Human Factors in Artificial Intelligence and Social Computing, Vol. 203, 2026, 13–23
https://doi.org/10.54941/ahfe1007313

Deliver Cross-Process Automation Across Finance, HR, Procurement by Orchestrating Actions Across Diverse Systems-Powered by AI & Governed Workflows

Elizabeth Koumpan[1], Laurentiu Gabriel Ghergu[2], Łukasz Strąk[3], and Grzegorz Jurek[3]

[1]IBM, BTO Business Consulting Services, ON, Canada
[2]IBM Consulting, Bucharest, B 060201, Romania
[3]IBM BTO Business Consulting Services Sp.z.o.o. Kraków, Poland

ABSTRACT

The rapid diffusion of data driven automation and agentic AI systems is reshaping the foundations of work, decision making, and human–technology interaction. As organizations move toward Society 5.0-Japan's vision for a human-centered "super smart" society in which cyber-physical intelligence augments human capability across economic and social systems, there is an urgent need for operational architectures that are not only technologically capable but also fundamentally humancentric. This paper presents an applied model using Smarter Operations framework that integrates agentic AI, enterprise data fabric, human-in-the-loop governance, secure multisystem orchestration, and enterprise digital twins that simulate processes and operational states for context-aware decision support. The result is an adaptive sociotechnical system that enhances human decision-making rather than replacing it, while simultaneously enabling automation at operational scale. This research demonstrates that when designed with ergonomics, human values, and sociotechnical principles at the center, agentic AI become powerful enablers of humancentric, resilient, and adaptive enterprises.

Keywords: Agentic AI, Intelligent operations, Socio-technical systems, Human-in-the-Loop governance, Enterprise automation, Society5, Digital twins, Workflow orchestration

INTRODUCTION

In the rapidly evolving landscape of enterprise operations, the integration of artificial intelligence (AI) and machine learning (ML) has emerged as a transformative force, driving unprecedented efficiency and innovation across industries. Organizations are increasingly moving beyond simple robotic process automation (RPA) toward "Smarter Operations," where generative AI, Large Language Models (LLMs), and autonomous AI agents are embedded directly into core enterprise resource planning (ERP) systems (Balla, 2025; Sridharan, 2025).

We are entering an era where multiagent AI systems (Figure 1) become the norm, with autonomous agents capable of contextual reasoning and

Received March 2, 2026; Revised April 7, 2026; Accepted April 21, 2026; Available online July 20, 2026

dynamic tool selection to achieve complex objectives. As we advance toward the Society 5.0 paradigm, we will see the emergence of new interaction models powered by networks of intelligent, interconnected AI agents, fundamentally reshaping work, engagement, and digital ecosystems (Koumpan, Topol, 2023).

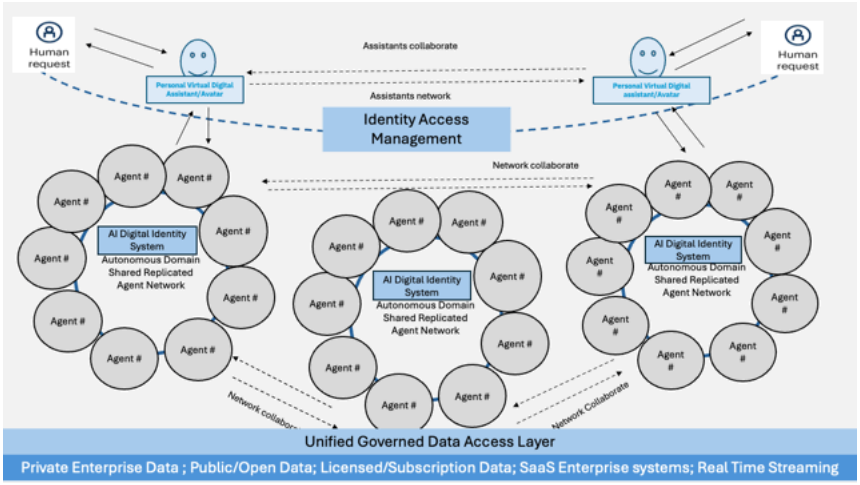

Figure 1: Interconnected multi-agent AI systems (IBM, 2024).

In Finance, HR, and Procurement, new technologies deliver major gains in productivity, document accuracy, and real-time decision support (Devaraju, 2024). AI-powered HR systems streamline hiring and workforce optimization, while Procurement benefits from faster supplier onboarding and higher master-data quality (Szelagowski et al., 2022; Devaraju, 2024).

Yet new technology is bringing new challenges. The designers of automated solutions need to think about end-to-end processes, understanding dependencies, interconnections, rather than in terms of tasks. Such transformation journey requires establishing a solid foundation, breaking down disconnected processes and data silos (Koumpan, Topol, 2021 to deliver human-centered, real-time, and collaborative services in secure enterprise. Limited data accessibility, weak auditability, and insufficient explainability of AI agents, combined with the inability to grant agents or users **precisely the right permissions at the right time**, remain major barriers to realizing fully connected agentic ecosystems. Addressing these gaps is essential to unlock the humancentric, trusted AI environments envisioned in Society 5.0.

DATA IS THE FOUNDATION FOR AGENTIC AI

Agentic AI can only reach its full potential when it is built on a strong data foundation-data that is connected, secured, governed, contextualized, and trusted.

In today's landscape, foundation models often operate as opaque systems, making it difficult to trace how inputs translate into outputs or why an

agent took a particular action. Their complex architectures, multilayer reasoning, and dynamic tool usage introduce challenges for explainability, interpretability, accountability, and auditability.

As agentic ecosystems grow, so does the volume of sensitive information they process, from personal data to proprietary and financial records, making robust identity, authentication, authorization essential. Agentlevel digital identity, finegrained permissions, and infrastructurelevel security controls must become standard to protect against misuse, breaches, or unauthorized autonomy.

Data is the fuel of agentic intelligence, but it comes in two essential forms: the data ingested into models for training, and the data that remains outside the model and must be accessed dynamically at runtime. Both are critical yet serve different purposes.

Training data shapes the model's foundational knowledge, while external operational data supplies the realtime context required for accurate reasoning, compliant actions, and effective decisionmaking. Raw data in either form does not create value on its own. Value emerges only when training data and runtime data are transformed into clean, connected, contextualized data products with clear lineage and governance.

The quality and integrity of both types of data directly determine how well AI agents can interpret situations, collaborate across systems, and execute business processes safely and reliably. Without this foundation, even the most advanced agentic systems will struggle to reason consistently or act with confidence across domains.

This becomes especially clear when looking at finance processes, where disconnected data silos frequently create friction and risk:

- Order-to-Cash (O2C): begins with customer orders and spans fulfillment, shipping, invoicing, payment collection, and cash application.
- Source-to-Pay (S2P): manages the full procurement lifecycle, need identification, sourcing, contracting, requisitions, POs, goods receipt, invoicing, and supplier payments.
- Record-to-Report & Reconciliation (R2R): involves collecting, matching, validating, and reporting financial data across ledgers, contracts, bank statements, and subledgers.

These processes are deeply interdependent. For example, delays or mismatches in O2C impact liquidity forecasts used in S2P planning. Contract terms in S2P influence reconciliation rules in R2R. And revenue and expense patterns identified in O2C and S2P inform anomaly detection and compliance workflows in R2R.

Connecting insights across O2C, S2P, and R2R is essential to achieve cohesive financial management and operational efficiency. It requires unified data models, crossprocess lineage, and agentic AI that can reason across endpoints-automating reconciliation, detecting anomalies, enriching context, and enabling predictive decision-making.

To build a fully connected digital ecosystem, where agents seamlessly coordinate across legacy and modern platforms, we need unified data and

integration capabilities and a governance framework rooted in the following principles:

- **Transparency & Explainability** - Agents must provide the rationale behind decisions and actions.
- **Ethics & Regulatory Compliance** - Automated checks to prevent bias and ensure policy alignment.
- **Security & Privacy** - Anonymization, masking, agent digital identity, and secure interagent communication.
- **Agent Identity & Permissioning** - Standardized impersonation when appropriate, with dynamic, timebound access controls. Logging on when it happens to ensure an audit chain of responsibility not lost.
- **CrossNetwork Authentication & Authorization** - Consistent, reusable mechanisms across environments.
- **Observability** - Ability to trace what agents did, why, and what data influenced their reasoning.
- **Autonomy Boundaries** - Clear definitions of operational limits, humanintheloop requirements, and accountability.
- **Interoperability & Orchestration** - Open standards (e.g., MCP) enabling agents to collaborate safely and effectively.
- **Continuous Monitoring & Feedback Loops** - Ongoing model improvement and safety assurance.
- **Data Strategy & Governance** - Strong quality, lineage, and confidentiality controls to ensure reliable insights.

Building a connected data ecosystem- grounded in a unified data model and robust integration fabric spanning both legacy and modern systems- is a foundational decision for any organization aspiring to develop human-centered, responsible AI.

In a world where multiple AI agents collaborate across domains, data connectivity, governance, explainability, security, and interoperability are not technical enhancements; they are the essential prerequisites. When these elements are in place, agentic AI can operate safely, transparently, and intelligently across endtoend financial and operational processes, bringing to life the connected enterprise envisioned in Society 5.0.

This is where the next evolution emerges: a Smarter Operations Hub, acting as the orchestration layer that binds data, systems, and AI agents together. It becomes the nerve center that turns connected data into coordinated action, enabling enterprisewide automation, augmentation, and intelligent decision-making.

Let's explore how this hub transforms operations and unlocks the full potential of agentic AI.

SMARTER OPERATIONS HUB FOR CROSS PROCESS AUTOMATION

Digital Twins

AI agents today function as autonomous digital workers, capable of managing business processes inside ERP ecosystems using advanced communication

protocols and distributed consensus for crossmodule coordination (Sridharan, 2025). In platforms like SAP, collaborative Joule agents already integrate procurement, finance, and logistics to boost efficiency, decision accuracy, and cross-functional synergy (Eyo et al., 2025). Generative AI and LLMdriven cognitive process automation further push ERP systems beyond rigid workflows, enabling contextaware reasoning, explanation generation, and continuous adaptation (Balla, 2025; Zuo, 2025). Oracle Cloud follows a similar trajectory, applying AI to data mapping, document intelligence, and workflow automation, with roadmaps increasingly oriented toward agentdriven, cross-functional orchestration (Gopalaswamy, 2025).

In parallel, digital twins have matured as executable, synchronized models of enterprise processes, assets, supply flows, and resource allocations. They enable realtime simulation, optimization, and intelligent planning, especially when paired with multi-agent ecosystems (Galuzin et al., 2022; Roxanne et al., 2025; Stokes, 2025). For example, the Autonomous Digital Twin of Enterprise (ADTE) framework demonstrates how ontological knowledge models and multi-agent systems can adaptively coordinate tasks, competencies, and resource plans across dynamic environments (Galuzin et al., 2022).

However, despite their complementary strengths, digital twins and agentic AI have largely evolved in isolation. Digital twins excel at simulation and scenario testing within their respective domains. Agentic AI excels at autonomous execution and crosssystem task completion. What is missing is the connective tissue that unifies insight (digital twins) with action (agentic AI).

This is precisely where the Smarter Operations Hub (SOH)(Figure 2) becomes transformative.

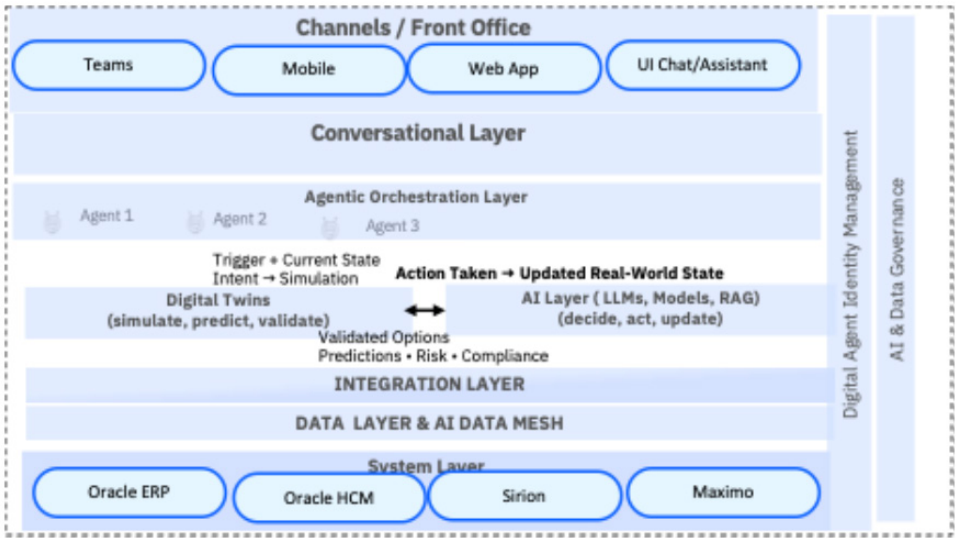

Figure 2: Smarter operations hub high level diagram (IBM, 2025).

Acting as a crossprocess orchestration layer, the SOH connects operational digital twins with enterprise workflows, integrates data fabrics and knowledge graphs, governs inter-agent communication and identity, and orchestrates

agentic automation across ERP, procurement, finance, supply chain, HR, and service operations.

The SOH enables **closed-loop operations**, where digital twins and agentic AI work together in a continuous, closed-loop decision cycle rather than as static layers.

- Agentic AI detects triggers in the environment and sends relevant state and intent information to the digital twin, which runs simulations to predict performance, cost, risk, and compliance outcomes.
- The digital twin then returns validated recommendations, allowing the agent to decide and execute actions such as issuing purchase orders, reallocating resources, adjusting forecasts, or triggering fulfillment.
- After acting, the agent updates the digital twin with the new operational state so the twin remains an accurate, living representation of reality.

This creates a dynamic feedback loop in which the digital twin provides predictive intelligence and guardrails, while the agent performs real-world actions and keeps the simulation current-ensuring decisions are both context-aware and continuously optimized.

SOH provides governance, supervision, and traceability, ensuring explainability, observability, and cross-system accountability. The result is a safer, smarter, continuously optimizing enterprise, one where simulation and execution work together in real time, reducing operational risk while maximizing efficiency and agility.

Let's look in the details of our Smarter Operations Hub.

Smarter Operations Hub Detailed Framework

The Smarter Operations Hub (SOH) is an emerging architectural layer that sits between the front office (sales, customer service, revenue generation) and the back office (ERP, finance, HR, compliance). Its purpose is to unify and orchestrate the increasingly complex operational landscape created by AI, automation, and multiagent systems. The SOH functions as the intelligent connective tissue of the enterprise, ensuring that data, workflows, agents, and systems operate as one coordinated whole.

The SOH provides capabilities that traditional integration layers and workflow systems cannot deliver:

- Unified data access across legacy and modern platforms with shared semantics and consistent lineage
- Crossdomain workflow orchestration, connecting processes that span multiple enterprise systems
- Multiagent coordination using open, interoperable protocols like MCP
- Finegrained identity, authentication, and permissioning for human and AI agents
- Endtoend process intelligence, including monitoring, anomaly detection, and simulation
- Governed automation aligned to organizational policies, controls, and compliance requirements

In highly regulated and dataintensive industries-such as energy & utilities, air transportation, financial services, insurance, healthcare, and manufacturing-the SOH becomes the operating system of operations, harmonizing digital twins, agentic AI, and enterprise systems into one coherent and trustworthy ecosystem. If the front office engages customers and the back office executes core processes, the SOH is what makes everything work together intelligently.

The SOH serves as the operational core, and consolidates multi-step workflows, case management, and integration orchestration across systems such as Oracle ERP, Workday, Maximo, CLM platforms, ServiceNow, and industry-specific operational tools. By unifying data into a shared semantic layer, the SOH allows intelligent agents to:

- Monitor processes in real time
- Detect anomalies or compliance deviations
- Trigger cross-system remediations
- Simulate impacts before execution (via digital twins)
- Execute actions securely with full traceability

This transforms previously disconnected workflows into continuous, closed-loop operational processes with increased transparency, reduced manual effort, and proactive control.

The SOH also incorporates conversational and natural-language interfaces, enabling workers to initiate, query, and manage workflows simply by describing what they need. Agentic automation then executes tasks across systems securely, respecting identity, permissions, and policy constraints. A core principle of the SOH is controlled autonomy: all automated decisions and agent-driven actions must pass through standardized controls, approvals, and audit trails. The unified data layer enforces consistent data access governance so that only authorized agents-or humans-can interact with sensitive information. All integrations with systems of record run through enterprisegrade authentication, encryption, and monitoring.

SOH Architecture

A modern Smarter Operations Hub (Figure#3) is composed of the following layers:

- Engagement Layer - Interfaces tailored to user personas, including AI assistants and copilots, sidebars embedded in enterprise apps, web portals and mobile apps. This multimodal layer provides intuitive access for employees, managers, and decisionmakers.
- SOH Core (Orchestration and Governance layer) - connecting the frontend and backend, built using orchestration frameworks. This may include tools such as Watsonx Orchestrate, the Lang* suite, etc. It is responsible for A2A (application-to-application) collaboration and humanintheloop governance.
- AI and automation ecosystem -managed by the SOH Core, this ecosystem consists of AI agents, RPA bots, LLMs reasoning engines, and legacy

applications accessible via APIs. Agents use MCP or similar standards to collaborate and invoke tools safely

- Data layer – a unified data foundation, consisting of domain data models, canonical schemas, knowledge graphs, historical data, event logs, digital twins' signals. This layer enables contextual reasoning, lineage tracking and high quality decisioning.
- Integration layer -designed with an APIfirst approach, supporting API Gateways, Event /streaming connectors, ETL/ELT for legacy systems, direct adaptors for on- premises platforms.
- Enterprise applications layer -Composed of all applications the company uses to execute its core business processes, such as Oracle, SAP, Workday, ServiceNow, Maximo, CLM, CRM, SCM, MRO platforms, Industry-specific applications.

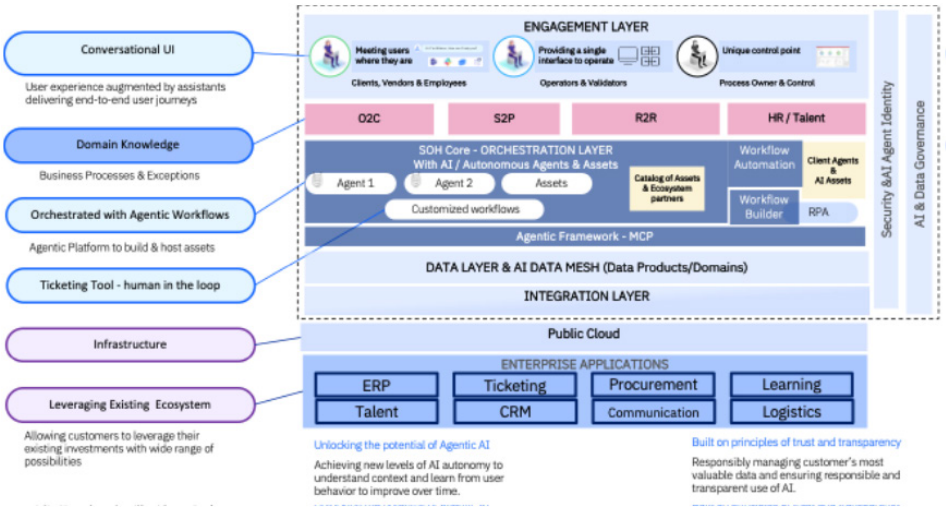

Figure 3: Smarter operations hub architecture (IBM, 2025).

To ensure secure operation, **identity propagation** must be maintained endtoend, where each agent action carries the initiating user identity. This enables granular authorization, zero-trust access, traceability, and aligned audit trails. Observability tooling captures every agent execution, prompt, decision, and system call

The **Smarter Operations Hub** is not just another integration platform-it is the **enterprise nervous system** for the AIdriven era. It unifies digital twins, agentic AI, enterprise workflows, and distributed data into a single, intelligent operational fabric. SOH provides the governance, orchestration, security, and interoperability required for safe, scalable automation across every business domain while ensuring consistent, explainable decision-making.

The recommended deployment for SOH is a hybridcloud architecture that balances vendor neutrality, operational efficiency, and security. In this model, the hub spans cloud and onprem environments, allowing organizations to run workloads where they perform best while still maintaining unified control. Key components of this deployment include:

- Kubernetes-based compute clusters for agent orchestration, workflow execution, and scalable service hosting.
- Managed AI and data services from major cloud providers, enabling flexible use of the best-in-class model hosting, vector search, and data processing capabilities.
- A unified Web Application Firewall (WAF) to secure all external access points and enforce consistent threat protection across environments.
- Secure internal communication using MCP (Model Context Protocol) to standardize agent-to-agent and agent-to-tool interactions, ensuring consistent governance, observability, and interoperability across the ecosystem.

This deployment approach ensures that SOH can evolve with new AI capabilities, scale across business functions, and operate securely across a heterogeneous technology landscape.

Case studies

Contract Leakage

Contract leakage (15–35%) often stems from inconsistent item definitions across ERP, sourcing, and supplier systems. Identical parts appear under multiple names, causing:

- Orphaned purchases (buyers can't find the right contract)
- Price drift (incorrect or noncontracted rates slip through)
- Supplier fragmentation

Using the SOH as a semantic intelligence layer, NLP/ML unifies all item descriptions into canonical golden records. This links purchases to contracts, validates pricing in real time, blocks noncompliant transactions, and supports tiered/indexbased contract structures, reducing leakage by 20–30% ($10-17M per $500M spend).

Talent

Probationary employees are managed in a thirdparty HR system (SAP/FiberGlass), while Company A runs Workday, Maximo, and ServiceNow. When probation ends, HR must manually transfer identity, performance, and employment data, an errorprone process with compliance risks. With the SOH, a secure crosscompany data transfer workflow is established:

- SOH pulls probation completion and performance data from Company B using scoped, timelimited access
- Enforces data minimization, mapping, and regulatory compliance
- Automatically harmonizes the data into Workday and Maximo
- Ensures all transfers are encrypted, logged, and stored in Microsoft Fabric for audit readiness.

This removes manual effort, improves accuracy, and ensures secure, compliant onboarding of probation workers.

CONCLUSION

This paper demonstrates that realizing the Society 5.0 vision requires architectures that are both technologically advanced and deeply humancentric. The proposed framework brings together multiagent systems, enterprise data fabrics, and secure orchestration to overcome fragmented systems, rising cognitive load, and inconsistent processes. By using agentic AI as secure, contextaware digital workers-with humanintheloop governance-the approach augments rather than replaces human judgment. The resulting hybrid sociotechnical system reduces cognitive burden, increases transparency, and enables people to shift toward highervalue oversight and decisionmaking roles.

ACKNOWLEDGMENT

The authors would like to acknowledge Jean-Stephane Payraudeau, IBM Managing Partner; Kyle Brown, IBM Fellow; Ndu Emuchay, IBM Fellow

REFERENCES

Balla, Ravi Teja (2025). "Enabling Cognitive Process Automation Using LLMs in ERP Systems with Generative AI," World Journal of Advanced Engineering Technology and Sciences, vol. 15, no. 3, pp. 1467-1474. DOI: 10.30574/wjaets.2025.15.3.1045

Devaraju, S. (2024). "AI-Powered HRM and Finance Information Systems for Workforce Optimization and Employee Engagement," Turkish Journal of Computer and Mathematics Education.

Eyo et al. (2025). "Agentic AI in SAP: Collaborative joule agents across procurement, finance and logistics," Global Journal of Engineering and Technology Advances, vol. 24, no. 2. DOI: 10.30574/gjeta.2025.24.2.0236

Galuzin, V., Galitskaya, A., Grachev, S., Larukhin, V., Novichkov, D., Skobelev, P. and Zhilyaev, A. (2022). "Autonomous Digital Twin of Enterprise: Method and Toolset for Knowledge-Based Multi-Agent Adaptive Management of Tasks and Resources in Real Time," Mathematics, vol. 10, no. 10, 1662. DOI: 10.3390/math10101662

Gopalaswamy, D. K. H. (2025). "AI and Human AI Collaboration in Oracle Cloud Technologies for Integration and Process Automation," European Journal of Computer Science and Information Technology, vol. 13, no. 8, pp. 107–128. DOI: 10.37745/ejcsit.2013/vol13n8107128 "Hybrid Machine Learning–Rule-based Decision Support Framework for Enhancing Operational Intelligence and Business Process Automation in ERP Systems" (2024). IEEE Xplore.

Koumpan, E., Topol, A. (2023). Value-driven architecture enabling new interaction models in Society 5.0. In: Tareq Ahram (eds) Human Factors in Software and Systems Engineering. AHFE (2023) International Conference. AHFE Open Access, vol 94. AHFE International, USA. http://doi.org/10.54941/ahfe1003765

Koumpan, E., Topol, A.W. (2021). Promoting Economic Development and Solving Societal Issues Within Connected Industries Ecosystems in Society 5.0. In: Ahram, T.Z., Karwowski, W., Kalra, J. (eds) Advances in Artificial Intelligence, Software and Systems Engineering. AHFE 2021. Lecture Notes in Networks and Systems, vol 271. Springer, Cham. https://doi.org/10.1007/978-3-030-80624-8_22

Roxanne et al. (2025). "Digital Twin Models for AI-Driven Fulfillment and Process Simulation."

Sridharan, V. (2025). "AI Agents & Digital Workers in ERP: Transforming Autonomous Decision-Making Across Industry Verticals," Journal of Information Systems Engineering and Management, vol. 10, no. 58s. DOI: 10.52783/jisem.v10i58s.12752

Stokes, G. (2025). "Enterprise Digital Twins in Financial Services: Convergence of Architecture, Operations, and Engineering," TechRxiv. DOI: 10.36227/techrxiv.175624549.98427022/v1

Szelagowski, M., Lupeikiene, A. and Berniak-Woźny, J. (2022). "Drivers and Evolution Paths of BPMS: State-of-the-Art and Future Research Directions," Informatica.

Zuo, A. (2025). "The Rise of Autonomous AI Agents: Automating Complex Tasks," International Journal of Artificial Intelligence for Science, vol. 1, no. 2. DOI: 10.63619/ijai4s.v1i2.007