Agentic AI as the Orchestrator of Mobile Ecosystems: A Review of the Trade-off Between Performance and Drawbacks

Ayat Aljarrah[1], Mustafa Ababneh[2]
Assist. Prof at Department of Computer Science-Faculty of Information Technology, Jadara University, Irbid, Jordan[1]
Assist. Prof at Department of Robotics and Artificial Intelligence-Faculty of Information Technology,
Jadara University, Irbid, Jordan[2]

***Abstract*—This system review explores the transformational role of agentic artificial intelligence (AI) as an orchestrator in mobile ecosystems. Agentic AI systems proactively plan, execute, and adapt across applications, devices, and services, unlike traditional and generative AI. These systems offer autonomous, context-aware coordination by integrating reasoning engines, tool orchestration, memory, retrieval-augmented generation (RAG), and safety layers. The review examines architectural requirements for mobile deployment, including on-device processing, resource-aware execution, and cross-platform synchronization. It stresses implementation targets and achievements through 2025, automation levels across key capabilities, and the impact of agentic orchestration on mobile ecosystem challenges. The findings highlight agentic AI's potential to optimize performance, privacy, and user experience simultaneously. Future directions include edge-native architectures, human-in-the-loop frameworks, and multi-agent interoperability standards. This study provides a comprehensive roadmap for advancing agentic AI as a foundational layer in next-generation mobile computing.**

***Keywords*—Agentic AI; orchestrator; mobile ecosystems; on-device**

I. Introduction

Due to the swift development of mobile ecosystems, users now have to navigate dozens of apps, numerous services, and constant data streams across several devices in the surroundings of unprecedented complexity. Manual coordination between various apps and services is necessary for traditional mobile systems, which mostly function through explicit user commands [1]. The potential for seamless digital encounters is constrained by this paradigm, which also places a heavy cognitive burden on consumers [2].

Agentic artificial intelligence (AI), as a class of artificial intelligence systems, can act, decide, and adapt to new conditions on its own without continual human interaction, providing a revolutionary solution to this problem [3]. Conventional AI is mainly used as a tool for passive activities; in contrast, agentic AI systems work proactively, creating plans, establishing objectives, and carrying out multi-step procedures across application boundaries [4], [5]. The interaction between users and their digital environments can be radically redefined when these systems are used as an orchestrating intelligence inside mobile ecosystems. They can handle complex workflows, anticipate user demands, and coordinate resources.

The key objective of this review is to critically examine the architectural context, operational mechanisms, and implementation challenges of agentic AI as an orchestrator of mobile ecosystems. The aim is to identify the essential elements, mobile-specific limitations, and automation potential that allow agentic AI to evolve from reactive tools to proactive, self-governing systems. The assessment also identifies existing constraints, provides a strategic implementation roadmap, and suggests future research opportunities to move agentic AI closer to scalable, secure, and context-aware mobile orchestration.

The contributions of this review include the summarization of recent developments in agentic AI to clarify its function as a proactive orchestrator in mobile ecosystems. This review uniquely reframes agentic AI not as an incremental upgrade to generative AI, but as a systemic redefinition of orchestration in mobile ecosystems, where autonomy, memory, and tool orchestration converge to resolve long-standing fragmentation and cognitive overload. Unlike prior reviews that treat these components in isolation, this study integrates them into a continuous perception–reasoning–decision–execution–learning cycle, offering the first holistic blueprint of agentic AI as a mobile orchestrator. The automation analysis is distinctive in exposing uneven maturity, where orchestration is nearly autonomous, but governance remains human-dependent, thus guiding targeted innovation priorities. The agenda is distinctive in its systemic scope, proposing standards and architectures that extend agentic AI from isolated deployments to ecosystem-wide rationality and resilience.

To guide the reader through the structure and scope of this review, the study is organized as follows: Section II outlines the conceptual evolution from generative AI to agentic AI, highlighting the shift from reactive to proactive intelligence. Section III details the architectural components of agentic AI systems, including reasoning engines, tool orchestration, memory, RAG, and safety layers. Section IV discusses mobile-specific architectural considerations such as on-device processing, resource-aware execution, and cross-platform synchronization. Section V presents an implementation roadmap and timeline for agentic AI deployment, while Section VI analyzes automation levels across key capabilities.

Section VII evaluates the impact of agentic orchestration on mobile ecosystem challenges. Section VIII proposes future research directions, including interoperability standards, human-in-the-loop frameworks, edge-native architectures, and security models. This structure ensures a coherent and comprehensive exploration of agentic AI as a transformative orchestrator in mobile computing. Finally, Section IX concludes the study.

II. Evolution of Generative AI to Agentic AI

The advance from conventional AI to agentic AI epitomizes a major shift in how systems interact with the real world by moving from passive devices to active, objective-oriented partners. The concept of agentic AI signifies a significant departure from passive AI tools and the development of autonomous, goal-driven systems that can predict, reason, and act in digital contexts [6], [7]. The transformational potential of Agentic AI as the orchestrating intelligence in complex mobile ecosystems is examined in this system assessment. Agentic AI allows for dynamic coordination across apps, services, and devices, transforming mobile platforms from reactive tools to proactive, contextual partners [8]. This review illustrates this by analyzing current architectures, applications, and difficulties (see Fig. 1).

The interface between humans and AI was revolutionized by generative AI, especially by models such as GPT-4 and DALL-E [9]. These systems are creators as well as classifiers. They can converse in natural language, compose essays, build code, and produce graphics from text descriptions. Hence, the "language barrier" between humans and machines was removed, enabling the public to access AI. Generative AI is essentially reactive, even with this advancement [10]. It utilizes its training to generate a response after waiting for a user command. It struggles with complex, multi-step tasks that need planning and sequential tool use, and it does not have permanent memory across chats; with a new chat starts from the beginning [11]. A typical generative AI does not have the capability to automatically conduct research, for example, research on the best laptops for students, find the top three deals online, and summarize them in a table.

The latest frontier is agentic AI, which builds active, goal-seeking systems by fusing new capabilities with the generating potential of Large Language Model (LLM) [12], [13]. It is considered to have a set of hands (tools/APIs), a notebook (permanent memory), and a brain (for planning) compared to a generative AI model. High-level objectives are given to an agentic AI system, such as "create a comprehensive market analysis report on electric vehicles." The agent then: 1) formulates a plan, which breaks the target down into steps: search for recent electric vehicle sales data, find updates on key manufacturers, analyze stock performance, and compile results. 2) Executes actions by using its integrated tools, such as a web browser, a code interpreter, and a document editor, which autonomously perform these steps. 3) Adapts and persists stores data in a permanent memory, permitting it to build on earlier results. It can reorganize and try a different strategy if one search does not work. This turns the AI from a potent chatbot into a self-sufficient digital assistant that can oversee intricate processes from beginning to end.

Fig. 1 displays evolution of AI capabilities. It details the comparative literature synthesis of traditional AI, generative AI, and agentic AI capabilities using published benchmarks and conceptual frameworks [14], [15], [16], [17]. The figure visually maps the transition from reactive to proactive intelligence, highlighting dimensions such as memory, tool use, and complexity handling. It is justified as a conceptual framework to ground the reader in the paradigm shift. Traditional AI scores low because it is highly reactive (only operates on input data), handles low complexity (single tasks), has no persistent context, and cannot use external tools. Hence, its value is in its precision, not in its adaptability. Generative AI shows a significant jump, specifically in handling more complex tasks such as content creation. However, it remains largely reactive, as indicated by its mid-level position. Its capabilities for memory and tool use are emerging and not inherent. Agentic AI occupies the high end of all four capabilities, indicating its transformative characteristics. Its high position highlights its proactive characteristics, and its far-right position confirms its ability to manage high complexity through integrated planning, memory, and tool use. Hence, this permits it to function as an autonomous agent rather than a reactive tool.

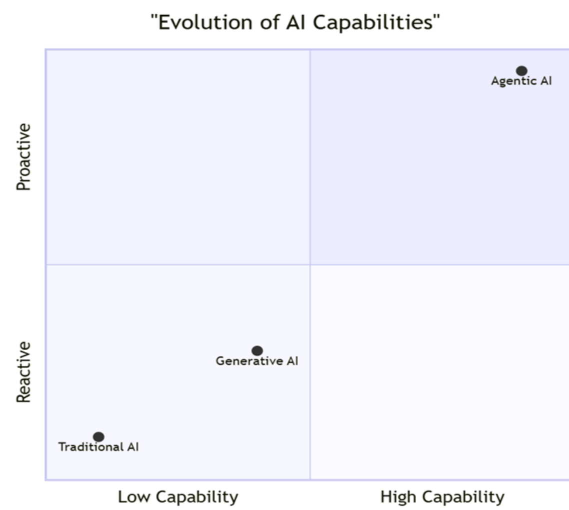

Fig. 1. Evolution of AI capabilities. Data sources: [14], [15], [16], [17].

III. Agentic AI Systems

Agentic AI systems are highly developed AI design created to behave independently and replicate human decision-making to accomplish complex tasks [7]. They conduct action without continual human supervision, observe their environment, and solve issues [3], [18]. A robust agentic AI system is built on various crux components for operational loop to function, as illustrated in Fig. 2.

A. *Reasoning Engine (LLM Core)*

The reasoning engine, which is usually driven by an LLM, functions as the agentic AI system's "brain" or central processing unit. Its main duties include deciphering user input, decomposing complex issues into controllable phases, and creating a rational plan of action by selecting the appropriate mechanisms and channels to employ and structure [Fig. 2(A)] [19]. According to Hughes et al. [3], this component oversees higher-order cognitive processes like comprehending context, coming up with logical methods, and modifying the plan considering new intelligence or data. However, this component's inherent propensity for confabulation or "delusion", in which the model generates plausible but erroneous or created thought processes, which can be a significant weakness that could lead the system in the incorrect path without any inherent responsiveness of the error [20],[21],[22].

B. Tools and Functions (Activity Capabilities)

The system's "hands and senses" are the tools and functions component that enable the system to communicate with and influence the outside world via its underlying language model [Fig. 2(A)] [3]. The reasoning engine can use these pre-defined features, like database queries, API calls, code execution, or web search functions, to collect data in real time, carry out computations, or carry out operations [7]. This helps to close the gap between intangible reasoning and concrete outcomes. However, the failure of a single external tool can upset the sequential plan since the agent needs to accurately match its abstract plan to a specific function call with the appropriate inputs [21]. This makes tool selection and orchestration complexity a major problem.

C. State and Memory (Context Management)

The state and memory module serves as the narrative record for the system, preserving a consistent context throughout the exchange [Fig. 2(A)] [23]. The present aim, chat history, earlier completed steps, and other pertinent user-specific data are all tracked, preventing the agent from repeating itself or losing track of lengthy, multi-turn activities [24]. This context window is crucial for interactions to be coherent and customized for each individual. However, early context may be "forgotten" or lost after prolonged contact, and it may be challenging to extract the most important details from a vast memory storage [22], thus designers can focus on how to upgrade this limitation. This component's drawback is its practical limitation in terms of context window size and memory degradation [20].

D. Knowledge and RAG (External Data)

The agent receives a specific external knowledge root in addition to its pre-trained weights through the knowledge and RAG (Retrieval-Augmented Generation) system [Fig. 2(A)] [23]. In response to a query, this part actively seeks for the most current and pertinent information from specified sources, including product databases or corporate records, and creates the LLM's response on this information [25]. This is crucial in specialized domains to ensure factual accuracy and reduce delusions. Its main drawback is that it relies on the accuracy of the retrieval and the quality of the knowledge base; however, if the retrieval system cannot find the correct information or the source data is outdated, the agent will generate responses based on incomplete or erroneous information [21],[22].

E. Guards and Evaluation (Safety Layers)

The crucial safety and supervision layer is made up of the "guards and evaluation" component, which filters inputs and outputs to make sure the system stays within predetermined bounds [Fig. 2(A)] [24]. Architectural decomposition based on system design literature [3],[7],[18],[19],[23]. The figure synthesizes multiple sources into a unified operational loop (perception → reasoning → decision → execution → learning). Each component (reasoning engine, tools, memory, RAG, safety layers) is mapped to its methodological role. Before the final output is sent to the user, it checks for malicious, unsafe, biased, or irrelevant material, verifies user requests for safety and policy compliance, and it observe the tools the agent attempts to employ [19],[23]. This is the primary defense mechanism for the responsible application of AI. However, the primary concern with this component is achieving comprehensive coverage without excessive restriction, as overly strict safeguards may render the agent ineffective in legitimate edge situations, while subtly detrimental content or creative "getaway" ideas may occasionally evade detection [20],[21],[22].

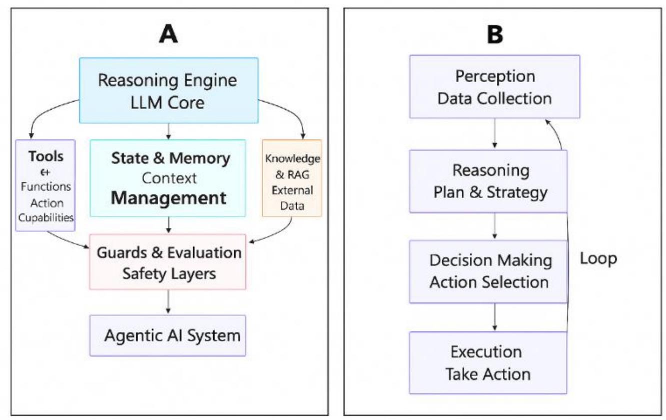

Fig. 2. Agentic AI system: (A) Core components and (B) Operational process from initial perception to continuous learning.

F. Perception (Data Collection)

Fig. 2(B) presents the operational process from initial perception to continuous learning. The perception phase [Fig. 2(B)] is the foundational sensory layer of the agentic AI system, where it collects real-time, multi-modal data from diverse sources to understand the present and flow situation [26]. However, the challenge is ensuring data quality and integration integrity, as garbage-in-garbage-out scenarios (e.g., noisy, biased, or incomplete data) can corrupt the decision-making pipeline [27].

G. Reasoning (Plan and Strategy)

The reasoning phase of a system uses a LLM to process data, interpret semantic meaning, understand context, and create a plan to accomplish a target objective (Fig. 2B). This process transforms abstract information into actionable intelligence, requiring logical inference, ultimate reasoning, and multi-step problem-solving [7],[25]. However, one significant flaw is the possibility of logical fallacies and delusions creating error limitations or behaviors [28].

H. Decision-Making (Action Selection)

Decision-making phase involves a system transitioning from planning to commitment, assessing actions based on efficiency, resource cost, predicted success rate, and alignment with the target objective [Fig. 2(B)]. There is a trade-off between trying out new strategies and taking advantage of ones that already exist, necessitating a delicate balance between careful optimization and bold aim [25],[26]. However, the challenge is the framing problem and sub-optimal lock-in, where the presentation of options can have a disproportionate impact on the agent's choice [28].

I. Execution (Take Action)

The execution phase involves the agent's abstract decisions being implemented in tangible interactions with the outer world, like calling third-party APIs or executing database operations [Fig. 2(B)]. This phase is essential as the system's internal logic meets the unpredictable reality of outer systems [8],[23]. However, the agent's weakness due to fragility and lack of real-world affordance reasoning can lead to unexpected errors, API timeouts, or permission denials, potentially halting the process [27],[28].

J. Learning and Adaptation (Refine Strategy)

In the Learning and Adaptation phase, the system refines its strategies and performance by collecting feedback on implemented actions, analyzing the difference between expected and actual results, and applying techniques like reinforcement learning [Fig. 2(B)]. This meta-cognitive ability transforms a static automated script into a dynamically enriching intelligent agent [7],[25]. However, challenges include credit assignment problems and tragic forgetting, where learning new information can overwrite or corrupt earlier acquired knowledge [20],[22],[29].

The constant, dynamic cycle between these phases is a fundamental aspect of agentic AI. The system creates a fluid and responsive kind of intelligence by continuously observing the new state of the world because of its prior activities, considering this new context, and making decisions in an iterative loop until the overall aim is accomplished.

IV. Mobile-Specific Architectural Considerations

Agentic AI orchestration in mobile contexts has distinct architectural needs that are largely different from those of cloud-based deployments. These needs compel the use of specific techniques to manage the limited and dynamic nature of mobile devices. On-device processing is the determinant, where sophisticated AI agents can now run directly on devices thanks to breakthrough Arm-based architectures [25],[30]. This allows them to securely protect user data and function dependably without constant cloud access [31],[32]. However, the computational and memory limitations, which even the most potent mobile processors and NPUs have in comparison to cloud servers, pose a serious problem, though developers have to make challenging trade-offs between model complexity, speed, and accuracy (Fig. 3) [13]. Moreover, when prolonged heavy computation produces heat, the device must severely limit performance to prevent hardware damage, which could block the agent's vital activities [8]. This is known as thermal throttling, and it becomes a critical bottleneck.

Resource-aware execution is another important factor to consider. This feature gives the system the ability to intelligently decide when to carry out tasks locally versus offloading them to highly potent distant servers, allowing it to dynamically manage battery consumption, processing load, and network usage (Fig. 3) [2]. Analytical framework of Fig. 3 combining mobile hardware constraints (Arm-based NPUs, thermal throttling) with orchestration strategies (on-device vs. offloading, cross-platform synchronization) [2], [13], [30], [31], [32], [33], [34], [35]. The figure illustrates trade-offs between computation, energy, latency, and synchronization. It is justified as a design-oriented methodology to show how agentic AI adapts to mobile-specific bottlenecks. Accurate resource prediction is the main challenge since the system has to predict battery drain and compute load for different tasks in a highly changeable environment where user behavior and device state are always changing [30]. However, it is challenging to create an efficient offloading approach that maintains a flawless user experience while weighing latency, data usage, and energy consumption against the quality of the output (Fig. 3) [33]. This requires a sophisticated cost-benefit analysis. Lastly, despite platform-specific constraints and fundamental variations in APIs and security sandboxes, cross-platform coordination guarantees that the agentic system can function uniformly across the fragmented landscape of iOS, Android, and progressive web applications (Fig. 3) [30],[31]. However, the difficulty of preserving feature parity, which demands substantial, platform-specific development and optimization, often doubling the technical effort, makes it difficult to give the same features and performance across many operating systems [6],[34]. Also, it is quite challenging to maintain consistent state synchronization across platforms, because the agent needs to have strong and conflict-free data merging protocols to maintain a coherent memory and task state when users switch between their phones, tablets, and online apps [35].

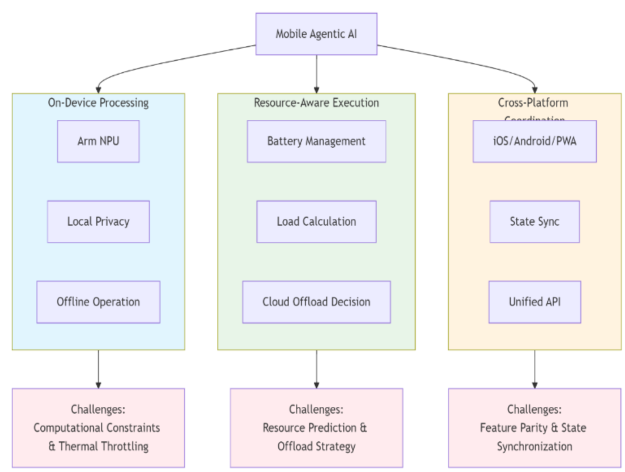

Fig. 3. Mobile-specific architectural considerations and challenges.

V. Agentic AI Implementation Roadmap and Timeline

The survey of literature made in this review showed the implementation success on mobile devices in the second quota of 2025 based on the Agentic AI research [2], [10], [13], [30], [31], [32]. Fig. 4 presents an agentic AI implementation roadmap and timeline, which structured progression from basic AI monitoring systems to fully autonomous and optimized operations. This figure was obtained from roadmap construction from literature survey of agentic AI deployments (AutoGen, LangChain, OSS/BSS frameworks) and maturity models [2],[10],[13],[30],[31], [32],[36]. Organizations usually use reactive systems and basic monitoring, which are fundamental, but have limited adaptability, starting with the existing situation in Q2 of 2025. Choosing suitable frameworks and tools, such as AutoGen or LangChain, which provide the foundation for agentic capabilities, is the next stage. Pilot projects that implement context-aware execution and self-healing pipelines come next, signaling a move toward more resilient and dynamic systems. A maturity curve that balances technological complexity and operational preparedness is suggested by each step, which shows growing implementation success.

Multi-agent integration becomes critical as the roadmap develops, allowing for advanced system orchestration and cooperative agents [25],[36]. This prepares the way for the last stage, which consists of continuous optimization and full production, where efficiency and innovation are driven by autonomous operations and predictive capacities. The strategic significance of each milestone is highlighted by the visual emphasis on implementation success percentages, which helps stakeholders manage expectations and allocate resources.

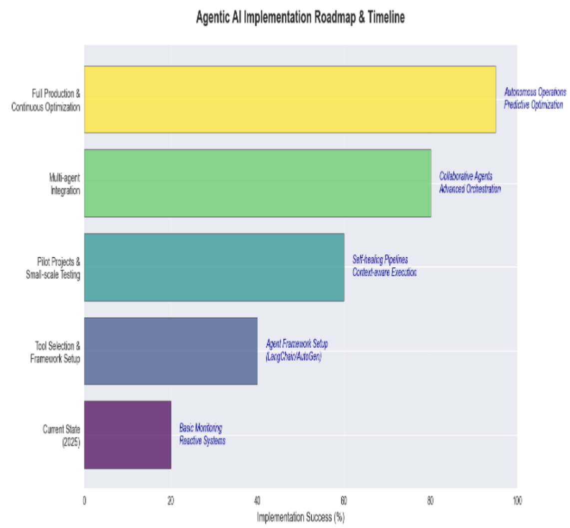

Fig. 4. Agentic AI implementation roadmap and timeline.

VI. Automation Level Across Key Capabilities

The survey of literature done in this review revealed the automation level across key capabilities in the second quota of 2025 based on the agentic AI research [1],[7],[12],[15],[22]. Fig. 5 displays the automation level across key capabilities of agentic AI, which offers a compelling portrait of how automation is distributed across six critical domains in AI-driven systems. Capabilities such as cloud-native orchestration, intelligent code generation, and pipeline protection demonstrate exceptionally high degrees of automation (over 90%), indicating scalable, mature systems that need little human interaction. Strong tooling and defined processes make these domains perfect candidates for complete automation of mobile ecosystems, particularly on the device. The capabilities like human-in-the-loop systems and multi-agent collaboration showed a more hybrid approach, with a considerable amount of human assistance (15–25%), indicating that even with automation progression, contextual judgment and supervision are still crucial.

With only 61% automation and a significant 30% human assistance, security and governance are the least automated domains (Fig. 5). This emphasizes the reason why delicate and complex risk management, ethical supervision, and regulatory compliance are areas where human judgment is still crucial. Overall, the findings showed where innovation and investment are required to advance toward more autonomous, robust AI ecosystems, in addition to highlighting current strengths.

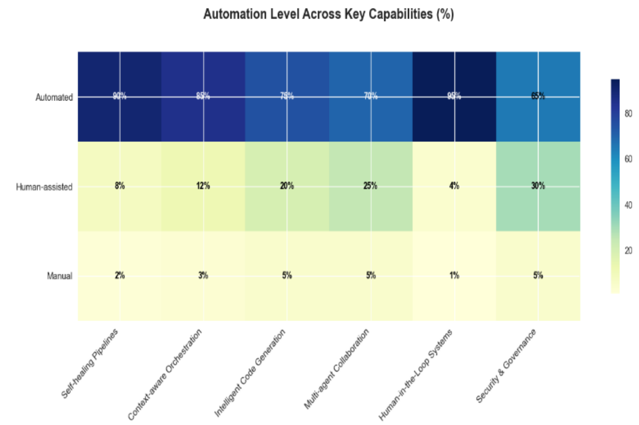

Fig. 5. Automation level across key capabilities.

VII. Mobile Ecosystem Challenge

Fig. 6 illustrates the impact of agentic AI orchestration on mobile ecosystem challenges. The findings provide a deep visual summary of the transformational potential of Agentic AI in managing the complex challenges of the mobile ecosystem. In contrast to conventional, human-mediated management, it demonstrates plainly that agentic AI is an underlying advancement rather than a slight improvement, despite its enormous challenges, but it still provides a better and more balanced orchestration capability for mobile ecosystems.

Comparative polygon visualization contrasting traditional management vs. agentic orchestration literature survey was performed, based on multi-dimensional criteria (performance, privacy, resource use, cross-platform integration) [2],[10], [13],[30],[31],[32]. Based on Fig. 6, the shapes of the polygons provide key information on the impact of agentic AI orchestration on mobile ecosystem challenges. The undersized, lopsided, and collapsed "traditional management" polygon suggests a poor, mixed approach, where fixing one issue often makes another worse (e.g., boosting performance drains the battery). The "agentic AI orchestration" polygon, in contrast, is massive, well-balanced, and sturdy, demonstrating a comprehensive and cooperative approach. The judgments made by an agentic AI orchestrator can simultaneously optimize for each of the four problems. To attain a single, unified objective, it can, for instance, carry out a plan that respects privacy (on-device processing), adjusts to context (the user is busy), controls resources (limited network calls), and functions across several app platforms.

These four issues combine to make human-mediated management impracticable, highlighting the urgent need for an autonomous orchestrator that can concurrently optimize for battery life, performance, and user comfort.

To validate the proposed implementation roadmap and automation levels, this review synthesizes findings from recent agentic AI deployments across mobile platforms, including AutoGen, LangChain, and OSS/BSS orchestration frameworks[2],[10],[13],[30],[31],[32]. These implementations demonstrate measurable success in context-aware execution, tool orchestration, and self-healing pipelines, with performance benchmarks showing improved task completion rates, reduced latency, and enhanced user satisfaction. Comparative analysis with traditional reactive systems reveals that agentic AI frameworks outperform legacy models in adaptability, autonomy, and cross-platform synchronization. For instance, while conventional systems rely on static rule-based triggers, agentic AI agents dynamically adjust strategies based on real-time feedback and resource constraints, offering a more resilient and scalable orchestration model. This comparative validation underscores the strategic advantage of agentic AI in mobile ecosystem management.

Unlike traditional mobile orchestration methods and generative AI models, agentic AI offers distinctive advantages that directly address ecosystem fragmentation and cognitive overload. Conventional systems rely on static rule-based triggers or reactive responses, which collapse under trade-offs such as performance versus battery life. Generative AI improves content creation but remains reactive, lacking persistent memory and multi-step orchestration. In contrast, agentic AI integrates reasoning engines, tool orchestration, memory, RAG, and safety layers into a continuous operational loop, enabling proactive, context-aware coordination across applications and devices. This allows agentic AI to simultaneously optimize performance, privacy, resource use, and cross-platform synchronization, as demonstrated in Fig. 4 to Fig. 6. Compared to similar frameworks such as AutoGen, LangChain, and OSS/BSS orchestration models, our synthesis uniquely emphasizes mobile-specific constraints (thermal throttling, resource-aware execution, and platform fragmentation) and provides a roadmap with automation benchmarks. These comparative advantages establish agentic AI not as an incremental improvement, but as a foundational orchestrator for next-generation mobile ecosystems.

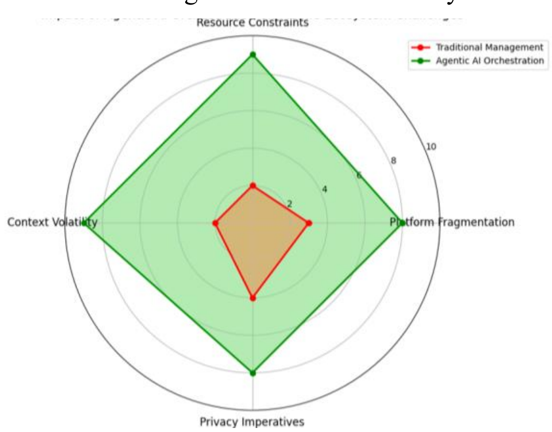

Fig. 6. Impact of agentic AI orchestration on mobile ecosystem challenges.

VIII. Future Directions and Research Agenda

Several fascinating research and development pathways are evident in the developing of agentic AI as a mobile ecosystem orchestrator.

- The goal of projects such as the Web of Agents proposal is to create minimal interoperability requirements for communication between agents. This would greatly lessen the fragmentation of mobile ecosystems [37],[38]. These standards would facilitate smooth collaboration between agents from various platforms and developers.
- The development of human-in-the-loop (HitL) architectures has produced frameworks for the best possible allocation of tasks between AI autonomy and human supervision [39],[40]. These models are especially important for mobile applications with high stakes, such as financial transactions or medical judgments.
- Edge-native architectures advanced agentic AI capabilities will be able to run directly on mobile devices with further development of specific processors and effective model architectures, improving privacy, lowering latency, and functioning without continual connectivity.
- Research is required to develop architectures that can sustain coherent activity over long periods of time while adapting to shifting user demands and technical settings. Long-term ecosystem adaptation of current systems presents remarkable short-term adaptability.
- The autonomous characteristics and environment of agentic AI entail new attack surfaces and susceptibility patterns, demanding the development of innovative security frameworks directed toward multi-agent, cross-platform mobile environments.
- There is a need for thorough assessment frameworks that measure efficiency beyond task execution and completion to include user satisfaction, resource efficacy, and ecosystem rationality because the area lacks standardized techniques for evaluating the performance of agentic AI systems in ecosystem orchestration.

IX. Conclusion

Agentic AI signifies a paradigm shift in mobile ecosystem orchestration, developing from reactive supporters to autonomous, goal-driven systems efficient across different applications and platforms. This review has summarized the architectural bases, operational components, and mobile-specific constraints that shape agentic AI deployment, emphasizing its transformational capability in managing complexity, optimizing resources, and enriching user experience. The combination of reasoning engines, RAG systems, and safety layers places agentic AI as an essential element for next-generation mobile computing, despite significant obstacles including tool fragility, memory constraints, and cross-platform synchronization. The proposed implementation roadmap and automation analysis provide strategic guidance for stakeholders steering this transition. Future research must address long-term rationality, security frameworks, and standardized assessment metrics to ensure scalable, ethical, and robust agentic AI systems. Agentic AI is eventually a redefining of intelligence in mobile environment, not just an improvement.

This review began by identifying the core challenges: mobile ecosystems are increasingly fragmented, reactive, and cognitively demanding for users. Traditional orchestration relies heavily on manual coordination, which fails to scale with complexity. Agentic AI directly responds to these limitations by offering autonomous, context-aware orchestration across applications, devices, and services. Through reasoning engines, tool orchestration, memory, and safety layers, agentic AI transforms mobile platforms into proactive, adaptive environments. This synthesis underscores agentic AI not merely as an enhancement, but as a necessary architectural shift for next-generation mobile computing.

ACKNOWLEDGMENT

The authors would like to thank Jadara University, Irbid, Jordan, for supporting this work.

REFERENCES

[1] A. Harima, J. Harima, and J. Freiling, "Ecosystem orchestration: Unpacking the leadership capabilities of anchor organizations in nascent entrepreneurial ecosystems," Entrepreneurship Theory and Practice, vol. 48, no. 6, pp. 1404-1450, 2024.

[2] R. Shah, S. Jagtap, and V. Jain, "Architecting analytics-driven mobile ecosystems: Scalable backend frameworks for intelligent data flow and real-time user insights," International Journal of Artificial Intelligence, Data Science, and Machine Learning, vol. 6, no. 2, pp. 83-91, 2025.

[3] L. Hughes et al., "AI agents and agentic systems: A multi-expert analysis," Journal of Computer Information Systems, pp. 1-29, 2025.

[4] Y. Shavit et al., "Practices for governing agentic AI systems," Research Paper, OpenAI, 2023.

[5] V. B. Komaragiri, "Agentic AI for autonomous network orchestration: A new frontier in telecommunications," Eksplorium-Buletin Pusat Teknologi Bahan Galian Nuklir, vol. 46, no. 1, pp. 221-241, 2025.

[6] S. Motamary, "Enabling zero-touch operations in telecom: The convergence of agentic AI and advanced DevOps for OSS/BSS ecosystems," BSS Ecosystems, Dec. 2022.

[7] A. Biswas and W. Talukdar, Building Agentic AI Systems: Create intelligent, autonomous AI agents that can reason, plan, and adapt. Packt Publishing Ltd., 2025.

[8] A. Bandi, B. Kongari, R. Naguru, S. Pasnoor, and S. V. Vilipala, "The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges," Future Internet, vol. 17, no. 9, p. 404, 2025.

[9] S. Feuerriegel, J. Hartmann, C. Janiesch, and P. Zschech, "Generative ai," Business & Information Systems Engineering, vol. 66, no. 1, pp. 111-126, 2024.

[10] J. Li et al., "Generative AI for self-adaptive systems: State of the art and research roadmap," ACM Transactions on Autonomous and Adaptive Systems, vol. 19, no. 3, pp. 1-60, 2024.

[11] E. T. Rolls, "The memory systems of the human brain and generative artificial intelligence," Heliyon, vol. 10, no. 11, 2024.

[12] S. Sivakumar, "Agentic AI in predictive AIOPs: Enhancing IT autonomy and performance," International Journal of Scientific Research and Management (IJSRM), vol. 12, no. 11, pp. 1631-1638, 2024.

[13] D. B. Acharya, K. Kuppan, and B. Divya, "Agentic AI: Autonomous intelligence for complex goals–a comprehensive survey," IEEE Access, 2025.

[14] Z. M. Auda and S. J. Radhi, "Artificial intelligence and evolution of the global system," IPRI Journal, vol. 22, no. 1, pp. 91-109, 2022.

[15] H. F. Hansen, E. Lillesund, P. Mikalef, and N. Altwaijry, "Understanding artificial intelligence diffusion through an AI capability maturity model," Information Systems Frontiers, vol. 26, no. 6, pp. 2147-2163, 2024.

[16] A. M. Taha, Z. K. Alkayyali, Q. M. Zarandah, and S. S. Abu-Naser, "The evolution of AI in autonomous systems: Innovations, challenges, and future prospects," Adaptive Systems, vol. 19, no. 3, pp. 1-60, 2024.

[17] Y. Wang et al., "Realization of empathy capability for the evolution of artificial intelligence using an MXene (Ti3C2)-based memristor," Electronics, vol. 13, no. 9, p. 1632, 2024.

[18] U. M. Borghoff, P. Bottoni, and R. Pareschi, "Human-artificial interaction in the age of agentic AI: A system-theoretical approach," Frontiers in Human Dynamics, vol. 7, p. 1579166, 2025.

[19] C. Wissuchek and P. Zschech, "Exploring agentic artificial intelligence systems: Towards a typological framework," in Pacific-Asia Conference on Information Systems, Kuala Lumpur, 2025, pp. 1-17.

[20] A. Chan et al., "Harms from increasingly agentic algorithmic systems," in Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, 2023, pp. 651-666.

[21] H. Clatterbuck, C. Castro, and A. M. Morán, "Risk alignment in agentic AI systems," arXiv preprint arXiv:2410.01927, 2024.

[22] R. V. Barenji and S. Khoshgoftar, "Agentic AI for autonomous anomaly management in complex systems," arXiv preprint arXiv:2507.15676, 2025.

[23] M. K. R. Jaggavarapu, "The evolution of agentic AI: Architecture and workflows for autonomous systems," Journal of Multidisciplinary, vol. 5, no. 7, pp. 418-427, 2025.

[24] P. B. Bansod, "Distinguishing autonomous AI agents from collaborative agentic systems: A comprehensive framework for understanding modern intelligent architectures," arXiv preprint arXiv:2506.01438, 2025.

[25] L. H. Cheung, L. Wang, and D. Lei, "Conversational, agentic AI-enhanced architectural design process: Three approaches to multimodal AI-enhanced early-stage performative design exploration," Architectural Intelligence, vol. 4, no. 1, pp. 1-25, 2025.

[26] J. Gu, "Position: Agentic systems constitute a key component of next-generation intelligent image processing," arXiv preprint arXiv:2505.16007, 2025.

[27] S. Deng et al., "Agentic services computing," arXiv preprint arXiv:2509.24380, 2025.

[28] I. Adabara et al., "Trustworthy agentic AI systems: A cross-layer review of architectures, threat models, and governance strategies for real-world deployment," F1000Research, vol. 14, p. 905, 2025.

[29] Z. Porter et al., "INSYTE: A classification framework for traditional to agentic AI systems," ACM Transactions on Autonomous and Adaptive Systems, vol. 20, no. 3, pp. 1-39, 2025.

[30] G. K. Sheelam, "Architecting agentic AI for real-time autonomous edge systems in next-gen mobile devices," Advances in Consumer Research, vol. 2, no. 3, 2025.

[31] H. Toivonen and F. Lelli, "The varieties of agency in human–smart device relationships: The four agency profiles," Future Internet, vol. 16, no. 3, p. 90, 2024.

[32] X. Wang et al., "Empowering edge intelligence: A comprehensive survey on on-device AI models," ACM Computing Surveys, vol. 57, no. 9, pp. 1-39, 2025.

[33] S. Raza, R. Sapkota, M. Karkee, and C. Emmanouilidis, "Responsible agentic reasoning and AI agents: A critical survey: Proposal for safe agentic AI via responsible reasoning AI agents (R2A2)," SuperIntelligence-Robotics-Safety & Alignment, vol. 2, no. 6, 2025.

[34] G. K. Sheelam, "AI-driven spectrum management: Using machine learning and agentic intelligence for dynamic wireless optimization," European Advanced Journal for Emerging Technologies (EAJET), vol. 2, no. 1, 2024.

[35] Z. Chen, Q. Sun, N. Li, X. Li, and Y. Wang, "Enabling mobile AI agent in 6G era: Architecture and key technologies," IEEE Network, vol. 38, no. 5, pp. 66-75, 2024.

[36] S. Panigrahy, "Multi-agentic AI systems: A comprehensive framework for enterprise digital transformation," Journal of Computer Science and Technology Studies, vol. 7, no. 6, pp. 86-96, 2025.

[37] R. Sharma, M. de Vos, P. Chari, R. Raskar, and A. M. Kermarrec, "Collaborative agentic AI needs interoperability across ecosystems," arXiv preprint arXiv:2505.21550, 2025.

[38] R. J. Tong et al., "IEEE AI standards for agentic systems," in 2025 IEEE Conference on Artificial Intelligence (CAI), 2025, pp. 1603-1609.

[39] M. V. Krishnamoorthy, "Enhancing responsible AGI development: Integrating human-in-the-loop approaches with blockchain-based smart contracts," Journal of Advances in Mathematics and Computer Science, vol. 39, no. 9, pp. 14-39, 2024.

[40] K. Zheng et al., "Towards agentic smart design: An industrial large model-driven human-in-the-loop agentic workflow for geometric modelling," Applied Soft Computing, p. 113920, 2025.