Journal of Advances in Developmental Research (IJAIDR)

E-ISSN: 0976-4844 ● Website: www.ijaidr.com ● Email: editor@ijaidr.com

A Systematic Review of Generative Artificial Intelligence and Agentic Commerce in Digital Retail: Theoretical Foundations, Operational Implementations, and Autonomous Ecosystems

Prikshit Gupta[1], Akanksha Negi[2]

[1]Research Fellow, Himachal Pradesh University, Shimla
[2]ResearchFellow, Himachal Pradesh University, Shimla

Abstract:
The integration of Artificial Intelligence (AI) into digital commerce has historically centered on predictive analytics, classification algorithms, and rules-based recommendation engines. However, the advent of Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) has catalyzed a profound paradigm shift, transitioning the industry from human-mediated digital transactions toward autonomous, AI-mediated economic ecosystems. This comprehensive review synthesizes the rapidly expanding body of literature on GenAI applications in e-commerce, spanning empirical research published between 2020 and 2026. By systematically examining foundational theoretical frameworks—including the Technology Acceptance Model (TAM), the Unified Theory of Acceptance and Use of Technology (UTAUT), the Technology-Organization-Environment (TOE) framework, and the Stimulus-Organism-Response (S-O-R) paradigm—this paper elucidates the socio-technical drivers of AI adoption among global enterprises and consumers. Furthermore, the review critically analyzes the operational deployment of GenAI across intelligent demand forecasting, adaptive pricing mechanisms, sub-process workflow automation, and hyper-personalized marketing paradigms, notably examining the MARK-GEN framework (Garg et al., 2025). Crucially, this paper explores the nascent frontier of "Agentic Commerce," characterized by machine-to-machine (M2M) transactions, autonomous shopping agents, and emerging smart contract protocols. Through an exhaustive evaluation of the technological readiness, systemic security risks, algorithmic biases, and governance imperatives associated with these autonomous digital agents, this review charts future research directions, offering actionable strategic guidance for practitioners and rigorous theoretical grounding for scholars navigating the future of digital retail (Halat, 2026).

Keywords: AI, Generative Artificial Intelligence, Technology Acceptance Model, Technology-Organization-Environment & algorithmic biases

Introduction to the Generative Paradigm Shift

Since the introduction of transformative multimodal transformer models such as BERT and GPT-3, Generative Artificial Intelligence (GenAI) has emerged as the leading technological catalyst reshaping global industries, with electronic commerce experiencing one of the most profound ontological shifts in its history (Baig & Yadegaridehkordi, 2026). Unlike traditional, legacy AI systems designed primarily to analyze, classify existing data, detect anomalies, or forecast trends based purely on historical patterns, GenAI employs deep learning, complex neural networks, unsupervised learning, and advanced probabilistic modeling to autonomously create entirely novel and contextually relevant content (Shevchyk, 2024). This capability spans the generation of text, images, video, and audio, drastically reducing the marginal cost of content creation while enabling unprecedented levels of hyper-personalization in digital retail (Baig & Yadegaridehkordi, 2026).

The literature delineates a clear historical evolution of artificial intelligence in e-commerce, progressing through several distinct developmental phases. Early implementations in the 2000s relied on basic, rule-based recommendation engines that offered highly limited, static personalization based on rudimentary collaborative filtering (Baig & Yadegaridehkordi, 2026). The introduction of Generative Adversarial Networks (GANs) represented a critical inflection point, enabling sophisticated visual merchandising applications such as virtual try-on tools and dynamic product rendering (Goodfellow et al., 2014). Most recently, the widespread enterprise deployment of generative models like ChatGPT, DALL-E, and Midjourney has revolutionized content marketing, search engine optimization (SEO), and dynamic pricing strategies, allowing real-time adaptation to macroeconomic volatility (De Cremer et al., 2023).

However, the most disruptive evolution currently documented in the empirical literature is the transition toward "Agentic Commerce." Driven by the advanced reasoning capabilities of modern large language models, agentic commerce represents a transition from AI-assisted product discovery toward fully autonomous purchasing environments (Halat, 2026; Chainlink, 2024). In these ecosystems, AI agents operate as independent, active economic actors capable of executing multi-step workflows: searching for products, assessing supplier reliability, negotiating transaction terms, executing payments, coordinating logistics, and managing post-purchase dispute resolution without continuous human intervention (Birch & Gamble, 2025; Halat, 2026). Despite immense industry enthusiasm and rapid venture capitalization, the academic literature regarding these advancements remains heavily fragmented across the disparate disciplines of information systems, behavioral marketing, cybersecurity, digital governance, and algorithmic law. This paper addresses this fragmentation by offering an integrated, systematic synthesis of theoretical models, empirical applications, and the structural governance necessary for the large-scale, sustainable adoption of generative and agentic systems in e-commerce.

Review Design and Literature Identification Strategy

To ensure a rigorous and replicable synthesis of the fragmented literature surrounding GenAI and agentic systems, contemporary systematic reviews in this domain increasingly rely on standardized reporting guidelines. Notably, recent scoping reviews analyzing GenAI in digital retail have adopted the PRISMA-ScR (Preferred Reporting Items for Systematic Reviews and Meta-Analyses Extension for Scoping Reviews) framework, originally developed by Tricco et al. (2018) to enhance the quality and transparent literature synthesis (Tricco et al., 2018). Unlike traditional PRISMA frameworks constrained to highly specific clinical meta-analyses, PRISMA-ScR accommodates broader research questions and the heterogeneous inclusion criteria necessary for assessing rapidly evolving technological paradigms (Tricco et al., 2018).

Systematic investigations into agentic commerce and GenAI applications have predominantly targeted major academic databases, including Scopus, Web of Science, IEEE Xplore, ScienceDirect, SpringerLink, and the ACM Digital Library, supplemented by grey literature from Google Scholar and arXiv to capture cutting-edge preprints (Halat, 2026). The literature synthesis systematically categorizes evidence using frameworks such as the AI-Commerce Integration Assessment Framework (ACIAF), which interprets data across four distinct analytical dimensions: technological readiness, operational feasibility, system performance, and regulatory maturity (Halat, 2026). By leveraging these structured methodologies, the subsequent sections present an objective, bias-mitigated analysis of the theoretical and empirical realities of AI adoption in contemporary digital commerce.

Theoretical Frameworks Underpinning AI Adoption in Digital Retail

To comprehend the efficacy, organizational acceptance, and consumer diffusion of Generative AI within e-commerce, behavioral researchers and information systems scholars have adapted several foundational theoretical frameworks to fit the generative context. The complex interplay between consumer psychology, organizational resource readiness, and inherent technological capabilities dictates the ultimate success or failure of AI-generated content (AIGC) in competitive markets.

The Technology Acceptance Model (TAM) and UTAUT

The Technology Acceptance Model (TAM), initially proposed by Davis (1989), alongside the Unified Theory of Acceptance and Use of Technology (UTAUT) developed by Venkatesh et al. (2003), remain the preeminent frameworks for assessing both organizational and consumer receptivity to novel digital interfaces (Davis, 1989; Venkatesh et al., 2003). TAM postulates that the adoption rate of a new technology is fundamentally driven by two primary cognitive constructs: perceived usefulness (PU) and perceived ease of use (PEOU) (Davis, 1989). In the contemporary context of e-commerce, if a GenAI-driven conversational agent or an autonomous product recommendation system is highly intuitive (PEOU) and effectively reduces the consumer's cognitive load and search time (PU), technological acceptance will inevitably scale (IBM, 2026).

Recent empirical literature significantly extends traditional TAM by integrating contextual moderating variables specific to generative outputs, including perceived anthropomorphism, inherent trust, security risk evaluation, and data privacy concerns (Baig & Yadegaridehkordi, 2026). Empirical studies indicate that consumers are markedly more likely to accept and engage with AI-generated content if it reflects their personal values, exhibits a high degree of perceived authenticity, and originates from a historically trustworthy brand source (Shevchyk, 2024). Conversely, highly generic, irrelevant, or hallucinatory AI outputs generate immediate skepticism and behavioral resistance (Kumari & Gupta, 2025). An advanced derivative of these behavioral models is the Disclosure and Involvement Moderation Model, which quantitatively analyzes how the explicit disclosure of AI involvement (e.g., legally mandated watermarks labeling content as "Generated by AI") interacts with a consumer's level of product involvement to moderate downstream trust and final purchase intent (Baig & Yadegaridehkordi, 2026). Furthermore, the AI Device Use Acceptance (AIDUA) model underscores that hedonic motivation and anthropomorphism are primary predictors of a user's willingness to engage with tools like ChatGPT for complex commercial inquiries (Yadegaridehkordi et al., 2026).

The Technology-Organization-Environment (TOE) Framework

While TAM and UTAUT predominantly focus on the end-user's micro-level cognitive mechanisms, the

Technology-Organization-Environment (TOE) framework provides a macro-level, structural lens for examining the enterprise-level organizational adoption of Generative AI. A rigorous, multi-case qualitative study conducted among diverse e-commerce enterprises in Anhui Province, China, effectively bridged the TOE framework with TAM to demonstrate that enterprise adoption of AI is highly stratified based on precise organizational contexts (Zhu & Abd Rozan, 2026).

The TOE framework systematically deconstructs technological adoption into three interrelated dimensions:

1. **Technological Context:** The availability, interoperability, and architectural characteristics of AI tools, ranging from open-source foundational LLMs to expensive, proprietary enterprise models (Israfilzade, 2025).
2. **Organizational Context:** The internal operational capabilities, financial resource availability, firm size, and the baseline digital AI literacy of the existing workforce (Israfilzade, 2025).
3. **Environmental Context:** The external macro-pressures driving adoption, including hyper-competitive market dynamics, shifting consumer demands for instantaneous fulfillment, and complex regulatory policy guidance (Israfilzade, 2025).

Qualitative findings derived via the TOE framework reveal that while GenAI applications theoretically span the entire e-commerce value chain—including content generation, advertising placement, predictive data analysis, customer service routing, and logistics scheduling—a profound digital divide persists across the industry (Israfilzade, 2025). Small and micro-enterprises face severe limitations in technical depth and custom engineering capabilities, forcing them to rely on rigid "platform-bound" or highly generic tool-augmented models (Israfilzade, 2025). In contrast, large enterprises leverage massive capital to engineer self-developed, fine-tuned proprietary models (Israfilzade, 2025). Furthermore, the empirical evidence underscores the dominance of human-machine collaboration in current enterprise settings, describing the operational standard as an "AI pre-processing + human fine-tuning" model. Within this paradigm, an enterprise's internal organizational capability and baseline AI literacy serve as far more critical determinants of successful GenAI deployment than any external policy support or generic technological availability (Israfilzade, 2025).

S-O-R Theory and the Value-Satisfaction Paradigm

The Stimulus-Organism-Response (S-O-R) theory has been elegantly adapted to evaluate the profound impact of Generative AI on digital consumer engagement and the modern prompt-to-purchase journey (Giri & Biswas, 2025). Under this behavioral paradigm, GenAI-driven hyper-personalization, tailored conversational support, and dynamically generated aesthetic designs act as the environmental *stimulus*. This stimulus triggers complex cognitive and affective processing within the consumer (the *organism*), ultimately driving heightened brand loyalty, extended session durations, and definitive purchase intention (the *response*) (Giri & Biswas, 2025).

Complementing the S-O-R framework is the theoretical Value-Satisfaction Paradigm, which posits that Generative AI directly influences holistic customer satisfaction across four distinct, measurable dimensions of value:

- **Functional Value:** Evaluated through the measurable reduction of search friction, improved system reliability, and the provision of highly accurate, context-aware information retrieval (Baig & Yadegaridehkordi, 2026).

- **Emotional Value:** Created through sophisticated anthropomorphic design in conversational interfaces, establishing deep psychological trust and positive emotional affect during human-AI interactions (Israfilzade, 2025).
- **Promotional Value:** Achieved through the dynamic, real-time personalization of targeted advertising and customized discount generation (Baig & Yadegaridehkordi, 2026).
- **Innovative Value:** Recognized as highly critical for distinguishing digital retail offerings and maintaining a competitive market advantage in an increasingly saturated digital landscape (Baig & Yadegaridehkordi, 2026).

Theoretical Framework	Primary Analytical Focus	Key Predictive Factors & Operational Mechanisms in E-Commerce
TAM / UTAUT	Micro-Level End-User Adoption	Perceived Usefulness, Perceived Ease of Use, Trust, Hedonic Motivation, Anthropomorphism, Privacy Concerns (Davis, 1989).
TOE Framework	Macro-Level Enterprise Adoption	Technological readiness, firm size, internal AI literacy, market pressure, platform architecture (Zhu & Abd Rozan, 2026).
S-O-R Theory	Consumer Behavioral Psychology	Stimulus (AI Content) Organism (Cognitive/Affective Engagement) Response (Purchase) (Giri & Biswas, 2025).
Value-Satisfaction	Holistic Customer Experience	Functional, Emotional, Promotional, and Innovative Value derivation (Baig & Yadegaridehkordi, 2026).

Operational Transformations: Workflow Automation and Intelligent Optimization

Systematic literature reviews mapping GenAI applications identify several major vectors of operational transformation within retail environments: enhancing customer purchase intention, driving backend optimization strategies, supporting executive decision-making, improving semantic information retrieval, generating programmatic advertising, and refining technical implementation pipelines (Baig &

Yadegaridehkordi, 2026).

Sub-Process Mapping and E-Commerce Workflow Automation

A critical, pragmatic realization in contemporary e-commerce literature is that broad conceptualizations such as "AI for e-commerce" are strategically insufficient; they are too ambiguous to be effectively governed, engineered, or measured for return on investment (ROI) (LeewayHertz, 2024). The true enterprise value of Generative AI is unlocked exclusively through rigorous sub-process mapping, wherein AI capabilities are embedded deeply into specific, repeatable workflows tied directly to existing systems of record (e.g., ERP and CRM platforms) (LeewayHertz, 2024).

Traditional machine learning algorithms were largely confined to predicting, scoring, or classifying data based strictly on historical patterns. Conversely, modern GenAI models can ingest, summarize, draft, compare, synthesize, and transform unstructured information into highly usable operational formats (LeewayHertz, 2024). This paradigm shift is particularly transformative for both document-heavy and narrative-heavy retail workflows. Within the **Product and Catalog Operations** function, for instance, GenAI engines can automatically ingest raw, unformatted manufacturer data sheets, supplier invoices, and rigid Stock Keeping Unit (SKU) attribute schemas to instantly draft hundreds of unique, SEO-optimized product detail pages (PDPs) that adhere to specific brand voice guidelines (LeewayHertz, 2024).

Similarly, in **Customer Service Operations**, rather than relying on frustrating, decision-tree chatbots that alienate consumers, sophisticated GenAI models can instantly ingest a specific customer's entire CRM history—including lifetime order records, return merchandise authorizations (RMAs), and previous sentiment-scored service interactions. The model then synthesizes this data to generate a concise summary for a human customer service specialist and pre-drafts a highly personalized, empathetic resolution email (LeewayHertz, 2024). This specific architecture adheres strictly to the "human-on-the-loop" governance philosophy. The AI operates as an advanced intelligence assistant—retrieving evidentiary data, highlighting potential fraud risks, and drafting the final output—but a named, accountable human reviewer must explicitly approve any customer-facing message or risk-bearing financial action before systemic execution (LeewayHertz, 2024).

Intelligent Demand Forecasting and Adaptive Dynamic Pricing

Generative AI significantly augments intelligent demand forecasting and adaptive pricing models, providing a definitive, mathematically measurable competitive edge in volatile retail sectors (Krishnakumar, 2022). Traditional, linear forecasting models historically struggle with sudden macroeconomic volatility, localized supply chain disruptions, and black-swan market fluctuations because they rely excessively on historical time-series data. GenAI, particularly when hybridized with advanced reinforcement learning mechanisms, excels in navigating these multidimensional data spaces (Krishnakumar, 2022).

By synthesizing vast quantities of unstructured, real-time environmental data—including global macroeconomic news sentiment, localized social media velocity, geographic weather patterns, and real-time competitor pricing files—GenAI enables retailers to adjust prices dynamically and predict demand at hyper-granular scales (Krishnakumar, 2022). This continuous, real-time algorithmic optimization drastically improves inventory management, minimizes costly warehouse holding times, prevents stockouts of high-velocity items, and maximizes revenue realization. However, scholars caution that these adaptive pricing models must be strictly governed by transparent ethical parameters to prevent algorithmic price gouging, which can rapidly erode consumer trust and trigger regulatory intervention (Zhu & Abd Rozan, 2026).

Innovations in Digital Marketing: Hyper-Personalization and the MARK-GEN Framework

The discipline of digital marketing is arguably the domain most visibly and immediately impacted by the advent of GenAI, undergoing a rapid transition from human-intensive, slow-cycle creative processes to automated, high-velocity content generation (LeewayHertz, 2024). To systematize this transition, researchers have proposed structural methodologies such as the MARK-GEN framework. This framework is particularly salient for businesses operating in highly diverse, multilingual emerging markets, providing a strategic blueprint to harness GenAI across advertising copywriting, hyper-personalization, and vernacular search engine optimization (Garg et al., 2025).

Generative AI fundamentally alters retail marketing through several sophisticated mechanisms:

- **Hyper-Personalization at Unprecedented Scale:** Utilizing advanced LLM tools (e.g., Jasper AI, Cohere, OpenAI), marketers can auto-generate thousands of A/B test variations tailored to the precise behavioral and psychographic insights of micro-segments. This includes dynamically inserting localized cultural references, individualized behavioral triggers, and context-specific promotional offers into email subject lines and social media reels in real time (Garg et al., 2025).
- **Synthetic Influencers and Advanced Visual Media:** The rapid rise of AI-generated avatars and purely synthetic virtual influencers (such as the globally recognized Lil Miquela) allows digital brands to maintain absolute, risk-free control over their brand ambassadors (Garg et al., 2025). This eliminates the severe reputational risks, behavioral unpredictability, and exorbitant financial costs traditionally associated with human influencers (Garg et al., 2025). Furthermore, fashion retailers are aggressively deploying multimodal transformers and diffusion models to enable photorealistic virtual try-ons, drastically enhancing both the aesthetic and functional value of the digital shopping experience (Garg et al., 2025).
- **Vernacular SEO and Deep Localization:** In complex, multi-linguistic emerging markets (e.g., India), GenAI is heavily utilized to auto-generate long-tail keywords, localized blog posts, and voice-search-friendly FAQs in regional dialects (e.g., Hindi, Tamil, Bengali). This capability drastically reduces customer acquisition costs (CAC) in Tier-2 and Tier-3 cities, enabling SMEs to achieve multilingual reach that was previously cost-prohibitive (Garg et al., 2025).

Despite these immense operational advantages, qualitative interviews conducted with senior marketing executives reveal a cautious, nuanced approach to strategic implementation. While GenAI is universally praised for backend operational execution and high-volume content scaling, many senior strategists express profound hesitation regarding its use in fundamental brand decision-making and high-level campaign ideation. Executives cite the severe risks of brand genericization, the erosion of authentic "human touch," potential copyright infringement, and the jeopardy of long-term brand equity if consumers perceive the brand as artificially sterile (Aw et al., 2025).

The Paradigm Shift Toward Agentic Commerce and Autonomous Ecosystems

The most novel, highly disruptive, and intellectually fertile horizon identified in the recent literature spanning 2025 to 2026 is the definitive transition from GenAI-assisted commerce to "Agentic Commerce."

Agentic commerce is formally defined as a new economic paradigm wherein autonomous, highly context-aware AI agents negotiate, execute, and manage complex commercial transactions directly on behalf of human users or organizational entities without requiring synchronous human input (Chainlink, 2024; Halat, 2026).

Defining the Agentic Ecosystem and Autonomous Purchasing

In traditional, legacy e-commerce paradigms, a human user is inherently constrained by the manual necessity of executing localized search queries, evaluating highly fragmented retail catalogs, cross-referencing prices across multiple browser tabs, inputting shipping logistics, and navigating friction-heavy, multi-step checkout gateways (Chainlink, 2024). Agentic commerce entirely removes the human from this immediate, manual execution loop. Under this model, a consumer provides a high-level intent or semantic goal—for example, a prompt stating, "Monitor my household cleaning supplies and autonomously restock them when inventory is low, ensuring the total cost remains under $50 with next-day delivery"—and the autonomous agent programmatically handles the entirety of the subsequent lifecycle (Airwallex, 2024; J.P. Morgan, 2024).

The agentic transaction lifecycle involves the AI simultaneously querying dozens of disparate marketplaces in milliseconds, pulling raw product specifications via API, calculating real-time landed costs (strictly inclusive of localized taxes, dynamic shipping rates, and user-specific loyalty discounts), and executing the final transaction securely utilizing tokenized biometric credentials (Airwallex, 2024). Crucially, because AI agents are algorithmic constructs devoid of emotion, they are entirely immune to traditional marketing psychological triggers, visual FOMO (Fear Of Missing Out), artificial scarcity tactics, and impulse buying cues (Airwallex, 2024). Consequently, structural impulse sales are projected to decline significantly as agentic commerce scales. However, retailers offset this loss through access to highly qualified, deterministic traffic and near-perfect transaction success rates, achieved by bypassing the massive cart abandonment rates that plague human-centric checkout funnels (Airwallex, 2024).

Industry analysts and technical researchers categorize the maturity of these agentic systems into four distinct evolutionary tiers:

- **Level 1 (Information Retrieval & Discovery):** Agents act as advanced, conversational search engines, retrieving, synthesizing, and summarizing product choices based on complex natural language queries.
- **Level 2 (Workflow Automation):** Agents autonomously navigate a single merchant's digital infrastructure, configuring complex products and populating a cart, but stopping to wait for final human authorization.
- **Level 3 (Autonomous Checkout):** Agents are granted financial autonomy to securely execute the transaction on a pre-trusted site using securely stored payment tokens, requiring no human intervention.
- **Level 4 (Multi-Agent Orchestration & Resolution):** A single meta-agent splits highly complex orders across multiple unaffiliated merchants, dynamically coordinates logistics with independent delivery carriers, and autonomously resolves shipping disputes, returns, or refunds in real time (Airwallex, 2024).

Maturity Level	Agent Capability	Human Intervention Requirement	Primary Technological Mechanism
Level 1	Retrieval & Discovery	High (Evaluates and selects options)	RAG (Retrieval-Augmented Generation), LLM Summarization (Airwallex, 2024)
Level 2	Workflow Automation	Moderate (Final cart approval required)	Web-crawling, DOM manipulation (Airwallex, 2024)
Level 3	Autonomous Checkout	Low (Pre-authorized budget limits)	Tokenized credentials, Secure APIs (Airwallex, 2024)
Level 4	Multi-Agent Orchestration	None (Full autonomous lifecycle)	M2M Protocols, Smart Contracts, Meta-Agent Routing (Airwallex, 2024)

Machine-to-Machine (M2M) Transactions: "Robots Paying Robots"

The fundamental foundation of true agentic commerce necessitates a complete, structural re-architecture of global digital payment infrastructure to facilitate high-velocity Machine-to-Machine (M2M) transactions, colloquially termed in the literature as "robots paying robots" (Birch & Gamble, 2025). Legacy human-centric payment gateways—which rely on CAPTCHAs, SMS two-factor authentication (2FA), and manual credit card data entry—are fundamentally incompatible with programmatic, microsecond agentic workflows.

To bridge this infrastructural gap, a new foundational layer of open, cryptographic protocols is rapidly emerging. Systems such as the Agentic Commerce Protocol (ACP), Anthropic's Model Context Protocol (MCP), and the Agent Payments Protocol (AP2) are being engineered to allow software to securely hold financial value, cryptographically sign transactions, and interact seamlessly with decentralized smart contracts (Goodfellow et al., 2014). In these environments, "smart wallets"—digital wallets operated entirely by robotic agents—act as the central financial orchestration mechanism. This architecture enables instantaneous, frictionless settlement and is projected to revive historically unviable economic models, such as fractional micropayments for API calls and highly dynamic, real-time supply-chain currencies (Birch & Gamble, 2025; Infosys, 2024).

Agentic network configurations can take multiple structural forms:

- **Agent-to-Site:** A consumer's personal AI agent interacts programmatically with a traditional retailer's existing, rigid web infrastructure (Infosys, 2024).
- **Agent-to-Agent (A2A):** An autonomous buyer agent negotiates pricing, volume discounts, terms, and delivery schedules directly with an autonomous seller agent representing the merchant. A prime example includes a personal healthcare agent collaborating seamlessly with a pharmacy's AI to process complex insurance claims and reorder critical prescriptions without any human friction (Halat, 2026).

Generative Engine Optimization (GEO) and Catalog Readiness

The rapid rise of the autonomous agentic buyer fundamentally alters the disciplines of digital marketing and search engine optimization. If the primary economic purchaser is no longer a human reading emotional copy and viewing aesthetic imagery, but rather an LLM-powered reasoning engine parsing JSON files and API endpoints, merchants must urgently shift their strategy from traditional SEO to Generative Engine Optimization (GEO) (Airwallex, 2024).

GEO prioritizes rendering product catalogs entirely machine-readable through strictly enforced Schema.org JSON-LD markup and structured metadata. Dense, factual specifications—such as precise Global Trade Item Numbers (GTINs), exact dimensional data, compliance certifications, and live API-driven inventory counts—will heavily outperform creative, descriptive marketing prose when AI agents evaluate products for a human user (Airwallex, 2024). Merchants who fail to expose clean, structured catalog APIs will simply be rendered invisible, systematically bypassed by agentic discovery engines in favor of technically optimized competitors (Airwallex, 2024). Furthermore, researchers utilizing the AI-Commerce Integration Assessment Framework (ACIAF) have noted that a successful enterprise transition to agentic commerce hinges critically on four dimensions: readiness (structured data), feasibility (payment APIs), performance (latency and transaction success rates), and maturity (regulatory and legal compliance) (Halat, 2026).

Ethical Governance, Systemic Security, and Adoption Barriers

Despite the extraordinary operational potential of GenAI and the transformative promise of agentic commerce, profound structural adoption barriers persist. These challenges are not merely technical hurdles solvable by greater computational power; rather, they encompass highly complex legal, ethical, and sociological dimensions that demand rigorous governance frameworks.

Algorithmic Bias, Consumer Trust, and Data Privacy

A pervasive and critical challenge in the deployment of GenAI is the ongoing mitigation of algorithmic bias and the preservation of fragile consumer trust. Large Language Models are trained on massive, historically scraped datasets that inherently contain human prejudices. Consequently, AI-generated marketing content, autonomous product recommendations, and dynamic pricing algorithms can unintentionally manifest and amplify racial, gender, or socioeconomic biases, leading to severe reputational damage and consumer alienation (Baig & Yadegaridehkordi, 2026).

Furthermore, data privacy remains an acute friction point. E-commerce platforms leveraging GenAI require massive, continuous data ingestion to achieve hyper-personalization, putting these systems in direct, inevitable tension with stringent global regulatory frameworks, most notably the European Union's General Data Protection Regulation (GDPR) and India's Digital Personal Data Protection (DPDP) Act of

2023 (Baig & Yadegaridehkordi, 2026). Transparency regarding exactly how consumer behavioral data is utilized, stored, and processed to train generative models is paramount. Empirical evidence demonstrates that a failure to explicitly disclose AI data practices routinely leads to a psychological phenomenon termed "online shopping fatigue," characterized by heightened consumer inertia, choice overload, and a definitive rejection of digital retail platforms (Birch & Gamble, 2025).

Systemic Security Vectors in Autonomous Agentic Ecosystems
As digital commerce structurally shifts toward fully autonomous Agent-to-Agent (A2A) ecosystems, the cybersecurity attack surface broadens exponentially. Academic research mapping the security framework for autonomous LLM agents in commerce and finance identifies several severe, cross-layer attack vectors resulting directly from the unprecedented delegation of financial authority to machine actors (Israfilzade, 2025).

Key systemic security and governance threats include:

1. **Agent Integrity and Prompt Injection:** Malicious actors can theoretically manipulate an agent's foundational reasoning layer through sophisticated prompt injection attacks hidden within seemingly benign product descriptions or review text. This could trick a buyer's agent into executing unauthorized purchases, exfiltrating personal data, or authorizing premium pricing against the user's intent (Israfilzade, 2025).
2. **Transaction Authorization and Smart Contract Vulnerabilities:** Because agentic commerce relies heavily on tokenized cryptographic credentials and immutable smart contracts for execution, any underlying code vulnerability in the cryptographic architecture can result in massive, rapid financial exposure without the possibility of a traditional bank chargeback (Birch & Gamble, 2025).
3. **Market Manipulation and Algorithmic Collusion:** Swarms of autonomous agents could theoretically engage in algorithmic collusion at millisecond speeds, artificially inflating demand signals, executing wash trades, or manipulating dynamic pricing models to the severe detriment of human consumers (Israfilzade, 2025).
4. **Legal Liability and the Definition of Consent:** The ontological shift to autonomous shopping poses a profound, currently unresolved legal question in digital jurisprudence: What constitutes legal consumer consent in a "human-not-present" transaction? If an AI agent hallucinates or misinterprets a natural language prompt (e.g., ordering 1,000 units of a product instead of 10), the legal and financial liability between the human consumer, the LLM developer, the retail merchant, and the payment gateway remains entirely ambiguous (Halat, 2026; J.P. Morgan, 2024).

Consequently, scholars emphasize that the large-scale, mainstream adoption of agentic commerce will be gated not by the raw technical capability of the AI models, but by the arduous establishment of robust governance conditions. This requires verifiable digital identity infrastructures, universally interoperable dispute resolution mechanisms, and transparent, Explainable AI (XAI) frameworks (Halat, 2026).

Future Research Directions
Based on the exhaustive synthesis of contemporary literature spanning 2020 to 2026, several urgent, high-impact pathways for future academic inquiry emerge:

1. **Explainable AI (XAI) and Algorithmic Interpretability:** To alleviate ethical concerns and

satisfy impending regulatory requirements, future empirical studies must focus on engineering transparent neural networks and feature attribution models. These models must clearly delineate the exact algorithmic pathways GenAI utilizes to arrive at specific dynamic pricing decisions or highly personalized product recommendations (Halat, 2026).

2. **Integration of Reinforcement Learning in Volatile Markets:** Researchers must actively explore the hybridization of Generative AI with advanced reinforcement learning in highly volatile, real-time e-commerce environments. Such studies are necessary to stress-test the scalability, resilience, and mathematical responsiveness of autonomous inventory optimization and adaptive pricing models during macroeconomic shocks (Baig & Yadegaridehkordi, 2026).
3. **Privacy-Preserving AI Architectures:** Future empirical studies should rigorously evaluate the commercial viability, computational overhead, and latency impacts of utilizing decentralized federated learning and differential privacy within e-commerce LLMs. These architectures are critical to securing consumer data while maintaining the high efficacy of hyper-personalization (Baig & Yadegaridehkordi, 2026).
4. **Immersive Commerce and the Multimodal Metaverse:** The technical intersection of GenAI with Augmented Reality (AR), Virtual Reality (VR), and extended reality (XR) warrants significantly deeper investigation. Researchers should quantify how AI-generated, immersive 3D retail environments impact consumer user experience (UX), functional value derivation, and long-term aesthetic brand perception (Baig & Yadegaridehkordi, 2026).
5. **Agentic Interoperability and Legal Frameworks:** There is an acute, immediate need for policy-oriented legal research focusing on the standardization of Machine-to-Machine (M2M) payment protocols, the precise legal definitions of algorithmic agency and liability, and the development of cost-effective, lightweight GenAI frameworks to ensure equitable technological access for Small and Medium-Sized Enterprises (SMEs) (Halat, 2026).

Conclusion

The integration of Generative Artificial Intelligence into electronic commerce represents an evolutionary leap from human-directed, rule-based systems to highly adaptive, cognitively complex digital ecosystems. By synthesizing foundational theoretical frameworks such as TAM, TOE, and S-O-R, it becomes empirically evident that the successful deployment of GenAI hinges on meticulously balancing technological capability with organizational readiness, consumer trust, and perceived utility. Operationally, GenAI is currently revolutionizing the backend of digital retail by autonomously executing narrative-heavy sub-processes, enabling real-time adaptive pricing, and facilitating hyper-personalized marketing at unprecedented global scales.

Looking forward, the digital retail industry stands on the precipice of "Agentic Commerce." The impending transition toward autonomous AI shopping agents and highly secure, machine-to-machine payment infrastructures promises to radically compress digital sales funnels and functionally eliminate traditional human purchasing friction. However, this future necessitates a fundamental, structural re-architecture of commercial data—shifting away from human-readable SEO toward machine-readable Generative Engine Optimization (GEO). Furthermore, it demands the urgent deployment of rigorous, trustless governance protocols to mitigate systemic security risks, prompt injection vulnerabilities, and algorithmic bias. Ultimately, the sustainable, long-term development of AI-driven autonomous commerce will rely entirely on the coordinated evolution of robust technical standards, explainable algorithmic transparency, and ethical legal frameworks designed specifically to foster enduring consumer trust in an increasingly automated global marketplace.

REFERENCES:

1. Airwallex. (2024). *Agentic commerce*. Airwallex Blog. https://www.airwallex.com/en-us/blog/agentic-commerce
2. Aw, E. C. X., Anjomshoae, A., Blome, C., Cham, T. H., Chan, H. K., Eastman, J., Kar, A. K., Ki, C. W., Leong, L. Y., Mandal, S., Mishra, A., Ooi, K. B., Raman, R., Rana, N. P., & Tan, G. W.-H. (2025). Future research directions of generative AI in marketing and logistics. *Asia Pacific Journal of Marketing and Logistics*, 1–15. https://doi.org/10.1108/APJML-11-2024-1605
3. Baig, M. I., & Yadegaridehkordi, E. (2026). Generative AI in e-commerce: A literature review of current applications, gaps and future directions. *Asia-Pacific Journal of Business Administration*, 1–29. https://doi.org/10.1108/APJBA-08-2025-0575
4. Birch, D. G. W., & Gamble, D. (2025). Agentic commerce and payments: Exploring the implications of robots paying robots. *Journal of Payments Strategy & Systems, 19(1)*, 72–84. https://doi.org/10.69554/NGEA2302
5. Chainlink. (2024). *What is agentic commerce?* https://chain.link/article/agentic-commerce
6. Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly, 13(3)*, 319–340. https://doi.org/10.2307/249008
7. De Cremer, D., Bianzino, N. M., & Falk, B. (2023). How generative AI could disrupt creative work. *Harvard Business Review*. https://hbr.org/2023/04/how-generative-ai-could-disrupt-creative-work
8. Garg, A., Kaushik, R., Singh, H., Srivastava, J., Sinha, A., & Kumar, A. (2025). Innovating marketing strategies with generative AI: The MARK-GEN framework and its implementation. *Journal of Marketing & Social Research, 2(7)*, 97–115. https://doi.org/10.61336/jmsr/25-07-12
9. Giri, A., & Biswas, W. (2025). Prompt-to-Purchase: Leveraging generative AI to revolutionize E-Tail (Electronic–Retail) customer engagement. *Indian Journal of Marketing, 55(12)*, 59–70. https://doi.org/10.17010/ijom/2025/v55/i12/175246
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* (Vol. 27). https://papers.nips.cc/paper_files/paper/2014/hash/5ca3e9b122f61f8f3f6b2c8f5f4c5b72-Abstract.html
11. Halat, O. (2026). Agentic commerce: A systematic review of AI-driven autonomous shopping and emerging transaction protocols. *Global Prosperity, 6(3)*. https://gprosperity.org/index.php/journal/article/view/325
12. IBM. (2026). *What is agentic commerce?* https://www.ibm.com/think/topics/agentic-commerce
13. Infosys. (2024). *Agentic commerce: The future of autonomous shopping*. Infosys Blogs. https://blogs.infosys.com/emerging-technology-solutions/digital-experience/agentic-commerce-the-future-of-autonomous-shopping.html
14. Israfilzade, K. (2025). The role of generative artificial intelligence in e-commerce: Trends,

challenges and opportunities. *The Eurasia Proceedings of Educational and Social Sciences, 42*, 1–20. https://doi.org/10.55549/epess.906

15. J.P. Morgan. (2024). *Agentic commerce: How AI could change the way you shop*. J.P. Morgan Payments. https://www.jpmorgan.com/payments/newsroom/agentic-commerce-ai-future-shopping
16. Krishnakumar, A. (2022). Transforming e-commerce with generative AI: Toward intelligent demand forecasting and adaptive pricing. *International Journal of Scientific Research in Science and Technology, 9(5)*, 753–768. https://doi.org/10.32628/IJSRST2295157
17. Kumari, N., & Gupta, K. (2025). Consumer acceptance of AI generated content in e commerce. *International Journal for Multidisciplinary Research, 7(2)*. https://doi.org/10.36948/ijfmr.2025.v07i02.39569
18. LeewayHertz. (2024). *Generative AI in e-commerce: Mapping operating-model opportunities*. https://www.leewayhertz.com/generative-ai-in-ecommerce/
19. Shevchyk, Y. (2024). *Generative AI in e-commerce: A comparative analysis of consumer-brand engagement with AI-generated and human-generated content* [Doctoral dissertation, Vytautas Magnus University]. https://www.vdu.lt/cris/entities/etd/ae838c57-922d-45bd-90e8-91812d4ccd99
20. Tricco, A. C., Lillie, E., Zarin, W., O'Brien, K. K., Colquhoun, H., Levac, D., Moher, D., Peters, M. D. J., Horsley, T., Weeks, L., Hempel, S., Akl, E. A., Chang, C., McGowan, J., Stewart, L., Hartling, L., Aldcroft, A., Wilson, M. G., Garritty, C., ... Straus, S. E. (2018). PRISMA extension for scoping reviews (PRISMA-ScR): Checklist and explanation. *Annals of Internal Medicine, 169(7)*, 467–473. https://doi.org/10.7326/M18-0850
21. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly, 27(3)*, 425–478. https://doi.org/10.2307/30036540
22. Yadegaridehkordi, E., Foroughi, B., & Ghobakhloo, M. (2026). Factors affecting academic staff's willingness to use ChatGPT for teaching and learning: A PLS-SEM and ANN approach. *Innovative Higher Education, 51*, 1–28. https://doi.org/10.1007/s10755-025-09835-8
23. Zhu, T., & Abd Rozan, M. Z. (2026). AI adoption in e-commerce enterprises: Insights into current practices and future directions from an interview study. *PLOS ONE, 21(3)*, e0336416. https://doi.org/10.1371/journal.pone.0336416