WWW.JBDAI.ORG

ISSN: 2692-7977

JBDAI Vol. 4, No. 1, 2026 DOI: 10.54116/jbdai.v4i1.82

EDITORIAL

FROM NLP TO LLMS TO AGENTIC AI: THE EVOLUTION OF ARTIFICIAL INTELLIGENCE

Jim Samuel
Rutgers University
AIXosphere
jim.samuel@rutgers.edu

Rajiv Kashyap
William Paterson University
KashyapR@wpunj.edu

Vishal Misra
Columbia University
vishal.misra@columbia.edu

Abhishek Tripathi
The College of New Jersey
tripatha@tcnj.edu

Hieu Duc Nguyen
Rowan University
nguyen@rowan.edu

Mahmoud Daneshmand
Stevens Institute of Technology
mdaneshm@stevens.edu

Alexander Pelaez
Hofstra University
5E Analytics
alexander.pelaez@hofstra.edu

1. Introduction

Generative artificial intelligence (AI) has permeated global enterprises and societies with a velocity that is without precedent in the history of technology. Yet we must remind ourselves that speed is not a strategy. It is merely a condition. What truly defines this era is the fundamental shift in how these systems have transitioned from calculators of language to actors of intent. We have reached an inflection point with the emergence of agentic AI (AAI) built on fundamental and generative AI building blocks (Samuel et al. 2024; Acharya et al. 2025; Tripathi et al. 2025). This transition from models that merely respond to those that possess autonomous, tool-integrated architectures demands a rigorous rethinking of our sociotechnical structures. One thing is now certain: governance and accountability can no longer be treated as administrative afterthoughts or late-stage constraints. They must be designed into the “nervous systems” of future technologies.

Where previous efforts settled for the fragmented automation of discrete tasks, agentic systems leverage the triad of autonomy, planning, and memory to command entire processes from end to end. They transform isolated tools into unified agents of performance. However, this has led us to the generative AI paradox: we see widespread adoption, yet measurable enterprise-level impacts remain stubbornly uneven. Although organizations invest and experiment with enthusiasm, transformative returns remain elusive. We have yet to align this new computational capacity with the disciplined practice of management and the rigorous requirements of institutional stability.

2. The Generative AI Paradox: Why High Adoption Has Not Meant High Impact

Understanding why this gap exists matters before accepting AAI as its solution. Three structural explanations are worth noting. First, there is a *deployment* problem. Many organizations have integrated generative AI into narrow, bounded tasks (e.g., drafting emails, summarizing documents) without embedding it into the decision workflows in which business value actually compounds. This adoption pattern is consistent with evidence that firms tend to experiment with AI in isolated, peripheral use cases rather than integrate it into core processes (Raisch and Krakowski 2021). Second, there is a *measurement* problem: the productivity gains from reactive large language models (LLM) often manifest as cognitive offload and time savings for knowledge workers, which are difficult to capture in traditional key performance indicators (KPIs) and return on investment (ROI) frameworks. Research results show that AI-driven productivity improvements frequently take the form of reduced cognitive load and time savings that organizations struggle to quantify (Davenport and Ronanki 2018; Luo et al. 2019). This measurement gap is consistent with work that shows that AI can improve human performance by adaptively reshaping information facets, leading to benefits that are real but often invisible to standard KPIs (Samuel et al. 2022). Third, and fundamentally, there is an *architectural* problem. Reactive systems, by design, can only assist. They respond to prompts but cannot initiate or maintain context across sessions and certainly cannot orchestrate multistep workflows without continuous human re-engagement. These limitations are well documented (Shneiderman 2020; Bommasani et al. 2021) and show that each interaction is, in effect, a dead end. The marginal value added by reactive AI is inherently low because humans must still serve as the connective tissue between model outputs and organizational action.

AAI offers a path out of this paradox by transforming generative AI from reactive generative assistance into proactive, goal-oriented collaborators capable of executing complex workflows (Randazzo et al. 2025). In principle, agentic systems have the potential to improve operational efficiency by enabling parallel execution, real-time adaptability, personalization, elasticity, and resilience (Iansiti and Lakhani 2020). In domains such as supply chains, agents continuously ingest data, anticipate risks, dynamically replan workflows, and optimize decisions across time horizons (Sukharevsky et al. 2025), and this could be extended to areas such as health care, where generative AI promises gains across clinical and administrative workflows (Bhuyan et al. 2025).

The economic stakes are substantial. By 2030, the U.S. B2C retail market alone could see up to $1 trillion in orchestrated revenue from agentic commerce, with global projections reaching as high as $3 trillion to $5 trillion (McKinsey and Company 2024). KPMG similarly estimates that AAI will be key to unlocking a staggering $3 trillion in corporate productivity improvements annually (KPMG International 2025). Their argument can be distilled into four enterprise value pathways: expanding the scope of work that can be automated, enabling always-on operational capacity, improving organizational resilience through continuous optimization, and converting institutional knowledge into executable actions. However, we note that these projections originate largely from consulting firms with commercial interests in accelerating enterprise AI adoption and should be read as directional signals rather than forecasts. In this framing, agent-powered commerce is not merely a user-experience innovation, it represents a structural shift in market dynamics in which intelligent systems can anticipate user intent, evaluate alternatives, negotiate options, and execute transactions with minimal human intervention (Gates 2023). Commercial validation is already visible: companies such as Anysphere (creator of Cursor) and Perplexity are reported to be reaching major revenue milestones (Bain and Company 2025). These signals reflect AAI's expanding deployment across domains, disciplines, and global challenges, establishing AAI as a critical frontier in contemporary AI innovation.

3. From NLP to LLMs

Natural language processing (NLP) serves as the root of AAI, with LLMs that form a powerful linguistic substrate for agentic behavior. Early NLP efforts (1950s) prioritized natural language understanding more than generation, with machine translation emerging as a defining application area (Chomsky 2002). The earliest era of NLP (1950–1969) was dominated by rule-based systems, built on handcrafted linguistic rules and patterns. A landmark example was Weizenbaum (1966), which used pattern matching to simulate conversation, demonstrating natural language interaction potential, despite lacking true semantic understanding (Weizenbaum 1966). The ALPAC report (Automatic Language Processing Advisory Committee (ALPAC) 1966) concluded that rule-based translation was economically unviable, triggering an early AI winter for NLP. The late 1980s-1990s marked a paradigm shift toward statistical and machine learning (ML) approaches driven by computational power and text corpora growth (Masoumzadeh et al. 2023). IBM's statistical machine translation demonstrated probabilistic model effectiveness. Support vector machines (SVMs) and conditional random fields (CRFs) enhanced NLP capabilities, and ML approaches supplanted rule-based systems on ambiguity-heavy tasks such as word sense disambiguation due to superior scaling (Navigli 2009). Big data, as defined by volume, velocity, and variety (Sagiroglu and Sinanc 2013), further transformed NLP by enabling data-driven approaches that fueled capabilities from sentiment analysis to entity recognition. (Chatterjee et al. 2022). Deep learning further accelerated this trajectory: recurrent neural networks (RNNs) and long short term memory (LSTMs) enabled robust sequence modeling for text and speech, and helped power widely adopted assistants such as Siri, Alexa, and Google Assistant (Gudivada et al. 2015).

LLMs represent a transformative breakthrough enabled by advances in neural architectures and large-scale training. The introduction of attention mechanisms and the Transformer architecture replaced recurrence with self-attention, reshaping the foundation of language modeling (Masoumzadeh et al. 2023). Transformer-based pretraining underlies modern LLM families, including GPT and BERT. Performance scaling with model size, often framed through parameter growth and broader training has been associated with emergent capabilities such as in-context learning and multistep reasoning. Yet, LLMs face persistent challenges in computational efficiency, factual reliability, interpretability, bias mitigation, and safe deployment; issues that do not automatically disappear with scale (Patil and Gudivada 2024). Critically, even as LLMs expanded what systems could generate and infer, standalone LLMs remained reactive: powerful responders to prompts but constrained in goal pursuit, memory, execution control, and tool-mediated action (Fauscette 2025). Prompt engineering and retrieval-augmented generation (RAG) extended usefulness by improving controllability and grounding, but multistep, real-world workflows still demanded substantial human orchestration (Fauscette 2025), whether the goal is transparent university support via auditable RAG pipelines (Chidipothu et al. 2025), emotion-consistent multilingual sentiment analysis across machine translation (Anderson et al. 2024), or culturally adaptive, human-centered educational AI that must flex to context rather than produce one-size-fits-all outputs (Samuel et al. 2023).

The past 2 years witnessed AAI's emergence as a new paradigm that transcends LLM limitations, shifting "LLMs as assistants" to agents as actors. With the advent of agents, AI has evolved from passive response systems to goal-driven systems; workflows are increasingly moving from reactive to proactive (Clark 2025; Fauscette 2025). AAI systems are autonomous systems pursuing complex goals with minimal intervention, demonstrating adaptability and self-sufficiency (Acharya et al. 2025). Agentic behavior is not a sudden invention; it builds incrementally on decades of research (Botti 2025) and draws heavily on information systems concepts: autonomy, workflow logic, and tool integration to convert LLMs from generative engines into active problem-solvers (Bandi et al. 2025; Fauscette 2025). At a systems level, AAI extends LLM capability through state maintenance, persistent memory, goal definition, multistep reasoning, decision-making, tool use, and feedback loops (Bandi et al. 2025; Fauscette 2025). This is what distinguishes meaningful agency from simple orchestration and is central to the value creation narrative now emerging across sectors (McKinsey and Company 2024; Fauscette 2025). In this confluence, NLP enables interaction and task framing, LLMs provide linguistic intelligence and reasoning, and agentic architectures provide autonomy, a transition that can be framed as moving from language understanding to language reasoning to language-driven action (Fauscette 2025).

The rapid rise of AAI has been accelerated by modular frameworks such as LangGraph, CrewAI, AutoGen, Agno, SmolAgents, Mastra, Pydantic AI, and Atomic Agents. Examples include Autodata converting natural-language instructions into datasets (Ma et al. 2025); medical AI pipelines automating data processing (Shimgekar et al. 2025); AI Cosmologist performing autonomous research cycles (Moss 2025); and MLR-Copilot generating ML research ideas (Li et al. 2024). Yet the frontier is defined as much by risks as by capability.

Challenges persist: hallucination and reasoning failures (Hosseini and Seilani 2025; Raheem and Hossain 2025); opaque decision-making that can perpetuate bias and inequity (Raheem and Hossain 2025); emergent agency that creates safety and alignment risks (Raheem and Hossain 2025); governance gaps because systems with action-taking capability are deployed without sufficient oversight (Hosseini and Seilani 2025; Raheem and Hossain 2025); and insufficient robustness and safety features in complex, unpredictable environments (Raheem and Hossain 2025).

In looking ahead, AI agents can be expected to evolve as constrained societal actors, structured through frameworks such as FEEG (Finder, Evaluator, Explainer, Generator) as an intent framework that governs mode of operation, retrieval, and sourcing (Finder); criteria-based judgment (Evaluator); stepwise clarification (Explainer); and bounded creation (Generator), thereby calibrating verification depth, tool use, and escalation thresholds (Samuel et al. 2025b). Such agents could be used to tackle a range of informational challenges, such as fake news or AI phobia. When failures, uncertainty, or opaque actions intensify news-driven "AI phobia," agents can help detect fear sentiment signals in discourse and user interactions, surface early warnings, and support timely, evidence-based responses (Samuel et al. 2025a). There is a need to develop more trust-aware agents that monitor sentiment drift, offer transparent explanations, and escalate sensitive cases to human oversight when needed.

4. Risks: Why Agentic Failures Are Qualitatively Different

The risks outlined above are real but also generic, which are hallmarks of any sufficiently complex AI system. What distinguishes agentic systems is not simply the presence of risk but the structure of that risk. When organizations move from reactive models to autonomous, multistep agents, the nature of failure changes. It becomes compounded, distributed, and far more difficult to detect. Effective management depends on the clarity of responsibility and the integrity of the underlying system of action. When those foundations shift, so too does the character of managerial risk (Drucker 1974).

The most consequential challenge is *error propagation*. In a reactive LLM, a hallucination or reasoning lapse appears immediately in the response to a single prompt, where a human can intercept and correct it. But an agentic system operates across a chain of interdependent steps, retrieving information, planning actions, invoking tools, and interacting with external systems. An early-stage error can move silently through this chain, shaping downstream decisions long before a human has any opportunity to intervene. Recent analyses of multi-agent workflows show that a single hidden message can compromise an entire system of agents, cascading through planning and execution stages in ways that are difficult to trace or reverse (Schultz 2025). Similarly, retrieval-augmented and tool-using agents exhibit multistage vulnerabilities such as context manipulation and cross-context contamination that allow early errors to silently influence later actions (Ramakrishnan and Balaji 2025). The result is that the failure is no longer local but systemic: the error becomes embedded in the architecture of the workflow itself.

A second challenge concerns *accountability*. Traditional lines of responsibility are blurred when an agent acts on behalf of a user by submitting a transaction, sending a communication, or modifying a record. Managerial effectiveness rests on unambiguous responsibility. However, agentic systems are designed to be autonomous decision systems. They exercise judgment across models, interfaces, and organizational processes in ways that existing legal and managerial frameworks were never designed to accommodate. Responsibility becomes difficult to assign when autonomous systems execute multistep actions across legal, organizational, and technical boundaries (Gulyamov et al. 2026). The question "who is responsible" becomes harder to answer when the system itself exercises delegated discretion.

Third, *tools-using agents introduce adversarial attack surfaces that reactive systems simply do not face. A language model generating text can be manipulated at the prompt level*. But an agent empowered to search the web, call application programming interfaces (APIs), or interact with external environments inherits the vulnerabilities of those environments. Multimodal and cross-context prompt injection attacks can compromise agentic systems and evade existing guardrails (Lee and Tiwari 2024). As agents gain greater autonomy and environmental access, the attack surface expands accordingly. Prompt injection remains a fundamental vulnerability in LLM-based systems, especially when models are embedded in multi-agent workflows or are granted access to external tools.

5. The Path Ahead. . .

The evolution from computational linguistics through statistical methods, big data, machine learning, and LLMs to AAI represents progression from language processing to language reasoning to, ultimately, language-driven action. Each phase, computational linguistics providing theoretical foundations, data science enabling pattern learning, algorithmic topic modeling making insights accessible, machine learning achieving robust performance, LLMs unifying understanding and generation, prompt engineering and RAG extending capabilities, serves as an irreplaceable foundational building block. Yet, in many practical settings, the path ahead increasingly requires some form of agentic architecture. Whereas standalone LLMs and RAG excel at narrow, reactive tasks, modern enterprise and societal needs demand systems that can pursue goals, maintain context, reason across steps, and orchestrate tools dynamically.

AAI represents a fundamental paradigm shift in strategic operations. These systems integrate core information systems concepts such as autonomy and workflow logic with the advanced intelligence of LLMs. By incorporating the foundations of computational linguistics, AAI moves beyond narrow, single-step applications. The emerging ecosystem uses modular architectures that can range from lightweight agents for bounded workflows to heavyweight systems for enterprise-scale coordination. This transition moves organizations away from reactive assistance toward proactive collaboration. It replaces isolated tasks with end-to-end automation. Ultimately, this shift evolves organizational praxis from human-orchestrated pipelines to increasingly autonomous operations.

The convergence of computational linguistics, NLP, and generative AI within agentic frameworks represents a strategic evolution capable of capturing enormous value that discrete task automation cannot reach. The surrender of human oversight to autonomous agency introduces profound systemic vulnerabilities. In this high-stakes environment, interpretability loss and alignment drift cease to be mere technical glitches; they become existential threats to institutional integrity, mandating an uncompromising architecture of sociotechnical governance.

The strategic path forward lies in the synthesis of these technologies. We need agentic architectures that transform linguistic intelligence into reliable autonomous action. Systems that support and strengthen human judgment rather than replace it. In this paradigm, language serves as the interface, whereas agency functions as the operating system, translating prompts into plans, and intelligence into accountable, outcomes-driven action at scale. Language frames the intent while agency captures the value. Yet, as we turn prompts into plans and intelligence into autonomous action, the finish line is only a victory if ethically steered toward win-win outcomes. In an era of expanded systemic risk and distributed agency, success cannot simplistically be determined by the speed of execution or a "I got here first" mindset, but by innovative and sustainable value creation, human enhancive AI based governance and stakeholder co-creation, that ensure that the future we are developing will lead to a better world for all.

References

Acharya, D. B., K. Kuppan, and B. Divya. 2025. "Agentic AI: Autonomous Intelligence for Complex Goals." *IEEE Access* **13**: 18912–18936.

Anderson, R., C. Scala, J. Samuel, V. Kumar, and P. Jain. 2024. "Are Emotions Conveyed Across Machine Translations? Establishing an Analytical Process for the Effectiveness of Multilingual Sentiment Analysis with Italian Text." *Journal of Big Data and Artificial Intelligence* **2**, no. 1: 57–73. https://doi.org/10.54116/jbdai.v2i1.30

Automatic Language Processing Advisory Committee (ALPAC). 1966. *Language and Machines: Computers in Translation and Linguistics*. Washington, DC: National Academy of Sciences-National Research Council.

Bain and Company. 2025. "State of the Art of Agentic AI Transformation: Technology Report." Accessed March 25, 2026. https://www.bain.com/insights/state-of-the-art-of-agentic-ai-transformation-technology-report-2025/

Bandi, A., B. Kongari, R. Naguru, S. Pasnoor, and S. V. Vilipala. 2025. "The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges." *Future Internet* **17**, no. 9: 404. https://doi.org/10.3390/fi17090404

Bhuyan, S. S., V. Sateesh, N. Mukul, A. Galvankar, A. Mahmood, M. Nauman, et al. 2025. "Generative Artificial Intelligence Use in Healthcare: Opportunities for Clinical Excellence and Administrative Efficiency." *Journal of Medical Systems* **49**, no. 1: 10.

Bommasani, R., D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, et al. 2021. "On the Opportunities and Risks of Foundation Models." Preprint, arXiv:2108.07258.

Botti, V. 2025. "Agentic AI and Multiagentic: Are We Reinventing the Wheel?" Preprint, submitted June 2, 2025. https://doi.org/10.48550/arXiv.2506.01463

Chatterjee, S., R. Chaudhuri, and P. Mikalef. 2022. "Examining the Dimensions of Adopting NLP and Big Data." *IEEE Transactions on Engineering Management* **71**: 3001–3015. https://doi.org/10.1109/TEM.2022.3146430

Chidipothu, N., R. Anderson, J. Samuel, A. Pelaez, J. Esguerra, and M. N. Hoque. 2025. "Improving Large Language Model (LLM) Performance with Retrieval Augmented Generation (RAG): Development of a Transparent Generative AI University Support System for Educational Purposes." *Journal of Big Data and Artificial Intelligence* **3**, no. 1: 1–15.

Chomsky, N. 2002. *Syntactic Structures*. 2nd edition. Walter de Gruyter.

Clark, S. 2025. "What Are AI Agents? The Autonomous Software Changing How Work Gets Done." VKTR. Accessed March 25, 2026. https://www.vktr.com/digital-experience/what-are-ai-agents-the-autonomous-software-changing-how-work-gets-done/

Davenport, T. H., and R. Ronanki. 2018. "Artificial Intelligence for the Real World." *Harvard Business Review* **96**, no. 1: 108–116.

Drucker, P. F. 1974. *Management: Tasks, Responsibilities, Practices*. Harper and Row.

Fauscette, M. 2025. "Agentic AI vs. LLMs: Understanding the Shift from Reactive to Proactive AI." Arion Research. Accessed March 25, 2026. https://www.arionresearch.com/blog/agentic-ai-vs-llms-understanding-the-shift-from-reactive-to-proactive-ai/

Gates, B. 2023. "AI-Powered Agents Are the Future of Computing." Gates Notes. Accessed March 25, 2026. https://www.gatesnotes.com/ai-agents

Gudivada, V. N., D. Rao, and V. V. Raghavan. 2015. "Big Data Driven NLP Research." In *Handbook of Statistics*, edited by C. R. Rao and V. Govindaraju, vol. 33, 203–238. Elsevier. https://doi.org/10.1016/B978-0-444-63492-9.00008-9

Gulyamov, S., S. Gulyamov, A. Rodionov, R. Khursanov, K. Mekhmonov, D. Babaev, et al. 2026. "Prompt Injection Attacks in Large Language Models and AI Agent Systems: A Comprehensive Review of Vulnerabilities, Attack Vectors, and Defense Mechanisms." *Information* **17**, no. 1: 54. https://doi.org/10.3390/info17010054

Hosseini, S., and H. Seilani. 2025. "The Role of Agentic AI." *Array, Advance Online Publication* 100399. https://doi.org/10.1016/j.array.2025.100399

Iansiti, M., and K. R. Lakhani. 2020. "Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World." Harvard Business Review Press.

KPMG International. 2025. "The Agentic AI Advantage: How Autonomous AI Systems Are Reshaping Value Creation." Accessed March 25, 2026. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/10/agentic-ai-advantage-report.pdf.

Lee, D., and M. Tiwari. 2024. "Prompt Infection: LLM to LLM Prompt Injection Within Multi Agent Systems." Preprint, arXiv. 2410.07283. https://arxiv.org/abs/2410.07283

Li, R., T. Patel, Q. Wang, and X. Du. 2024. "Mlr-Copilot: Autonomous Machine Learning Research Based on Large Language Models Agents." Preprint, arXiv:2408.14033.

Luo, X., S. Tong, Z. Fang, and Z. Qu. 2019. "Frontiers: Machines vs. Humans: The Impact of AI on Productivity." *Marketing Science* **38**, no. 6: 937–947.

Ma, T., Y. Qian, Z. Zhang, Z. Wang, X. Qian, F. Bai, et al. 2025. "AutoData: A Multi-Agent System for Open Web Data Collection." Preprint, arXiv:2505.15859. https://arxiv.org/abs/2505.15859

Masoumzadeh, S., C. Bosman, and N. Osborne. 2023. "Becoming Walkable: Relational and Contextual Effects of Enhanced Walkability." *Journal of Urbanism: International Research on Placemaking and Urban Sustainability* **16**, no. 3: 341–357. https://doi.org/10.1080/17549175.2021.1963811

McKinsey and Company. 2024. "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value." Accessed March 25, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Moss, A. 2025. "The AI Cosmologist I: An Agentic System for Automated Data Analysis." Preprint. Accessed March 25, 2026. https://arxiv.org/abs/2504.03424

Navigli, R. 2009. "Word Sense Disambiguation: A Survey." *ACM Computing Surveys* **41**, no. 2: 1–69.

Patil, R., and V. Gudivada. 2024. "A Review of Current Trends, Techniques, and Challenges in Large Language Models (LLMs)." *Applied Sciences* **14**, no. 5: 2074. https://doi.org/10.3390/app14052074

Raheem, T., and G. Hossain. 2025. "Agentic AI Systems: Opportunities, Challenges, and Trustworthiness." In *Proceedings of the 2025 IEEE International Conference on Electro Information Technology (eIT). IEEE* 618–624. https://doi.org/10.1109/eIT64391.2025.11103638

Raisch, S., and S. Krakowski. 2021. "Artificial Intelligence and Management: The Automation–Augmentation Paradox." *Academy of Management Review* **46**, no. 1: 192–210.

Ramakrishnan, B., and A. Balaji. 2025. "Securing AI Agents Against Prompt Injection Attacks: A Comprehensive Benchmark and Defense Framework." Preprint, arXiv:2511.15759.

Randazzo, S., K. R. Mills, L. Peng, S. H. Lee, A. Bernstein, F. Dell'Acqua, et al. 2025. "Cyborgs, Centaurs and Self-Automators: The Three Modes of Human-GenAI Knowledge Work." Harvard Business School Working Paper, no. 26-036.

Sagiroglu, S., and D. Sinanc. 2013. "Big Data: A Review." In *2013 International Conference on Collaboration Technologies and Systems (CTS)*. 42–47. IEEE. https://doi.org/10.1109/CTS.2013.6567202

Samuel, J., R. Kashyap, Y. Samuel, and A. Pelaez. 2022. "Adaptive Cognitive Fit: Artificial Intelligence Augmented Management of Information Facets and Representations." *International Journal of Information Management* **65**: 102505.

Samuel, J., T. Khanna, J. Esguerra, S. Sundar, A. Pelaez, and S. S. Bhuyan. 2025a. "The Rise of Artificial Intelligence Phobia! Unveiling News-Driven Spread of AI Fear Sentiment using ML, NLP and LLMs." *IEEE Access* **13**: 125944.

Samuel, J., A. Tripathi, and E. Mema. 2024. "A New Era of Artificial Intelligence Begins. . . Where Will It Lead Us?" *Journal of Big Data and Artificial Intelligence* **2**, no. 1: 1–6.

Samuel, J., M. Ye, T. Khanna, Y. Yao, X. Lin, R. Anderson, et al. 2025b. "Finder, Evaluator, Explainer, Generator (FEEG): A Bloom's Taxonomy-Based Query Classification Framework for LLMs and Generative AI." *Expert Systems with Applications* **237**: 130755.

Samuel, Y., M. Brennan-Tonetta, J. Samuel, R. Kashyap, V. Kumar, K. K. M. Sri, et al. 2023. "Cultivation of Human Centered Artificial Intelligence: Culturally Adaptive Thinking in Education (CATE) for AI." *Frontiers in Artificial Intelligence* **6**: 1185745, ISSN=2624-8212

Schultz, J. S. 2025. "On the Potential of Agentic Workflows for Animal Training Plan Generation." *Frontiers in Veterinary Science* **12**: 1563233. https://doi.org/10.3389/fvets.2025.1563233

Shimgekar, S. R., S. Patil, P. Sharma, P. Khatri, R. Singh, and A. Gupta. 2025. "Agentic AI Framework for End-To-End Medical data Inference." Preprint, arXiv. Accessed March 25, 2026. https://arxiv.org/abs/2507.18115

Shneiderman, B. 2020. "Human-Centered Artificial Intelligence: Reliable, Safe and Trustworthy." *International Journal of Human–Computer Interaction* **36**, no. 6: 495–504. https://doi.org/10.1080/10447318.2020.1741118

Sukharevsky, A., D. Kerr, K. Hjartar, L. Hämäläinen, S. Bout, V. Di Leo, and G. Dagorret. 2025. "Seizing the agentic AI advantage. McKinsey and Company – QuantumBlack." Accessed March 25, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage

Tripathi, A., J. Samuel, M. Brennan-Tonetta, H. Nguyen, and E. Mema. 2025. "When Machines Create! Envisioning Our Future with the Transformative Power of Generative AI." *Journal of Big Data and Artificial Intelligence* **3**, no. 1: 16–24. https://doi.org/10.54116/jbdai.v3i1.55

Weizenbaum, J. 1966. "ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine." *Communications of the ACM* **9**, no. 1: 36–45. https://doi.org/10.1145/365153.365168