Review

Agentic AI: A Review, Applications, and Open Research Challenges

Omer Khalid , Ammad Ul Haq Farooqi , Muhammad Bilal *

Posted Date: 8 December 2025

doi: 10.20944/preprints202512.0592.v1

Keywords: agentic ai, agentic architecture; multi-agent llms

 doi:10.20944/preprints202512.0592.v1

Article

Agentic AI: A Review, Applications, and Open Research Challenges

Omer Khalid , Ammad Ul Haq Farooqi , Muhammad Bilal *

Department of Software Engineering, National University of Computer and Emerging Sciences, Islamabad, 44000, Pakistan
* Correspondence: muhammad.bilal@isb.nu.edu.pk

Abstract

Agentic Artificial Intelligence (AI) marks a shift from traditional AI systems that simply generate responses to autonomous systems that can independently plan to achieve goals with minimal human intervention. These models can do much more than just respond to prompts as they can observe, adapt, coordinate with other agents, and even refine their own outputs over time. This literature review draws insights from fifty-one recent empirical studies on various domains to understand how agentic AI is being built and used today. Agentic AI systems appear in the domains of healthcare, digital twin architectures, educational platforms, e-commerce applications, cybersecurity systems, and large-scale network management systems and they often improve efficiency, reduce manual workload, and help in making more informed decisions. However, this increased autonomy also raises new questions as well because autonomous systems that can act without human intervention must be reliable, explainable, secure, and aligned with human expectations otherwise they may cause great harm to humans. Many implementations of such systems are still in early stages, lacking standard evaluation methods and are facing challenges such as data access, ethical responsibility, and coordination among multiple agents. For clearer understanding, this review outlines a taxonomy of agentic AI and it portrays several of its current application domains, discusses common architectures and techniques, and highlights its limitations and future directions. The results of this review suggest that progress in governance, multimodal reasoning, and scalable coordination will be central to advancing safe and useful agentic AI systems.

Keywords: agentic ai; agentic architecture; multi-agent llms

1. Introduction

In recent years, we see that the artificial intelligence has evolved from simply classifying data or generating responses into autonomous systems that can interact, reason and make decisions within complex and changing environments [1–3]. This progress or advancement has given rise to what is currently known as agentic AI systems that do not just respond to commands but they can also function as autonomous agents as well [4–6]. When we compare the agentic AI to the the traditional AI models that follow a set of predefined rules, we see that the agentic systems are built to interpret goals, to devise plans, to assess the results and work alongside both humans and other AI models when needed [2,7]. There are also studies which bridge a gap by offering a comprehensive comparative analysis of leading frameworks such as CrewAI, LangGraph, AutoGen, Semantic Kernel and MetaGPT [8]. While some literature provided us with a comprehensive overview of the emerging field of AI agentic programming [9].

The need for this move towards agentic AI systems mainly comes from real world applications where the situations or scenarios shift way too much quickly for constant human oversight and it just cant simply keep up for example the healthcare operations, the coordination for logistical operations and monitoring cybersecurity opeartions [10–12]. In such environments the conventional AI systems fall short because they might depend on a fixed set of instructions or they might depend on narrow

 doi:10.20944/preprints202512.0592.v1

contextual windows or they might depend upon continuous human supervision [1,6]. As explained earlier that the need for the Agentic AI came from the real world applications, so this timeline in the figure 1 shows how AI evolved from traditional to Agentic AI.

The Agentic AI systems on the other hand are designed to adjust their behavior as circumstances change [13] but the same features that make agentic AI so promising also bring serious concerns such as the people worry about whether these systems can be relied on, how predictable their behavior really is and to what extent we should trust them in situations that matter [4,14,15]. As these agentic AI systems become more capable and more widely used, it becomes increasingly important for the researchers, developers, and organizations to understand what this shift truly means. Only then can we make sure that autonomous AI is introduced in a way that it is responsible, safe and aligned with human needs and expectations [4,14,16,17].

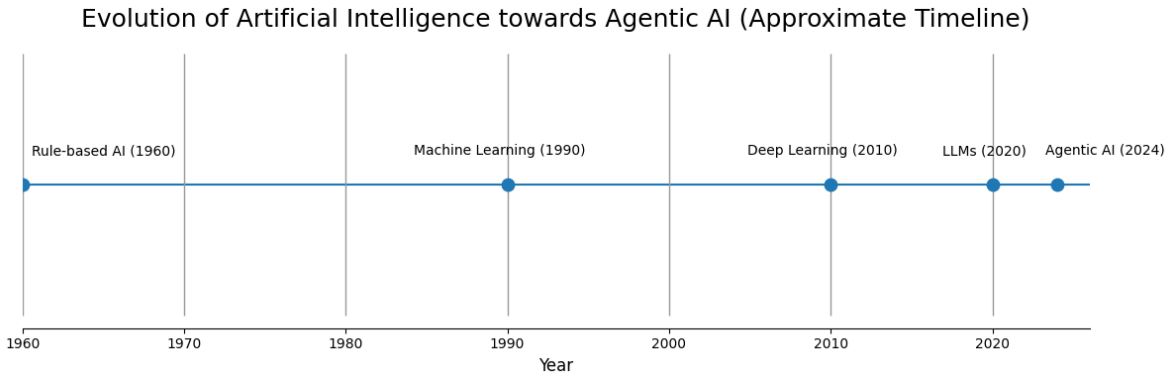

Figure 1. Timeline For Evolution of AI to Agentic AI.

In the literature we observe that the recent works of agentic AI approach and examine it from several different angles such as how these systems are designed, how they are coordinated, how multiple agents inside these systems work together and how these systems perform in the real world environments [18–20]. We observed that a growing number of studies highlight the use of both large and smaller language models to break down the reasoning tasks, to use external tools and collaborate with specialized agents when needed [20–22].

We also examined the frameworks such as LangGraph, AgentOps pipelines, intent-driven orchestration layers, and digital twin coordination platforms which show how the agentic AI can manage and execute the complex workflows in practice [18,23–25]. For example in the healthcare domain we observed the agentic systems are being developed to interpret clinical data, assist in the medical decision-making and track the patient conditions in real time [10,26–28], in the context of networking and large-scale infrastructure the agents in the agentic AI are used to automate the configuration tasks, identify security vulnerabilities and are used to manage the resources more efficiently [18,22,25,29]. There was some literature which provided us with critical analysis of the conceptual misuse of these terms such as Agentic AI and Multiagentic AI [30] while some literature identified the core challenges of LLMs such as challenges relating to security, privacy and trust, misinformation, misuse and bias, energy consumption, transparency and explainability, and value alignment, that can creep into Agentic AI systems [31].

In the educational focused domain we observed that the research is looking at how agentic AI can serve as a learning partner or peer facilitator in supporting the students rather than simply delivering information [32,33]. So despite the difference in domains we observed that the underlying aim remains the same and that is to move away from the systems which only provide fixed outputs and towards building such systems which are capable of making adaptive and context aware decisions [6,13]. Regardless of the ongoing progress in the field of agentic AI we see that many agentic AI implementations remain in early testing phases or they are limited to controlled environments. This

 doi:10.20944/preprints202512.0592.v1

shows that the field is still in a transitional phase and it is yet to reach a point of widespread and standardized deployment [5,14].

Across the finalized literature we have observed that although the progress in the field of agentic AI has been notable there are still many areas where the field lacks clarity and maturity. Firstly we would to add that there is no broadly accepted taxonomy or a shared terminology for describing the different forms of the agentic systems. This fact makes it difficult for the researchers to compare approaches or evaluate their effectiveness in a consistent way [1]. Addtionally while there are many studies which showcase the promising prototypes, there are far fewer studies which examine the long-term stability, the ethical oversight and how well these systems transfer across different application domains [3,13].

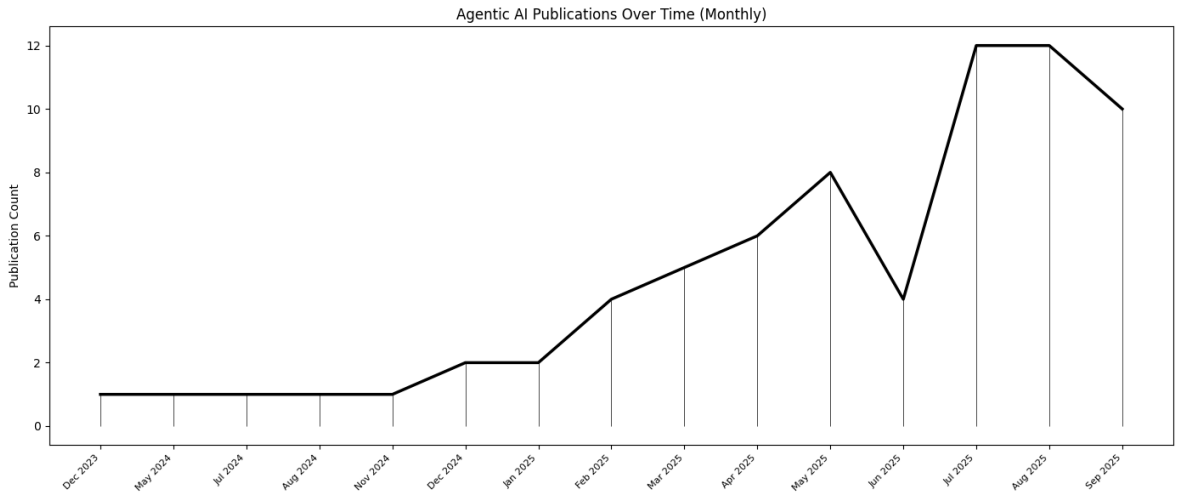

Figure 2. Distribution Of Papers Over Months

In addition, while there are many studies which showcase the promising prototypes far fewer studies examine the long-term stability, the ethical oversight or how well these systems transfer across different application domains [3,13]. We have also observed that many of the current implementations of agentic AI rely on narrowly tailored datasets, they rely on workflows which are carefully engineered and on controlled testing environments. Due to this the scalability and usefulness of such systems remain limited in real-world deployments [14,16,34]. We would also like to add that the increase in autonomy in such system raises questions about who is actually accountable for the decisions made by these systems and how reasoning in such systems can be understood and explained to the users and whether the users can trust these systems especially in the high-stakes contexts where the outcomes have real impact [27,32]. These mentioned constraints or challenges present a need for a more structured and comprehensive review which brings the existing research together in a clear way which explains that what are the research gaps that still require attention and how the agentic AI is currently being built and applied.

This study aims to provide the mentioned comprehensive review in a structured way as we have tried to map the current landscape of agentic AI such as we have tried to examine where and how it is being used as well as we have tried to identify the methods and limitations that are common across existing work of it and before the end of the study we also propose the research priorities to guide the development of autonomous systems that are safe, dependable and transparent in practice. There were also some studies which provided us with a comparative analysis across both AI agents and agentic AI paradigms. We observed that the application domains enabled by AI Agents such as customer support, scheduling, and data summarization are then contrasted with Agentic AI deployments in research automation, robotic coordination and medical decision support [35]. There were also some studies that had contributed to the understanding of agentic AI frameworks by offering practical insights and they set the ground work for further research into the applications of agentic AI in (small, medium,

 doi:10.20944/preprints202512.0592.v1

and micro-enterprises) SMMEs' dynamic and resource-constrained economic environments [36]. We also obeserved some literature addressing the ethical challenges posed by Agentic AI, in which they propose solutions for goal alignment, resource constraints and environmental adaptability [37]. In the literature we had also observed that there is enthusiasm for agentic models for their usage in medical radiology field and there is also a pressing need to optimize workflows, enhance clinical decision making, reduce radiologist cognitive load, and improve time-to-diagnosis for critical findings [38].

Many taxonomies and survey papers on agentic AI remain domain-specific and do not provide a cross-domain comparison of methods, domains and open challenges, limiting their usefulness for generalizable system designs [35,39,40]. These papers focus on edge intelligence, communications, elderly care, and scientific discovery but they only talk about agentic AI features itself and rarely analyze its orchestration, governance, or multi-agent coordination holistically [41–44]. Many of these papers highlight challenges like ambiguity in definitions, lack of protocol standardization, and fragmented architectures, yet they fail in synthesizing a consolidated state of the art view across application domains [45–47]. These surveys also emphasize the absence of benchmarking standards, safe operational practices, and end-to-end lifecycle governance [48,49]. Unlike these studies, our taxonomy integrates findings across domains, methods and techniques, datasets, limitations, and future work producing a comprehensive taxonomy synthesis that captures both technical and operational dimensions.

The diagram in figure 2 highlights the distribution of papers over the months of December 2023 to September 2025. This distribution is according to the currently available studies in our review/taxonomy.

The rest of the paper is arranged in a way that reflects how the taxonomy developed during the review process. In section 4, we describe how the fifty-one papers were gathered and analyzed. We essentially show which databases were searched, which keywords were used and which study counted as relevant enough to get included. In section 5 we then show the taxonomy itself and explain the main ideas or dimensions that separate different kinds of agentic systems from one another. From there, the discussion moves into what these systems look like in practice. In section 6 we look at where the agentic AI is currently being used for example in healthcare, networking and infrastructure management, e-commerce settings, digital twin systems and even education etc. In the section 7 we shift our focus to how these systems actually work under the hood such as what are the planning components, how the multiple agents coordinate with each other and how models connect with external tools. In section 8 we then review the datasets and data sources that the field of agentic AI relies on. In the section 9 we then look at the main challenges that keep appearing across the literature, the technical issues, the deployment difficulties and the concerns which are tied to ethics and trust. In section 10 we then offer some suggestions for where future research in the field of agentic AI could be headed, based on the gaps in the section 9. Lastly, in the section 11 we conclude the review.

2. Materials and Methods

The authors carried out this review using a practical, stepwise method to find and evaluate work on agentic AI. In all, 51 empirical papers were included. To collect those studies, they searched several familiar academic databases, IEEE Xplore, the ACM Digital Library, ScienceDirect, SpringerLink, arXiv, and Google Scholar, using search phrases that reflect how people in the field actually talk about these systems. Examples include agentic AI, autonomous AI systems, LLM multi-agent collaboration, tool-using AI, digital twin agents, and AI workflow automation. From each study, details were pulled that help show how agentic AI is currently being developed and used. This included the domain where the system was applied, the type of model or coordination framework used, the data sources supporting the system, and any strengths or limitations noted by the authors. The review also paid close attention to recurring concerns such as scalability, reliability, explainability, and the nature of interaction between humans and the system, since these issues appeared across many papers. Once collected, these details were compared and grouped so that common patterns could be seen more

 doi:10.20944/preprints202512.0592.v1

clearly. Those recurring patterns formed the basis of the taxonomy discussed later in the paper. The intention was not simply to summarize previous research, but to show how current agentic AI systems are being built, where they tend to overlap, and where substantial gaps or unanswered questions remain as the field continues to develop.

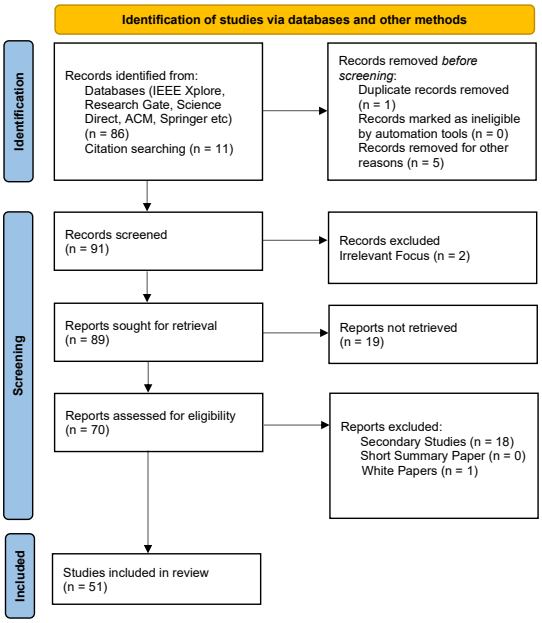

Figure 3. PRISMA Flow Diagram illustrating the study selection process.

The PRISMA diagram summarizes identification, screening, eligibility and inclusion steps and matches the counts reported in figure 3.

The paper selection followed a PRISMA-style [50] screening process, summarized as follows:

- Initial Search: Approximately 97 papers were identified through database searches and citation chaining.
- Screening by Title and Abstract: Papers unrelated to agentic AI, general machine learning, or purely theoretical models were removed.
- Full-Text Review: Papers that discussed autonomous agent behavior, multi-agent coordination, or AI-driven decision workflows were retained.
- Final Inclusion: 51 papers were included for full analysis and synthesis.

 doi:10.20944/preprints202512.0592.v1

3. Taxonomy of Agentic AI and Its Current State of the Art

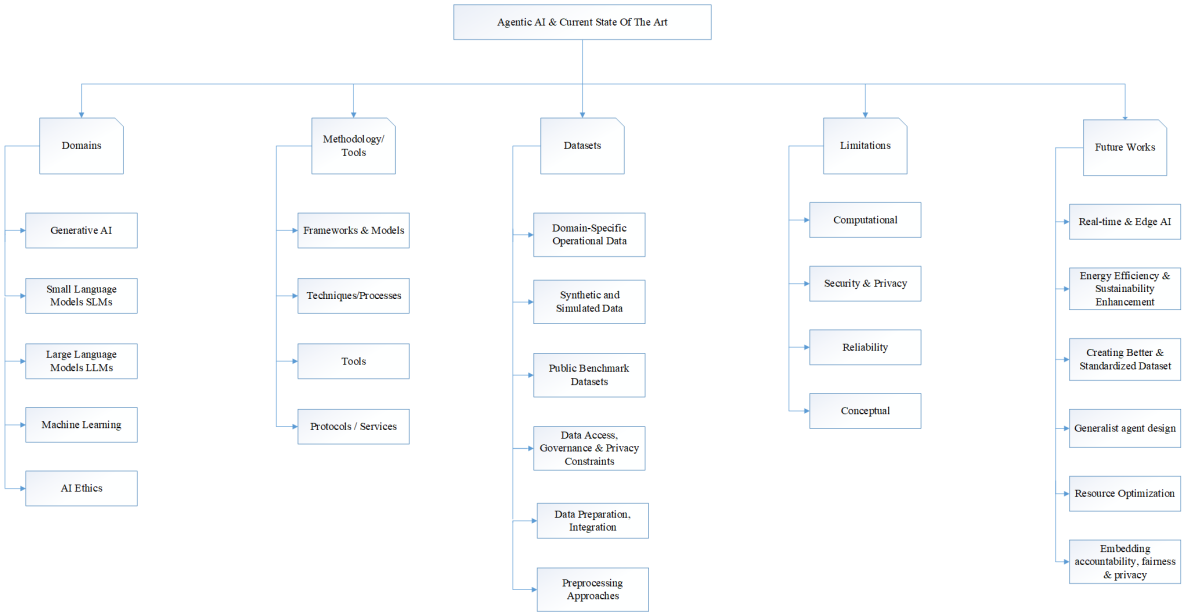

Figure 4. Proposed Taxonomy Diagram of Agentic AI.

The taxonomy groups the literature into five main dimensions: (1) Application Domains, (2) Methods and Techniques, (3) Data Sources and Datasets, (4) Limitations, and (5) Future Work. This has been illustrated in the figure 4.

4. Application Domains

Agentic AI is not a single-use idea it's appearing across many real-world fields where autonomy, continuous adaptation, and multi-step decision making are valuable. Below we grouped the 51 papers by their primary application domains, summarize what agentic systems are doing in each area, and highlight recurring benefits and domain-specific challenges. All syntheses below are drawn from the set of 51 studies highlighted above.

4.1. Healthcare (diagnostics, monitoring, clinical workflows)

Agentic systems are used to automate end-to-end clinical pipelines, detect early signs of illness from medical notes and sensors data, and support medical personnel with explainable outputs and workflow suggestions [10,26,27]. Examples include agentic inference pipelines for multimodal medical data, workflows for detecting cognitive concerns, and personalized transplant-monitoring agents [10,27], common benefits are faster sorting of and allocation of treatment to patients, continuous monitoring, and scalable deployment, while recurring concerns are privacy, regulatory compliance, and the need for clinician-in-the-loop validation [22,27,28].

4.2. Networking, Telecom and 6G (intent-driven orchestration and adaptive networks)

Papers show agentic AI acting as intent translators and auto-orchestrators for network infrastructure which helps in decomposing high-level operator goals into IaC or runtime actions, optimizing spectrum/ bandwidth, and autonomously reacting to faults [18,51–53]. AgentNet-style GFMs and intent-based orchestration demonstrate improved resource utilization and reduced OpEx [51,52], but real-time constraints, energy efficiency, and safe distribution of control remain open problems.

4.3. Cybersecurity and Network Monitoring

Agentic architectures (e.g., NetMoniAI) combine distributed micro-agents with centralized controllers for anomaly detection, DDoS classification, and layered responses [12,20,54]. These systems

 doi:10.20944/preprints202512.0592.v1

improve detection accuracy and interpretability compared to monolithic detectors, yet adaptive mitigation, adversarial robustness, and secure agent credentialing (A2A concerns) are highlighted as urgent gaps [12,54,55].

4.4. DevOps, Infrastructure-as-Code (IaC), and IT Operations (AIOps)

Multiple works apply agentic workflows to CI/CD pipelines, automated vulnerability detection and remediation, and predictive AIOps [14,18,22,56]. Here agents parse IaC templates, recommend or enact fixes, and reduce MTTR in live systems [18,22], the main trade-offs are latency and runtime cost versus gains in automation, along with governance/approval workflows to prevent unsafe automated changes [14,22,56].

4.5. E-commerce, Supply Chains and Business Process Automation

Agentic systems in commerce automate catalog management, customer interactions, demand forecasting, and end-to-end supply chain decisions [11,57]. Papers report substantial operational cost reduction and improved throughput when agents can act autonomously on structured business data [19,57], but they also flag dataset biases, limited adaptivity to rare events, and integration friction with legacy enterprise systems [19].

4.6. Digital Twins, Scientific Infrastructure and Industrial Control

Agentic LLMs are paired with digital twins to provide language-driven interfaces, select relevant sensors/tools, and orchestrate predictive maintenance and fleet management tasks [24,52]. Use cases include ship-fleet coordination and particle accelerator control [24,58], benefits include lowered cognitive burden and better situational awareness, while strict verification, safety constraints, and the need for digital-twin fidelity are practical barriers [24,52,58].

4.7. Education and Human-AI Collaborative Learning

Work on agentic educational agents frames them as co-learners or peer collaborators that supports student learning, adapt to learner profiles, and participate in collaborative tasks [32,33]. While promising personalization and engagement, concerns include dependency effects, ethical auditing, and designing interfaces that help students form accurate mental understanding of agent behavior [28,32,33]. Beyond education, agentic AI has also been explored in creative and design-oriented domains, where conversational and multimodal agents support early-stage architectural exploration and human–AI co-design [59].

4.8. Perception, Robotics and Multimodal Interaction

Papers such as Feature4X and embodied-agent studies show agentic systems combining vision, audio, and action to perform complex perception–reasoning–action cycles (e.g., scene reconstruction, robotic control loops) [3,34,60]. These multimodal agents broaden applicability to real-world embodied tasks but intensify compute, data, and safety requirements.

4.9. Cross-cutting / Governance and Socio-technical Applications

Several studies focus on governance, runtime instrumentation, and the human side of deploying agentic systems: ethics frameworks centered on interactional respect, MI9-style runtime governance, and organizational perceptions of ROI and responsibility [14,61,62]. These works underline that successful deployments hinge as much on governance and explainability as on pure technical capability of the agentic systems [14,61,62].

5. Methods and Techniques

This section summarizes the technical approaches used to build, coordinate, and evaluate agentic AI systems in the 51 papers. We have grouped techniques into following categories so we can see

 doi:10.20944/preprints202512.0592.v1

which ideas are recurring, which are experimental, and where engineering choices trade off cost, latency, or safety.

5.1. Orchestration Patterns

We observed that there are two orchestration families dominate: (a) LLM-centric prompt chaining and (b) separation-of-concerns architectures that split semantic reasoning from control logic [19,34]. Prompt chaining / LLM-driven pipelines are simple to prototype: an LLM decomposes a task into steps, calls tools, and iterations [19,20]. They work well for short-horizon, loosely structured tasks but are expensive such as they have many API calls in them. They also are fragile to prompt drift, and struggle with observability [19,20,63]. Several applied systems still rely on this pattern for rapid iteration [28]. Separation-of-concerns architectures (e.g., TB-CSPN and similar designs) place LLMs in the semantic layer while using formal coordination (Petri nets, FSMs, or token-driven protocols) for control and messaging [17,34]. This reduces LLM calls, improves verifiability, and makes orchestration more predictable under scale [17,34]. TB-CSPN explicitly reports large efficiency gains by offloading coordination to a formal control plane [19].

5.2. Multi-Agent Coordination and Protocols

Across the finalized literature we have observed that the coordination factor in a multi-agent systems is shifting from the improvised message passing to more structured and standardized communication protocols [7,20,34]. The protocols such as Agent-to-Agent (A2A) protocols and the frameworks such as the Model Context Protocol (MCP)-style frameworks define clear procedures for discovering other agents capabilities, for exchanging credentials and negotiating intentions. The mentioned factors play a key role in ensuring that the multi-agent systems remain both secure and scalable [13,55]. The research studies that implement the agent-to-agent (A2A) communication protocols emphasize that establishing the trust between agents requires more than just a simple message exchange. Checks and balancing factors such as the verified identities, proper authorization checks and safeguards against replayed or altered messages are consistently highlighted as the core requirements for maintaining a reliable coordination and preventing the manipulation within the system [55]. Across the finalized literature we also observed a hierarchical coordination method in which we have a flow that goes from supervisor agent to the consultant agent and from there it goes to the worker agent. The mentioned approach continues to work well for the tasks that require breaking down of the complex problems into smaller components such as repository management or infrastructure-as-code (IaC) remediation workflows. This hierarchal approach naturally limits the fault propagation and it establishes a more clear boundary of responsibility in the agentic systems [13,18,22].

5.3. Planning, Reasoning and Long-Horizon Control

In the finalized literature we have observed that the agentic AI systems use different types of planning strategies. These strategies are based on factors such as how long a task will take and how much risk is acceptable [3,16]. In some cases the agentic AI systems combine the classical planning with LLM-based reasoning where the large language models create the task-level plans which are then checked with symbolic rules or validations to prevent harmful actions [13]. In other system we had observed the agents use probabilistic or metacognitive checks which means that they measure uncertainty, select backup options or they ask humans for confirmation when unsure. This style of reasoning where the decisions are guided by probabilities and confidence estimates, is especially important in sensitive settings such as healthcare and other safety-critical domains where uncertainty must be handled carefully and the errors carry significant consequences [10,27,32]. In contrast the learning-based planning approaches such as reinforcement learning and imitation learning are used when the agents are able to interact with their environment and improve their decision-making through direct feedback over time. from their environment. These approaches require secure sandbox environments and accurate domain simulators to ensure safety [24,58,64].

 doi:10.20944/preprints202512.0592.v1

5.4. Model Stack: LLMs, SLMs, and Hybrid Deployments

A practical theme that we had observed in the literature is mixing of model sizes and specializations to balance cost and capability [21,25,32]. Small Language Models (SLMs) are recommended for repetitive, narrow tasks (low latency, cheaper inference). Papers show that a surprisingly large fraction of LLM-initiated operations can be handed off to SLMs after task clustering and fine-tuning [21]. LLMs as generalists/fallbacks: large models remain valuable for open-ended reasoning, out-of-distribution queries, and complex planning. Hybrid stacks route the common tasks to SLMs and escalate ambiguous or hard reasoning to LLMs [23,32].

5.5. Tooling and External Integrations (APIs, IaC, Edge Components)

Agentic agents typically require robust tool integration [2,18,29]. Tool adapters convert agent requests into concrete API calls (e.g., kubectl, Terraform, monitoring APIs). Effective adapters include schema validation, idempotency checks, and dry-run modes to reduce catastrophic changes [18,22]. Edge-aware orchestration: For latency-sensitive or privacy-sensitive tasks, authors recommend running SLMs or lightweight agents on edge devices while reserving heavy inference in the cloud [21,26,29]. Edge orchestration engines and containerized toolchains are common here in this usecase [29].

5.6. Formal Methods and Verifiability

We had observed in the finalized literature that to address the reliability and the safety needs several works combine the formal control methods with probabilistic reasoning [23,26]. We had also observed that the models such as the Petri nets, Colored Petri Nets (TB-CSPN) and FSMs were used to represent thecoordination workflows so they can be inspected, simulated and formally reasoned about [17,19,34]. The mentioned mechanisms and models help to connect tothe policies that humans can understand with the actual coordination behaviors which is carried out by agents. In systems like MI9, conformance engines and containment strategies actively monitor agents during their operation such as comparing their actions to predefined policy rules and stepping in when something starts to diverge [17]. Now instead of stopping the system outright they applied graduated interventions ranging from gentle corrections to stronger restrictions when needed. This approach strengthens oversight, its improves the transparency and it also helps to reduce the unexpected or unsafe behaviors from agents in agentic AI systems.

5.7. Observability, AgentOps and Runtime Governance

Across the finalized literature we had observed that the operational maturity in the agentic systems depends heavily on tools that can observe and interpret how agents behave in practice [19,23] such as for example the agentOps pipelines monitor the internal workflow of agents by recording things like call traces, memory usage and interactions with external tools. These can also track performance factors such as task success rates, reasoning depth and response time and these can automatically flag issues which arise during the execution [23,56]. However the observability cannot just stop at simply documenting inputs and outputs. For the effective monitoring of the agentic system we must also capture what is happening inside the agent which includes its reasoning steps and the semantic cues that shape its decisions [19,23]. For runtime governance we observed that the information this gathered from telemetry data and behavioral risk indicators data is fed into the runtime governance modules which look for signs of drift, misaligned goals or the formation of unintended subgoals. Systems such as MI9 show promise in this regard as it presents strong results in this area by combining semantic telemetry with causal trace analysis allowing this system to detect problematic behavior early and maintain a more reliable agent performance [17].

5.8. Security, Identity and Trust Mechanisms

Across the studies we had observed that the security factor appears time and time again as a core design priority rather than something which is added later in the development process [12,65]. Such as we had observed that many studies highlighted the importance of using strong credential

 doi:10.20944/preprints202512.0592.v1

management and secure communication channels for Agent-to-Agent (A2A) interactions. Without the verified identities and carefully scoped permissions the agent networks can become vulnerable to issues such as impersonation attacks or unauthorized privilege escalation [13,58]. Additionally the research also focuses on stress-testing these agentic AI systems through adversarial scenarios which includes prompt manipulation and data poisoning attacks. To address these types of security and trust risks, the reviewed work recommends the use of sandboxed execution environments, strict policy enforcement and rollback or fail-safe mechanisms that allow the systems to recover safely if something goes wrong [12,64].

5.9. Evaluation Metrics and Benchmarks

When we were going through the studies we had observed a limitation that was coming back again and again and it was lack of a consistent framework for evaluating agentic AI systems [3,4,26]. Another thing we had obersved was that the metrics that were used across the studies also varied widely from one project to another. Some of the works in the finalized studies measure performance in terms of task completion or success rates such as the execution metrics reported in EnvX [13] while other rely on familiar machine learning indicators such as precision, recall and F1 score etc particularly in detection-focused applications [26,34]. There were also some works that we had observed that concentrated on the operational performance using measures like mean time to recovery (MTTR), throughput, latency or cost per task to assess how well the system performs under real-world conditions [12,14,56]. In addition to this several works evaluated governance related outcomes such as how frequently the system is able to detect or prevent unsafe or unintended behaviors [14,17]. Despite all of these efforts we observed that the field of agentic AI lacks the shared benchmarks especially those which are designed to test multi-step reasoning, long-term coordination and failure safe decision-making. Developing such standards remains an important step for the advancement of reliable and a comparable evaluation of agentic AI.

6. Data Sources / Datasets / Evaluation Practice

We have reviewed a total of 51 finalized studies and across the 51 studies reviewed, we had observed that the LLM's play a central role in how the agentic AI systems are trained, how they are evaluated and how they are deployed. However unlike the traditional machine learning research where the standardized benchmark datasets are widely available the agentic AI systems often depend on domain-specific, proprietary or simulated data. This is largely because these systems are designed to operate in dynamic, real-world environments where the tasks require interaction rather than a static prediction against a query. As a result of this phenomenon the data used in these systems is frequently shaped by the environment settings itself regardless of the fact where it comes from such as it can come from the clinical sensors data or the network telemetry data or user workflows data or the artificially constructed simulations data. This section provides summary of the types of data used across the finalized studies, how these datasets are sourced in those said studies or how they are generated and the challenges that come in the evaluating and reproducing phase of the system performance.

6.1. Domain-Specific Operational Data

We observed across the studies that they are many agentic systems that operate inside the applied environments such as the healthcare facilities, network infrastructures, industrial plants, retail platforms and logistics systems [25,64]. In cases such as these we had observed that the datasets are typically native to the operational domain such for the healthcare domain as we have clinical monitoring data, the patient's sensor streams data, imaging reports, risk assessment logs and anonymized EMR data when allowed [10]. In the field of networking and telecommunication we have got the network telemetry data, the routing statistical data, event logs, congestion and spectrum usage traces data. For the field of cybersecurity we we have got anomaly detection logs data, the traffic signatures data, attack simulation outputs data and the data about intrusion patterns [12,20,25]. For the field of devOps or Infrastructure-as-Code we have got version histories data, data about the configuration

 doi:10.20944/preprints202512.0592.v1

repositories, deployment logs data, CI/CD job states data and the data about the observed system drift [18]. In the field of supply Chain and e-Commerce we got the data about transaction histories, inventory timelines data, vendor catalog data, and pricing/state transition records data [57,64]. These types of data is highly contextual and sensitive, which means access is often restricted and difficult to standardize across research groups.

6.2. Synthetic and Simulated Data

When real-world data is unavailable due to privacy, confidentiality, or safety constraints, studies frequently rely on simulation frameworks or synthetic datasets designed to replicate environmental conditions. Examples include network digital twins for routing, load balancing, and fault scenarios [19], simulated patient vital signal streams for controlled clinical testing, emulated cyber-attack networks to generate repeatable intrusion patterns [12,20,25], and synthetic task pipelines used to test orchestration and planning under controlled noise [13]. Synthetic data enables reproducibility and stress-testing, but may not fully capture real-world noise, irregularity, or adversarial behavior [12].

6.3. Public Benchmark Datasets

Public datasets are used less often because agentic tasks tend to be highly specialized. When they are used, these datasets mainly serve as baseline resources for evaluation. Examples include Feature4X and other multimodal perception datasets for vision-based reasoning, DAVIS datasets for scene understanding in embodied agent tasks [60], GitTaskBench for evaluating tool-use and coding agents [13], and standard machine learning datasets such as UCI or clinical and imaging archives are used when there are agent pipelines including classification components [10]. However, there are very few public benchmarks that evaluate multi-step decision-making, which makes it difficult to compare results across different studies.

6.4. Data Access, Governance and Privacy Constraints

A recurring theme across the papers is the difficulty of sharing or replicating datasets, due to patient privacy and clinical ethics constraints, proprietary logs from enterprise or industrial deployments, government and telecom regulation on network traces [12] and security sensitivities in cybersecurity and DevOps environments [12]. As a result, cross-institution reproducibility is limited, and evaluation tends to be local, within the context of each organization's environment.

6.5. Data Preparation, Integration and Preprocessing Approaches

Most agentic AI systems require structured data organization before agents can reason over it [10]. Common preparation strategies observed include schema alignment and metadata standardization, feature extraction pipelines for multimodal sensor or telemetry data [10,17], log parsing and normalization in DevOps and network operations, versioning of agent memory or replayable execution traces [17], and filtering and ethical review layers (especially in healthcare) [10]. These steps ensure the agent can interpret environment state, track temporal changes, and make informed decisions.

7. Limitations

Across the 51 papers, authors repeatedly surface a consistent set of limitations that slow the maturation of agentic AI from promising prototypes into reliable, widely adopted systems. We have categorized the studies's limitations and constraints into six categories such as technical, data and evaluation, governance and ethics, security and adversarial risk, human factors, and operational/economic limitations.

7.1. Technical Limitations

LLM-driven cost and latency: Many orchestration designs that depend heavily on repeated large language model (LLM) calls face issues like excessive token usage, high API expenses, and slow response times. These limitations make such systems unsuitable for real-time or high-throughput

 doi:10.20944/preprints202512.0592.v1

applications [12,19,51]. Studies that adopt a separation-of-concerns approach such as TB-CSPN highlight these issues as key reasons for shifting coordination responsibilities away from constant LLM calls [19]. Scalability of coordination: Formal coordination techniques and Petri-net-based controllers improve system predictability but encounter challenges when scaled to hundreds or thousands of agents. Ensuring that resource usage grows at a slower (sub-linear) rate as the number of agents increases remains an ongoing engineering challenge [19,34]. Long-horizon planning and brittleness: Agents designed to plan across many steps often fail when facing changing or unpredictable environments, such as digital twins, particle accelerators, or network control systems. In such cases, learning-based planners must be carefully sandboxed to prevent unsafe or unstable exploration behaviors [24,58].

7.2. Data, Benchmarking and Evaluation

Lack of standardized benchmarks: There are still very few widely accepted community benchmarks that evaluate multi-step reasoning, long-term planning, tool-use, or safe failure handling. Most studies rely on their own in-house or simulated datasets, which makes it difficult to compare results across different papers [2–4]. Proprietary and sensitive datasets: Fields such as healthcare and telecommunications often depend on sensitive data sources like system logs or electronic health records (EHRs). Because these cannot be openly shared, reproducibility is reduced and overall research progress in the community slows down [10,26,27]. Synthetic-to-real gaps: While simulators and synthetic datasets allow for safe and controlled testing, they often fail to capture real-world conditions such as noise, adversarial behavior, or subtle human interaction patterns that significantly influence how agents perform in practical scenarios [17].

7.3. Governance, Explainability and Ethics

One of the limitations that we observed was that even though frameworks like MI9 offer approaches for monitoring agent behavior and containing unwanted actions at runtime [17], building the governance systems which are both reliable and light weight remains difficult. This challenge becomes even more bigger when we work with black-box LLMs and external third-party tools, where visibility into internal reasoning is limited overall [14,17]. In addition to this the continuous monitoring can also add computational overhead and when the telemetry data is incomplete or even delayed it may introduce new points of failure rather than improving the stability of the system. Agentic AI systems also blurs the traditional lines of responsibility like when an AI agent takes an action it is not clear that where the responsibility of it lies, like does it lie with the developer, does it lie with the organization deploying the system, or does it lie with the the system itself. This raises some unresolved legal and contractual questions involving liability, intellectual property, and operational oversight [4,14,66]. Here a study says that this is a "moral crumple zone" where the accountability becomes diffused across multiple actors which makes it difficult to assign the responsibility when the errors occur [4]. The existing safety principles such as (HHH) in which we design the systems to be helpful, honest, and harmless offer a starting point but again this approach fails for agents which operate independently over a long period of time and for agents when they are interacting with users repeatedly. Recent studies show a trend leaning towards ethics in all aspects which focus on respect, relational awareness and the broader social impact of an agent's behavior. However making these ethical principles into clear and enforceable operational standards is still an ongoing and largely unresolved challenge.

7.4. Security and Adversarial Risks

Standards for securely managing agent identities, permissions, and message exchange are still in an early stage of development. Because of this, there are real risks, an agent could be impersonated, granted more access than intended, or even coordinate harmful actions with others. Without clear and widely adopted security protocols, it becomes difficult to maintain trust and control when multiple agents operate together in the same system. Studies examining Agent-to-Agent or A2A and related

 doi:10.20944/preprints202512.0592.v1

communication protocols highlight the importance of using signed identities and scoped authorizations to establish trust and ensure secure interactions among agents [13,55].

7.5. Human Factors and Usability

Users often build a wrong idea about how agentic systems actually work. They usually think of them as advanced search tools and fail to understand the confidence levels, reasoning steps, or how much control these systems really have [62]. Because of this, users can trust these systems too much or use them the wrong way. Current studies on showing the reasoning and confidence of such agents are still in early stages. In classrooms or work environments, using too much automation can make people too dependent on these systems or cause them to lose their own skills over time. Researchers suggest that users should sometimes work without the system ("unplugged" sessions) and that some decisions should still need human review or approval [4,32].

7.6. Operational and Economic Constraints

There are clear cost versus benefit trade-offs when using high-reasoning agents. In many business tasks, such advanced systems can be unnecessarily complex and expensive. Decision-makers need to carefully balance the cost of running and coordinating these models with the actual value they provide [14,57]. Several industry studies show that using a mix of smaller language models (SLMs) with large language model (LLM) backups is often a more cost-effective solution in real-world use [14,21,57].

7.7. Summary of Limitations

In short, agentic AI faces a many bottlenecks such as the technical stack (models, orchestration), the data and evaluation ecosystem (benchmarks and privacy), governance and legal frameworks, security posture, human-centric design, and practical deployment economics all need concurrent progress to remove limitations in agentic platforms and increase their adoption. Many papers in the set propose partial remedies such as SLM/LLM hybrids to lower cost, formal coordination layers (TB-CSPN) to stabilize orchestration, and runtime governance (MI9, AgentOps) to increase safety but no single study offers a complete solution [67]. Closing these gaps will require coordinated research, cross-industry data sharing agreements, standardized benchmarks, and regulatory/technical tooling for runtime oversight.

The severity graph in Figure 5 provides an overview of the fundamental challenges present in the Agentic AI domain. As depicted, the majority of these challenges are technical in nature, with LLM-driven cost and latency, limitations in coordination capabilities, and issues related to long-horizon planning and brittleness emerging as the primary concerns.

8. Future Work and Research Directions

The reviewed literature shows that agentic AI is advancing steadily, yet several open research directions remain critical for making these systems more dependable, scalable, explainable, and suitable for real-world deployment. The following future work areas reflect recurring suggestions and forward-looking proposals across the 51 studies.

8.1. Scalable, Efficient Agentic Architectures

There are many agentic AI systems that heavily depend upon large language models or LLMs to break up tasks into small steps, to choose different type of tools and to coordinate actions among agents [22,62]. As we move forward in technology, there is a need to explore more efficient hybrid designs which combine LLMs with smaller models which are more domain focused, components which do symbolic reasoning, and structured coordination layers [21]. The main goal is to lower the cost of inference, reduce memory usage and improve response speed all while still retaining the flexibility and adaptability which makes these systems effective [68–70]. Approaches such as model distillation, role specialization, compressing agent responsibilities, and caching stable or frequently used workflows appear to be promising directions for future research.

 doi:10.20944/preprints202512.0592.v1

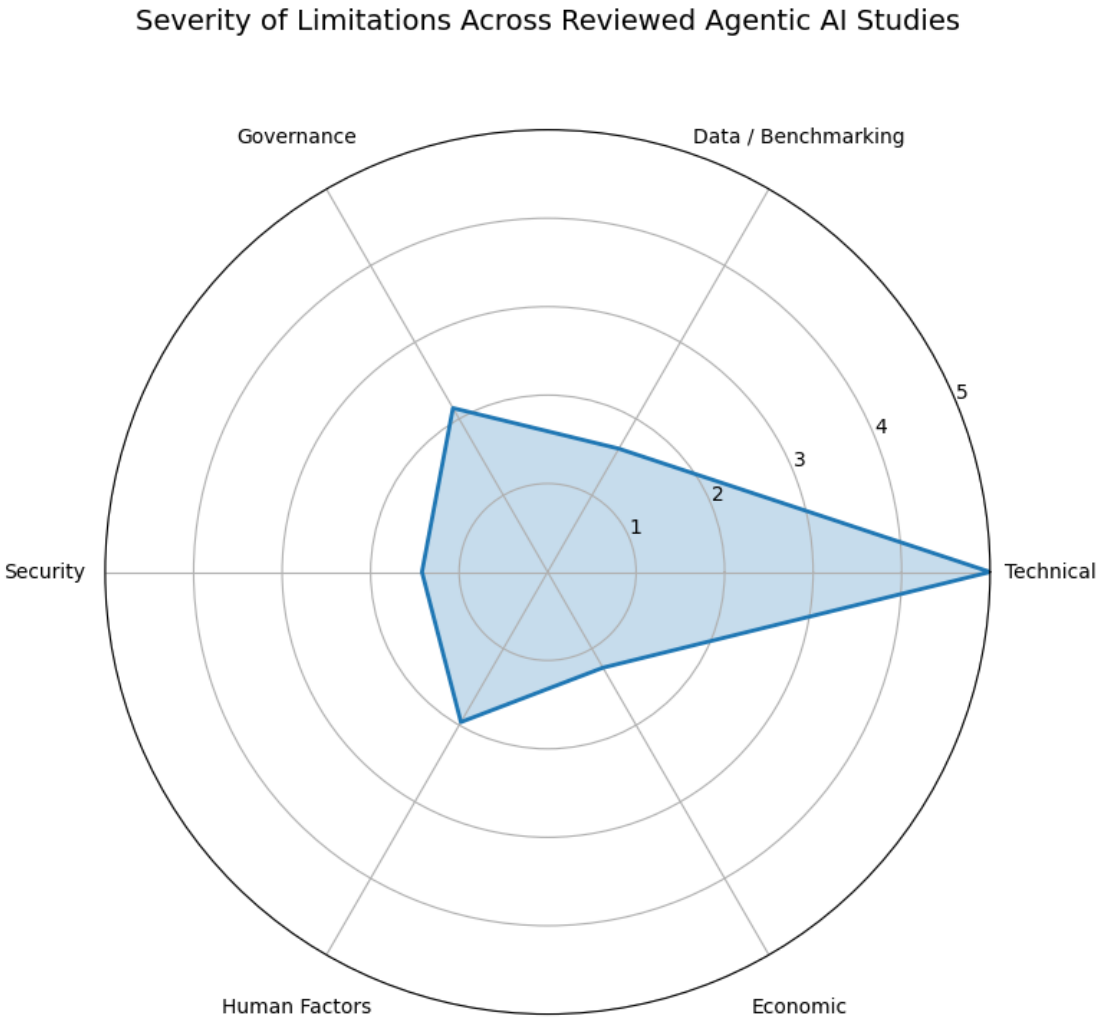

Figure 5. Core problems distribution in the Agentic AI.

8.2. Standardized Benchmarks and Evaluation Frameworks

While going through the studies we have observed a gap in which we saw that there is a lack of shared benchmarks for evaluating decisions which are made by following multiple steps, for evaluating collaborative workflows and for evaluating how the system handles failure safely [6],[68]. In order for the field of agentic AI to progress the future work should consider development of evaluation suites which are not tied to a specific domain and that can accurately reflect the autonomy of the agent, the planning ability of the agent and the reliability of the agent in using tools [13,62,71]. The mentioned benchmarks need more than just simple task completion metrics as they should also incorporate some of the governance measures as well. These governance measures can include measures such as how quickly misalignment is detected or how effectively the system can intervene to prevent or mitigate risky behavior.

8.3. Runtime Governance, Traceability and Oversight

We have observed that as the agent's autonomy increases the need for continuous oversight mechanisms also increases and it essentially becomes a must. The future systems will need to have detailed activity and reasoning logs which are also known as trace logs [17]. They will also need to have a continuous monitoring system for checking alignment [23] and ways or methods to manage unexpected or unwanted behaviour [17]. They should also have interfaces that let people review the decisions in a detailed or just overview form [34]. All of this highlights the need for research in field of

 doi:10.20944/preprints202512.0592.v1

"explainable autonomy frameworks" which should be capable of translating internal agent's reasoning into a human understandable form.

8.4. Security, Trust, and Multi-Agent Credential Infrastructure

In order to prevent issues such as impersonation, manipulation or unauthorized actions, the multi agentic systems need to have some measures such as verification, authorization, and communication protocols. While keeping these measures in mind, the future research can focus on implementing policy based message routing [29,55]. It can also focus on detecting adversarial or malicious behavior in real time [63]. All in all we can safely say that building frameworks in which we have Agent-to-Agent (A2A) trust, which are also scalable and resistant to tampering remains an important open challenge.

8.5. Human-Centered Interaction and Collaboration

Agentic AI is changing the way how humans interact with intelligent systems. This brings new risks or challenges of over-reliance, misinterpretation of confidence, or communication friction. To address these issues the future research can explore on designing interaction models which can help the users to form accurate mental models [53], introducing adjustable autonomy levels that allows for shared or negotiated control between humans and agents [4,16] and establishing safeguards that require human authorization for irreversible or high-impact decisions [4,14]. So, inshort more research is needed on cooperative task design, where humans and agents work together as adaptive partners rather than traditional tool–operator pairs.

8.6. Cross-Domain Generalization and Transfer

Right now, most agentic AI systems are still tightly tied to the specific domain they were designed for. They depend on specialized datasets, workflows, and assumptions about the environment, which makes it hard for them to operate anywhere else. A major direction for future research is to develop agents that can transfer what they've learned in one domain (for example, DevOps automation) and apply it in a different domain (such as IoT or healthcare) without needing to be retrained from scratch [26,64]. To Achieve this it requires advances in how agents form generalizable understandings of the world, how they link and reinterpret knowledge across different contexts, and how they adapt their reasoning when conditions change.

9. Conclusion

In this review study, we take a look at how agentic AI is evolving and where it will stand in the future. The big shift that we see is that the AI systems are no longer just recognizing patterns or giving answers when they are prompted, it is that these systems are beginning to plan, make decisions on their own without any human intervention, they are adapting to changes now, and they are starting to interact with people and other systems in more flexible ways. Agentic AI is already being tried out in domains and areas like healthcare, network control, cybersecurity, education, digital twins, and business automation. The results while they are promising less manual work, faster responses, more adaptive systems. They also reveal some real challenges as well. These challenges are scaling the systems in a reliable way, keeping the AI systems under oversight without slowing them down, evaluating them consistently and coordinating multiple agents securely.

In order to understand all of this we introduced a taxonomy which breaks the field of Agentic AI into key parts such as how agentic systems are built, how agents in these systems coordinate, what they can do, where they are used, what kinds of data and tools they depend on, how these agents and these systems they are in are governed and what are the limitations which is still hold these systems back. By organizing the research in this way we have shown how different approaches related and where the the field of agentic AI is going as a whole. One of the main things that we observed was that the technical progress alone is not enough, in order to move forward in this field of agentic ai, we need clear methods for human agent collaboration, stronger and more practical governance at runtime, shared evaluation standards that show what happens in real world conditions not just in controlled

 doi:10.20944/preprints202512.0592.v1

experiments. The future of the field of agentic AI depends on finding the right balance in giving the agents enough memory to be useful while also ensuring that they remain safe, accountable and aligned with human goals. Moving on ahead, the progress in field will require collaboration across multiple fields such as the fields of AI, security, ethics, human–computer interaction and systems design to make sure that the agentic AI becomes both, more effective and more trustworthy in practice.

References

1. Miehling, E.; Ramamurthy, K.N.; Varshney, K.R.; Riemer, M.; Bouneffouf, D.; Richards, J.T.; Dhurandhar, A.; Daly, E.M.; Hind, M.; Sattigeri, P.; et al. Agentic AI Needs a Systems Theory, 2025. arXiv:2503.00237 [cs], https://doi.org/10.48550/arXiv.2503.00237.
2. White, J. Building Living Software Systems with Generative & Agentic AI, 2024. arXiv:2408.01768 [cs], https://doi.org/10.48550/arXiv.2408.01768.
3. Garg, V. Designing the Mind: How Agentic Frameworks Are Shaping the Future of AI Behavior. *Journal of Computer Science and Technology Studies* **2025**, *7*, 182–193. https://doi.org/10.32996/jcsts.2025.7.5.24.
4. Mukherjee, A.; Chang, H.H. Agentic AI: Autonomy, Accountability, and the Algorithmic Society, 2025. arXiv:2502.00289 [cs], https://doi.org/10.48550/arXiv.2502.00289.
5. Wissuchek, C.; Zschech, P. Exploring Agentic Artificial Intelligence Systems: Towards a Typological Framework **2025**.
6. Porter, Z.; Calinescu, R.; Lim, E.; Hodge, V.; Ryan, P.; Burton, S.; Habli, I.; Lawton, T.; McDermid, J.; Molloy, J.; et al. INSYTE: A Classification Framework for Traditional to Agentic AI Systems. *ACM Trans. Auton. Adapt. Syst.* **2025**, *20*, 15:1–15:39. https://doi.org/10.1145/3760424.
7. de Curtò, J.; de Zarzà, I. LLM-Driven Social Influence for Cooperative Behavior in Multi-Agent Systems. *IEEE Access* **2025**, *13*, 44330–44342. https://doi.org/10.1109/ACCESS.2025.3548451.
8. Derouiche, H.; Brahmi, Z.; Mazeni, H. Agentic AI Frameworks: Architectures, Protocols, and Design Challenges, 2025. arXiv:2508.10146 [cs], https://doi.org/10.48550/arXiv.2508.10146.
9. Wang, H.; Gong, J.; Zhang, H.; Xu, J.; Wang, Z. AI Agentic Programming: A Survey of Techniques, Challenges, and Opportunities, 2025. arXiv:2508.11126 [cs], https://doi.org/10.48550/arXiv.2508.11126.
10. Shimgekar, S.R.; Vassef, S.; Goyal, A.; Kumar, N.; Saha, K. Agentic AI framework for End-to-End Medical Data Inference, 2025. arXiv:2507.18115 [cs], https://doi.org/10.48550/arXiv.2507.18115.
11. Pamisetty, A.; Farms, M. Application of agentic artificial intelligence in autonomous decision making across food supply chains **2024**. *1*.
12. Zambare, P.; Thanikella, V.N.; Kottur, N.P.; Akula, S.A.; Liu, Y. NetMoniAI: An Agentic AI Framework for Network Security & Monitoring, 2025. arXiv:2508.10052 [cs], https://doi.org/10.48550/arXiv.2508.10052.
13. Chen, L.; Peng, Z.; Yang, Y.; Wang, Y.; Tang, W.T.; Kobayashi, H.H.; Zhang, W. EnvX: Agentize Everything with Agentic AI, 2025. arXiv:2509.08088 [cs], https://doi.org/10.48550/arXiv.2509.08088.
14. Allam, H.; Dempere, J. Agentic AI for IT and Beyond: A Qualitative Analysis of Capabilities, Challenges, and Governance. *The Artificial Intelligence Business Review* **2025**, *1*. https://doi.org/10.64044/j63vmh26.
15. Alberts, L.; Keeling, G.; McCroskery, A. Should agentic conversational AI change how we think about ethics? Characterising an interactional ethics centred on respect, 2024. arXiv:2401.09082 [cs], https://doi.org/10.48550/arXiv.2401.09082.
16. Bedar, I.; Desroches, C. Agentic AI: Designing for Autonomy Without Losing Control.
17. Wang, C.L.; Singhal, T.; Kelkar, A.; Tuo, J. MI9 – Agent Intelligence Protocol: Runtime Governance for Agentic AI Systems, 2025. arXiv:2508.03858 [cs], https://doi.org/10.48550/arXiv.2508.03858.
18. Brodimas, D.; Birbas, A.; Kapolos, D.; Denazis, S. Intent-Based Infrastructure and Service Orchestration Using Agentic-AI. *IEEE Open Journal of the Communications Society* **2025**, *6*, 7150–7168. https://doi.org/10.1109/OJCOMS.2025.3600706.
19. Borghoff, U.M.; Bottoni, P.; Pareschi, R. Beyond Prompt Chaining: The TB-CSPN Architecture for Agentic AI. *Future Internet* **2025**, *17*, 363. Publisher: Multidisciplinary Digital Publishing Institute, https://doi.org/10.3390/fi17080363.
20. Qayyum, A.; Albaseer, A.; Qadir, J.; Al-Fuqaha, A.; Abdallah, M. LLM-Driven Multi-Agent Architectures for Intelligent Self-Organizing Networks. *IEEE Network* **2025**, pp. 1–10. https://doi.org/10.1109/MNET.2025.3605319.
21. Belcak, P.; Heinrich, G.; Diao, S.; Fu, Y.; Dong, X.; Muralidharan, S.; Lin, Y.C.; Molchanov, P. Small Language Models are the Future of Agentic AI, 2025. arXiv:2506.02153 [cs], https://doi.org/10.48550/arXiv.2506.02153.

 doi:10.20944/preprints202512.0592.v1

22. Toprani, D.; Madisetti, V.K. LLM Agentic Workflow for Automated Vulnerability Detection and Remediation in Infrastructure-as-Code. *IEEE Access* **2025**, *13*, 69175–69181. https://doi.org/10.1109/ACCESS.2025.3560911.
23. Moshkovich, D.; Zeltyn, S. Taming Uncertainty via Automation: Observing, Analyzing, and Optimizing Agentic AI Systems, 2025. arXiv:2507.11277 [cs], https://doi.org/10.48550/arXiv.2507.11277.
24. Timms, A.; Langbridge, A.; Antonopoulos, A.; Mygiakis, A.; Voulgari, E.; O'Donncha, F. Agentic AI for Digital Twin. *Proceedings of the AAAI Conference on Artificial Intelligence* **2025**, *39*, 29703–29705. https://doi.org/10.1609/aaai.v39i28.35373.
25. Chatzistefanidis, I.; Leone, A.; Nikaein, N. Maestro: LLM-Driven Collaborative Automation of Intent-Based 6G Networks. *IEEE Networking Letters* **2024**, *6*, 227–231. https://doi.org/10.1109/LNET.2024.3503292.
26. Karunanayake, N. Next-generation agentic AI for transforming healthcare. *Informatics and Health* **2025**, *2*, 73–83. https://doi.org/10.1016/j.infoh.2025.03.001.
27. Suura, S.R. Agentic AI Systems in Organ Health Management: Early Detection of Rejection in Transplant Patients. *Journal of Neonatal Surgery* **2025**.
28. Tian, J.; Wang, L.; Fard, P.; Junior, V.M.; Blacker, D.; Haas, J.S.; Patel, C.; Murphy, S.N.; Moura, L.M.V.R.; Estiri, H. An Agentic AI Workflow for Detecting Cognitive Concerns in Real-world Data, 2025. arXiv:2502.01789 [cs], https://doi.org/10.48550/arXiv.2502.01789.
29. Salama, A.; Nezami, Z.; Qazzaz, M.M.H.; Hafeez, M.; Zaidi, S.A.R. Edge Agentic AI Framework for Autonomous Network Optimisation in O-RAN, 2025. arXiv:2507.21696 [eess], https://doi.org/10.48550/arXiv.2507.21696.
30. Botti, V. Agentic AI and Multiagentic: Are We Reinventing the Wheel?, 2025. arXiv:2506.01463 [cs], https://doi.org/10.48550/arXiv.2506.01463.
31. Brohi, S.; Mastoi, Q.u.a.; Jhanjhi, N.Z.; Pillai, T.R. A Research Landscape of Agentic AI and Large Language Models: Applications, Challenges and Future Directions. *Algorithms* **2025**, *18*, 499. Publisher: Multidisciplinary Digital Publishing Institute, https://doi.org/10.3390/a18080499.
32. Paul, A.; Yu, C.L.; Susanto, E.A.; Lau, N.W.L.; Meadows, G.I. AgentPeerTalk: Empowering Students through Agentic-AI-Driven Discernment of Bullying and Joking in Peer Interactions in Schools, 2024. arXiv:2408.01459 [cs], https://doi.org/10.48550/arXiv.2408.01459.
33. Yan, L. From Passive Tool to Socio-cognitive Teammate: A Conceptual Framework for Agentic AI in Human-AI Collaborative Learning, 2025. arXiv:2508.14825 [cs], https://doi.org/10.48550/arXiv.2508.14825.
34. Borghoff, U.M.; Bottoni, P.; Pareschi, R. Human-artificial interaction in the age of agentic AI: a system-theoretical approach. *Frontiers in Human Dynamics* **2025**, *7*. Publisher: Frontiers, https://doi.org/10.3389/fhumd.2025.1579166.
35. Sapkota, R.; Roumeliotis, K.I.; Karkee, M. AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges. *Information Fusion* **2025**, *126*, 103599. arXiv:2505.10468 [cs], https://doi.org/10.1016/j.inffus.2025.103599.
36. Olujimi, P.A.; Owolawi, P.A.; Mogase, R.C.; Wyk, E.V. Agentic AI Frameworks in SMMEs: A Systematic Literature Review of Ecosystemic Interconnected Agents. *AI* **2025**, *6*, 123. Publisher: Multidisciplinary Digital Publishing Institute, https://doi.org/10.3390/ai6060123.
37. Acharya, D.B.; Kuppan, K.; Divya, B. Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access* **2025**, *13*, 18912–18936. https://doi.org/10.1109/ACCESS.2025.3532853.
38. Dietrich, N. Agentic AI in radiology: emerging potential and unresolved challenges. *British Journal of Radiology* **2025**, *98*, 1582–1584. https://doi.org/10.1093/bjr/tqaf173.
39. Zhang, R.; Liu, G.; Liu, Y.; Zhao, C.; Wang, J.; Xu, Y.; Niyato, D.; Kang, J.; Li, Y.; Mao, S.; et al. Toward Edge General Intelligence with Agentic AI and Agentification: Concepts, Technologies, and Future Directions, 2025. arXiv:2508.18725 [cs], https://doi.org/10.48550/arXiv.2508.18725.
40. Ogbu, D. Agentic AI in Computer Vision Domain -Recent Advances and Prospects. *International Journal of Research Publication and Reviews* **2023**, *Vol 4*, 5102–5120. https://doi.org/10.55248/gengpi.5.1124.3309.
41. Zhao, C.; Liu, G.; Zhang, R.; Liu, Y.; Wang, J.; Kang, J.; Niyato, D.; Li, Z.; Xuemin.; Shen.; et al. Edge General Intelligence Through World Models and Agentic AI: Fundamentals, Solutions, and Challenges, 2025. arXiv:2508.09561 [cs], https://doi.org/10.48550/arXiv.2508.09561.
42. Khalil, R.A.; Ahmad, K.; Ali, H. Redefining Elderly Care with Agentic AI: Challenges and Opportunities, 2025. arXiv:2507.14912 [cs], https://doi.org/10.48550/arXiv.2507.14912.

 doi:10.20944/preprints202512.0592.v1

43. Zhang, R.; Tang, S.; Liu, Y.; Niyato, D.; Xiong, Z.; Sun, S.; Mao, S.; Han, Z. Toward Agentic AI: Generative Information Retrieval Inspired Intelligent Communications and Networking, 2025. arXiv:2502.16866 [cs], https://doi.org/10.48550/arXiv.2502.16866.
44. Gridach, M.; Nanavati, J.; Abidine, K.Z.E.; Mendes, L.; Mack, C. Agentic AI for Scientific Discovery: A Survey of Progress, Challenges, and Future Directions, 2025. arXiv:2503.08979 [cs], https://doi.org/10.48550/arXiv.2503.08979.
45. Sapkota, R.; Roumeliotis, K.I.; Karkee, M. Vibe Coding vs. Agentic Coding: Fundamentals and Practical Implications of Agentic AI, 2025. arXiv:2505.19443 [cs], https://doi.org/10.48550/arXiv.2505.19443.
46. Bandi, A.; Kongari, B.; Naguru, R.; Pasnoor, S.; Vilipala, S.V. The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges, 2025.
47. Hosseini, S.; Seilani, H. The role of agentic AI in shaping a smart future: A systematic review. *Array* **2025**, *26*, 100399. https://doi.org/10.1016/j.array.2025.100399.
48. Jiang, F.; Pan, C.; Dong, L.; Wang, K.; Dobre, O.A.; Debbah, M. From Large AI Models to Agentic AI: A Tutorial on Future Intelligent Communications, 2025. arXiv:2505.22311 [cs], https://doi.org/10.48550/arXiv.2505.22311.
49. Ren, Y.; Liu, Y.; Ji, T.; Xu, X. AI Agents and Agentic AI-Navigating a Plethora of Concepts for Future Manufacturing, 2025. arXiv:2507.01376 [cs], https://doi.org/10.48550/arXiv.2507.01376.
50. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. Publisher: British Medical Journal Publishing Group Section: Research Methods & Reporting, https://doi.org/10.1136/bmj.n71.
51. Dev, K.; Khowaja, S.A.; Singh, K.; Zeydan, E.; Debbah, M. Advanced Architectures Integrated with Agentic AI for Next-Generation Wireless Networks, 2025. arXiv:2502.01089 [cs], https://doi.org/10.48550/arXiv.2502.01089.
52. Xiao, Y.; Shi, G.; Zhang, P. Towards Agentic AI Networking in 6G: A Generative Foundation Model-as-Agent Approach, 2025. arXiv:2503.15764 [cs], https://doi.org/10.48550/arXiv.2503.15764.
53. Sheelam, G.K.; Komaragiri, V.B. Self-Adaptive Wireless Communication: Leveraging ML And Agentic AI In Smart Telecommunication Networks. *Metallurgical and Materials Engineering* **2025**, pp. 1381–1401. https://doi.org/10.63278/mme.vi.1716.
54. Gosmar, D.; Dahl, D.A. Hallucination Mitigation using Agentic AI Natural Language-Based Frameworks, 2025. arXiv:2501.13946 [cs], https://doi.org/10.48550/arXiv.2501.13946.
55. Habler, I.; Huang, K.; Narajala, V.S.; Kulkarni, P. Building A Secure Agentic AI Application Leveraging A2A Protocol, 2025. arXiv:2504.16902 [cs], https://doi.org/10.48550/arXiv.2504.16902.
56. Sivakumar, S. Agentic AI in Predictive AIOps: Enhancing IT Autonomy and Performance. *International Journal of Scientific Research and Management (IJSRM)* **2024**, *12*, 1631–1638. https://doi.org/10.18535/ijsrm/v12i11.ec01.
57. Alecsoiu, O.R.; Faruqui, N.; Panagoret, A.A.; Ionuț, C.A.; Panagoret, D.M.; Nitu, R.V.; Mutu, M.A. EcoptiAI: E-Commerce Process Optimization and Operational Cost Minimization Through Task Automation Using Agentic AI. *IEEE Access* **2025**, *13*, 70254–70268. https://doi.org/10.1109/ACCESS.2025.3560549.
58. Sulc, A.; Hellert, T.; Kammering, R.; Hoschouer, H.; John, J.S. Towards Agentic AI on Particle Accelerators, 2025. arXiv:2409.06336 [physics], https://doi.org/10.48550/arXiv.2409.06336.
59. Cheung, L.H.; Wang, L.; Lei, D. Conversational, agentic AI-enhanced architectural design process: three approaches to multimodal AI-enhanced early-stage performative design exploration. *Architectural Intelligence* **2025**, *4*, 10. https://doi.org/10.1007/s44223-025-00092-5.
60. Zhou, S.; Ren, H.; Weng, Y.; Zhang, S.; Wang, Z.; Xu, D.; Fan, Z.; You, S.; Wang, Z.; Guibas, L.; et al. Feature4X: Bridging Any Monocular Video to 4D Agentic AI with Versatile Gaussian Feature Fields, 2025. arXiv:2503.20776 [cs], https://doi.org/10.48550/arXiv.2503.20776.
61. Ackerman, L. Perceptions of Agentic AI in Organizations: Implications for Responsible AI and ROI, 2025. arXiv:2504.11564 [cs], https://doi.org/10.48550/arXiv.2504.11564.
62. Brachman, M.; Kunde, S.; Miller, S.; Fucs, A.; Dempsey, S.; Jabbour, J.; Geyer, W. Building Appropriate Mental Models: What Users Know and Want to Know about an Agentic AI Chatbot. In Proceedings of the Proceedings of the 30th International Conference on Intelligent User Interfaces, New York, NY, USA, 2025; IUI '25, pp. 247–264. https://doi.org/10.1145/3708359.3712071.

 doi:10.20944/preprints202512.0592.v1

63. Atta, H.; Baig, M.Z.; Mehmood, Y.; Shahzad, N.; Huang, K.; Haq, M.A.U.; Awais, M.; Ahmed, K. QSAF: A Novel Mitigation Framework for Cognitive Degradation in Agentic AI, 2025. arXiv:2507.15330 [cs], https://doi.org/10.48550/arXiv.2507.15330.
64. Mishra, L.N.; Senapati, B. Retail Resilience Engine: An Agentic AI Framework for Building Reliable Retail Systems With Test-Driven Development Approach. *IEEE Access* **2025**, *13*, 50226–50243. https://doi.org/10.1109/ACCESS.2025.3552592.
65. Huang, K.; Narajala, V.S.; Yeoh, J.; Ross, J.; Raskar, R.; Harkati, Y.; Huang, J.; Habler, I.; Hughes, C. A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control, 2025. arXiv:2505.19301 [cs], https://doi.org/10.48550/arXiv.2505.19301.
66. Raju, N.V.D.S.S.V.P.; Faruqui, N.; Patel, N.; Alecsoiu, O.R.; Thatoi, P.; Alyami, S.A.; Azad, A. LegalMind: Agentic AI-Driven Process Optimization and Cost Reduction in Legal Services Using DeepSeek. *IEEE Access* **2025**, *13*, 126981–126999. https://doi.org/10.1109/ACCESS.2025.3586781.
67. Floridi, L.; Buttaboni, C.; Hine, E.; Morley, J.; Novelli, C.; Schroder, T. Agentic AI Optimisation (AAIO): what it is, how it works, why it matters, and how to deal with it, 2025. arXiv:2504.12482 [cs], https://doi.org/10.48550/arXiv.2504.12482.
68. Khamis, A. Agentic AI Systems: Architecture and Evaluation Using a Frictionless Parking Scenario. *IEEE Access* **2025**, *13*, 126052–126069. https://doi.org/10.1109/ACCESS.2025.3590264.
69. Pang, C. Toward Data Systems That Are Business Semantic Centric and AI Agents Assisted. *IEEE Access* **2025**, *13*, 113752–113762. https://doi.org/10.1109/ACCESS.2025.3583260.
70. Koh, N.T.; Sharma, A.; Xiao, J.; Siong Chin, C.; Jun Xing, C.; Lok Woo, W. Optimized Sequential Agentic AI-Guided Trainable Hybrid Activation Functions for Solar Irradiance Forecasting. *IEEE Access* **2025**, *13*, 149976–149990. https://doi.org/10.1109/ACCESS.2025.3602978.
71. Bentley, P.J.; Lim, S.L.; Ishikawa, F. Situating AI Agents in their World: Aspective Agentic AI for Dynamic Partially Observable Information Systems, 2025. arXiv:2509.03380 [cs], https://doi.org/10.48550/arXiv.2509.03380.