

Das, A., Gervet, T., Romoff, J., Batra, D., Parikh, D., Rabbat, M., & Pineau, J. (2019). *TarMAC: Targeted Multi-Agent Communication* (pp. 1538–1546). PMLR. https://proceedings.mlr.press/v97/das19a.html

Dautov, R., & Distefano, S. (2017). Distributed Data Fusion for the Internet of Things. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 10421 LNCS*, 427–432. DOI: 10.1007/978-3-319-62932-2_41

Ding, W., Jing, X., Yan, Z., & Yang, L. T. (2019). A survey on data fusion in internet of things: Towards secure and privacy-preserving fusion. *Information Fusion*, *51*, 129–144. DOI: 10.1016/j.inffus.2018.12.001

Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-Agent Systems: A Survey. *IEEE Access : Practical Innovations, Open Solutions*, *6*, 28573–28593. DOI: 10.1109/ACCESS.2018.2831228

Foerster, J., Assael, I. A., de Freitas, N., & Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. Advances in Neural Information Processing Systems, 29. https://proceedings.neurips.cc/paper_files/paper/2016/file/c7635bfd99248a2cdef8249ef7bfbef4-Paper.pdf

Gao, J., Li, P., Chen, Z., & Zhang, J. (2020). A Survey on Deep Learning for Multimodal Data Fusion. *Neural Computation*, *32*(5), 829–864. DOI: 10.1162/neco_a_01273 PMID: 32186998

Granatyr, J., Botelho, V., Lessing, O. R., Scalabrin, E. E., Barthes, J. P., & Enembreck, F. (2015). Trust and reputation models for multiagent systems. *ACM Computing Surveys*, *48*(2), 27. DOI: 10.1145/2816826

Halbouni, A., Gunawan, T. S., Habaebi, M. H., Halbouni, M., Kartiwi, M., & Ahmad, R. (2022). Machine Learning and Deep Learning Approaches for CyberSecurity: A Review. *IEEE Access : Practical Innovations, Open Solutions*, *10*, 19572–19585. DOI: 10.1109/ACCESS.2022.3151248

Hussain, M., O'Nils, M., Lundgren, J., & Mousavirad, S. J. (2024). A Comprehensive Review on Deep Learning-Based Data Fusion. *IEEE Access : Practical Innovations, Open Solutions*, *12*, 180093–180124. DOI: 10.1109/ACCESS.2024.3508271

Ignatious, H. A., El-Sayed, H., & Kulkarni, P. (2023). Multilevel Data and Decision Fusion Using Heterogeneous Sensory Data for Autonomous Vehicles. *Remote Sensing, 15*(9). DOI: 10.3390/rs15092256

Jennings, N. R. (2002). *Agent-Based Computing*. 17–30. DOI: 10.1007/978-0-387-35602-0_3

Jøsang, A., Ismail, R., & Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems*, *43*(2), 618–644. DOI: 10.1016/j.dss.2005.05.019

Kendall, A., & Gal, Y. (2017). What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? Advances in Neural Information Processing Systems, 30. https://proceedings.neurips.cc/paper_files/paper/2017/file/2650d6089a6d640c5e85b2b88265dc2b-Paper.pdf

Kulkarni, S., Dwivedi, J. N., Pramanta, D., & Tanaka, Y. (2024). Edge Computational Intelligence for AI-Enabled IoT Systems. *Edge Computational Intelligence for AI-Enabled IoT Systems*, 1–328. DOI: 10.1201/9781032650722

Kumar, P. S. (2025). Uncertainty Estimation in Deep Learning Based Computer Vision. *Applied Computer Vision through Artificial Intelligence*, 331–348. DOI: 10.1002/9781394272624.ch16

Landis, J. R., & Koch, G. G. (1977). The Measurement of Observer Agreement for Categorical Data. *Biometrics*, *33*(1), 159. DOI: 10.2307/2529310 PMID: 843571

Li, C., Lu, G., Chen, W., Xu, C., Li, L., & Wang, Z. (2025). Reinforcement Learning-Based Multimodal Data Fusion Timing Decision-Making. *2025 IEEE International Conference on Unmanned Systems (ICUS)*, 1401–1405. DOI: 10.1109/ICUS66297.2025.11295247

Li, J., Hong, D., Gao, L., Yao, J., Zheng, K., Zhang, B., & Chanussot, J. (2022). Deep learning in multimodal remote sensing data fusion: A comprehensive review. *International Journal of Applied Earth Observation and Geoinformation*, *112*, 102926. Advance online publication. DOI: 10.1016/j.jag.2022.102926

Li, X., Dunkin, F., & Dezert, J. (2024). Multi-source information fusion: Progress and future. *Chinese Journal of Aeronautics*, *37*(7), 24–58. DOI: 10.1016/j.cja.2023.12.009

Lin, Z., & Xie, J. (2021). Research on improved evidence theory based on multi-sensor information fusion. *Scientific Reports*, *11*(1), 9267. DOI: 10.1038/s41598-021-88814-3 PMID: 33927275

Mendelson, E., & Zwillinger, D. (2024). *Introducing Game Theory and its Applications*. Introducing Game Theory and Its Applications. DOI: 10.1201/9781003051275

Möller, D. P. F. (2023). Intrusion Detection and Prevention. *Advances in Information Security*, *103*, 131–179. DOI: 10.1007/978-3-031-26845-8_3

Murphy, K. P. (2022). Probabilistic Machine Learning: An Introduction. MIT Press. https://probml.github.io/pml-book/book1.html

Noh, S. (2020). Intelligent Data Fusion and Multi-Agent Coordination for Target Allocation. *Electronics, 9*(10), 1–13. DOI: 10.3390/electronics9101563

Oliehoek, F. A., & Amato, C. (2015). A Concise Introduction to Decentralized POMDPs. *Springer Briefs in Intelligent Systems*, 1–116. DOI: 10.1007/978-3-319-28929-8

Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative Agents: Interactive Simulacra of Human Behavior. *UIST 2023 - Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. DOI: 10.1145/3586183.3606763

Pinto Neto, E. C., Sadeghi, S., Zhang, X., & Dadkhah, S. (2023). Federated Reinforcement Learning in IoT: Applications, Opportunities and Open Challenges. *Applied Sciences, 13*(11). DOI: 10.3390/app13116497

Power, D. J. (2002). Decision Support Systems: Concepts and Resources for Managers. *Information Systems Management*, *20*(4), 80–81. DOI: 10.1201/9781482280180

Qian, H., Wang, M., Zhu, M., & Wang, H. (2025). A Review of Multi-Sensor Fusion in Autonomous Driving. *Sensors, 25*(19). DOI: 10.3390/s25196033

Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language Models Can Teach Themselves to Use Tools. *Advances in Neural Information Processing Systems*, *36*, 68539–68551. DOI: 10.52202/075280-2997

Senel, N., Kefferpütz, K., Doycheva, K., & Elger, G. (2023). Multi-Sensor Data Fusion for Real-Time Multi-Object Tracking. *Processes, 11*(2). DOI: 10.3390/pr11020501

Sharma, D. K., Chatterjee, M., Kaur, G., & Vavilala, S. (2022). Deep learning applications for disease diagnosis. *Deep Learning for Medical Applications with Unique Data*, 31–51. DOI: 10.1016/B978-0-12-824145-5.00005-8

Sukhbaatar, S., Szlam, A., & Fergus, R. (2016). Learning Multiagent Communication with Backpropagation. *Advances in Neural Information Processing Systems, 29*. https://proceedings.neurips.cc/paper_files/paper/2016/file/55b1927fdafef39c48e5b73b5d61ea60-Paper.pdf

Sun, P., Zhao, B., & Li, X. (2025). Research on multi-source heterogeneous data fusion method of substation based on cloud edge collaboration and AI technology. *Discover Applied Sciences, 7*(4), 262. DOI: 10.1007/s42452-025-06725-8

Sun, X., & Zeng, P. (2025). A DS Evidence Theory-Based Method for Multi-Source Information Fusion in Advanced Systems: Synergistic Intelligence. *Proceedings of 2024 4th International Conference on Big Data, Artificial Intelligence and Risk Management, ICBAR 2024*, 503–508. DOI: 10.1145/3718751.3718832

Sycara, K. P. (1998). Multiagent Systems. *AI Magazine*, *19*(2), 79–79. DOI: 10.1609/AIMAG.V19I2.1370

Tao, Z., & Chao, J. (2023). Financial Fraud and Anomaly Detection Techniques: A Literature Review. *ACM International Conference Proceeding Series*, 634–641. DOI: 10.1145/3644523.3644639

Topol, E. J. (2019). *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books.

Tsanousa, A., Bektsis, E., Kyriakopoulos, C., González, A. G., Leturiondo, U., Gialampoukidis, I., Karakostas, A., Vrochidis, S., & Kompatsiaris, I. (2022). A Review of Multisensor Data Fusion Solutions in Smart Manufacturing: Systems and Trends. *Sensors, 22*(5). DOI: 10.3390/s22051734

Turban, E., & Aronson, J. E. (1998). *Decision Support Systems and Intelligent Systems* (5th ed.). Prentice Hall.

van Kolfschooten, H., & van Oirschot, J. (2024). The EU Artificial Intelligence Act (2024): Implications for healthcare. *Health Policy (Amsterdam)*, *149*, 105152. DOI: 10.1016/j.healthpol.2024.105152 PMID: 39244818

Wang, H., Liu, J., Dong, H., & Shao, Z. (2025). A Survey of the Multi-Sensor Fusion Object Detection Task in Autonomous Driving. *Sensors, 25*(9). DOI: 10.3390/s25092794

Wooldridge, M. (2009). *An Introduction to MultiAgent Systems* (2nd ed.). Wiley.

Zaman, M., Puryear, N., Abdelwahed, S., & Zohrabi, N. (2024). A Review of IoT-Based Smart City Development and Management. *Smart Cities, 7*(3), 1462–1501. DOI: 10.3390/smartcities7030061

APPENDIX A (TABLE A1)

Key Terminology Used in This Review

Term	Definition as used in this paper
Agentic AI	Autonomous, goal-directed AI systems capable of closed-loop perception, reasoning, and action in dynamic environments
Information Fusion	The process of integrating heterogeneous multi-source data into a unified, decision-ready representation
Agentic DSS	A decision support system in which agentic AI serves as the fusion and reasoning engine, extending classical DSS with autonomy and closed-loop adaptation
Multi-Agent System	A collection of interacting autonomous agents; a structural substrate on which agentic AI operates
MARL	Multi-Agent Reinforcement Learning; agents learning joint policies through shared environmental reward signals
Trust Modeling	Dynamic assessment and weighting of data source or agent reliability within the fusion process
Decision Coupling	Tight integration of fusion and decision-making within a single closed-loop framework

Note. These definitions are applied consistently throughout the manuscript. Terms used interchangeably in the broader literature are assigned specific meanings as defined here.

APPENDIX B (TABLE B1). STANDARDIZED DATA EXTRACTION TEMPLATE

This template governs the data extraction protocol for all 380 studies included in the qualitative synthesis. Each study was independently coded by three expert raters across the twelve dimensions below. Inter-rater reliability was assessed using Cohen's kappa (composite k = 0.78-0.86). The template supports replication and transparency consistent with PRISMA 2020 standards.

Field Name	Description / Definition	Response Categories / Format	Notes
Study ID	Unique alphanumeric identifier assigned to each included study during the coding process.	e.g., S001, S002, ... S380	Assigned sequentially during data extraction
Authors and Year	Full citation reference for the included study, including all authors and publication year.	Last name(s) and year; e.g., Smith et al. (2023)	Follow APA 7th edition format
Application Domain	The primary organizational or technical domain in which the agentic fusion system is deployed or evaluated.	Healthcare / Autonomous Systems / Smart Cities / Cybersecurity / Enterprise Analytics / Other (specify)	Select the single most dominant domain; use Other if cross-domain
Architecture Type	The structural configuration describing how agents and fusion components are organized and distributed across the system.	Centralized / Distributed / Hierarchical / Federated / Hybrid Cloud-Edge / Other (specify)	Select all that apply; primary architecture listed first
Fusion Level	The stage at which data integration occurs within the system pipeline.	Data-Level / Feature-Level / Decision-Level / Multi-Level / Not Specified	Multi-Level if more than one level is integrated
Learning Paradigm	The machine learning approach employed by agents to adapt and improve fusion and decision strategies over time.	Reinforcement Learning / Multi-Agent RL / Self-Supervised Learning / Continual Learning / Supervised Learning / No Learning / Other (specify)	Select primary paradigm; note secondary paradigms in Notes column
Coordination Mechanism	The method by which multiple agents interact, share information, and align their actions toward collective organizational objectives.	Consensus-Based / Negotiation / Contract-Net / Emergent Communication / Market-Based / No Coordination / Other (specify)	Select all mechanisms present; note dominant mechanism first
Uncertainty Modeling	The framework or method used to represent, quantify, and manage uncertainty in data sources and fusion outputs.	Bayesian / Dempster-Shafer / Data-Driven / Hybrid / Implicit (attention/softmax) / None / Other (specify)	Note whether uncertainty is explicitly modeled or implicit
Trust Modeling	Whether the system incorporates explicit mechanisms to assess and update the reliability or credibility of data sources or agents.	Yes - Dynamic / Yes - Static / Partial / No	If yes, describe the trust mechanism type in Notes
Decision Coupling	The degree to which fusion and decision-making are integrated within a single closed-loop system rather than treated as sequential stages.	Sequential (post-fusion decision) / Perception-Reasoning-Action Loop / Closed-Loop Autonomous / Not Specified	Closed-Loop implies continuous feedback between fusion and action

continued on following page

This template governs the data extraction protocol for all 380 studies included in the qualitative synthesis. Continued

Field Name	Description / Definition	Response Categories / Format	Notes
Memory and Knowledge Representation	Whether the system maintains persistent internal representations of past states, learned knowledge, or organizational context to inform future decisions.	Yes - Explicit Knowledge Base / Yes - Learned Latent Representation / Partial / No / Not Specified	Describe representation type briefly in Notes if Yes
Empirical Validation	Whether the study provides empirical evidence or computational validation of the proposed agentic fusion system, as opposed to purely theoretical or conceptual contributions.	Yes - Real-World Deployment / Yes - Simulation / Yes - Benchmark Dataset / Partial / No (Conceptual Only)	Studies coded as No were excluded from the 380-study corpus per eligibility criteria

Note. The full list of 380 included studies is publicly available at https://doi.org/10.17605/OSF.IO/SQ87G. The full coded corpus is available from the corresponding author upon reasonable request.

Dr. Dawa Chyophel Lepcha holds a Bachelor of Engineering from the University of Mumbai, a Master of Engineering from Chandigarh University, Punjab, India, and a Master of Science from the University of South Wales, Cardiff, United Kingdom. He earned his PhD from Chandigarh University, where his doctoral research focused on advanced image processing algorithms at the intersection of computational imaging and machine learning. He is currently a Postdoctoral Research Fellow at the Graduate Institute of Biomedical Informatics, Taipei Medical University, Taiwan, where his research centres on trustworthy clinical AI for Alzheimer's disease neuroimaging. His work introduces quantitative frameworks for individual-level explainability validation, per-patient multimodal attribution, cross-site harmonisation robustness, and longitudinal biomarker trajectory prediction, forming a coherent research programme toward a deployable clinical AI system for real-time dementia staging with per-patient explainability. His broader research interests span artificial intelligence, deep learning, computer vision, medical image analysis, natural language processing, and generative AI. He has published extensively in high-impact SCI and SCIE indexed journals and serves as a reviewer for several prestigious peer-reviewed journals, reflecting his active contribution to the global research community.

Dr. Aaliya Ali is an interdisciplinary researcher and academician whose work integrates artificial intelligence with advanced biomedical and pharmaceutical applications. Her research primarily focuses on the design and development of innovative nanocarrier-based drug delivery systems alongside advanced medical imaging algorithms. In parallel, she actively investigates artificial intelligence and machine learning methodologies for healthcare applications, with particular emphasis on medical image analysis, computer vision, and AI-assisted disease detection. Her research aims to bridge pharmaceutical sciences with emerging AI-driven technologies to enhance diagnostic accuracy, enable personalized treatment strategies, and facilitate translational healthcare solutions. Dr. Ali has published extensively in reputed SCIE-indexed journals and holds patents in the domains of herbal nanoformulations and therapeutic innovations. Through her interdisciplinary approach, she is committed to developing next-generation intelligent healthcare systems that synergistically combine nanotechnology, natural therapeutics, and artificial intelligence to improve clinical outcomes and advance precision medicine.