

The (Agent-Orchestrated) Customer Journey: How AI Agent Optimization is Revolutionizing Branding

Steliana VASILEVA*

D. A. Tsenov Academy of Economics, Svishtov, Bulgaria

*Corresponding author, s.vasileva@uni-svishtov.bg

Abstract. Artificial intelligence (AI) agents (software entities that can perceive, decide, and act with minimal human oversight) are reshaping how consumers discover, evaluate, purchase, and use brands across digital channels. In online retail, these agents coordinate search and curation, personalize content, and accelerate transactions, thereby redistributing influence along the customer journey from retailers toward brands and agentic intermediaries. This paper synthesizes current evidence on AI agents' effects across pre-purchase, purchase, and post-purchase stages, highlighting: (a) hyper-personalized discovery and conversational assistance that compress consideration sets; (b) agent-enabled transaction support (e.g., dynamic pricing, streamlined checkout, real-time help) that lowers friction; and (c) post-purchase orchestration (recommendations, returns, and service) that shapes satisfaction, loyalty, and lifetime value. The paper interprets these developments against recent shifts in the information ecosystem, such as the introduction of AI Overviews in Google Search and the integration of community content via data-access agreements – which alter how brand information is surfaced and cited to consumers. Building on emerging work in generative/agentic search optimization, the paper argues that brands must complement traditional Search Engine Optimization (SEO) with AI-agent optimization (AAO) strategies (e.g., entity hygiene, structured data, citation seeding, and conversational answerability) to remain machine-retrievable in generative results. The paper also notes cross-model variation in what sources AI systems cite, implying fragmented visibility requirements for brands. Conceptually, the paper integrates insights from customer-journey research with recent AI-in-marketing evidence to propose a stage-wise framework and a research agenda focused on trust, transparency, and long-term brand equity effects.

Keywords: AI agents, customer journey, online retailing, branding, customer experience, AI-agent optimization.

Introduction

Digital customer journeys are undergoing a structural shift as answer-first interfaces – large-scale generative search experiences and agentic shopping assistants – mediate how consumers discover, evaluate, and purchase. In May 2024 Google mainstreamed AI Overviews in Search and expanded globally thereafter, placing model-generated answers above links and reframing what counts as visibility for brands. In parallel, Amazon and category platforms introduced embedded, conversational shopping agents (for example, Amazon's Rufus and Instacart's Ask Instacart), further compressing search, evaluation, and action into a single, model-mediated touchpoint. These moves reconfigure the practical “unit of competition” from pages and keywords to answers and citations within opaque model outputs.

While platforms report favorable engagement, early audits and news coverage document non-trivial accuracy variance and source-selection issues in generative answers, highlighting brand risks associated with misattribution and hallucination. Google publicly acknowledged post-launch fixes after high-profile errors; independent assessments continue to show uneven performance, with studies and industry analyses describing omission, consensus bias, and error rates across AI search services. These dynamics suggest that discoverability and trust now depend

not only on traditional ranking factors but also on how model pipelines sample, summarize, and cite brand-relevant information.

Classic customer-journey frameworks remain foundational but require reinterpretation under model mediation. The McKinsey consumer decision journey reconceived purchase as an iterative loop; subsequent work on omnichannel and retail experiences specified touchpoints and capability requirements. Google’s “messy middle” further described non-linear oscillation between exploration and evaluation. In a world where large language models (LLMs) synthesize answers atop SERPs or within retailer apps, this loop and “middle” are increasingly executed by agents, not users, thereby shifting where brands must intervene (Court et al., 2009; Grewal & Roggeveen, 2020).

Emerging technical work formalizes the optimization problem created by answer-first interfaces. The Generative Engine Optimization (GEO) literature shows that content style, structure, and evidence cues can measurably influence whether generative engines cite a source in their answers, implying a new optimization regime adjacent to – but distinct from – Search Engine Optimization (SEO). This reinforces a broader managerial point: as engines fuse retrieval and generation, firms must optimize model selection and citation rather than only for page ranking (Aggarwal et al., 2024).

To organize a brand response, AI-agent optimization (AAO) is defined as the cross-functional capability set that makes a brand correctly findable, faithfully represented, and action-ready inside model-mediated experiences across the customer journey. Specifically, AAO spans (a) content and evidence design for generative engines (e.g., claim-level citations and structured facts), (b) technical signals and interfaces (e.g., data feeds, schema, product knowledge accessible to agents), (c) measurement oriented to share-of-answer and citation recall rather than page rank, and (d) governance aligned with emerging safety and compliance regimes. AAO complements and extends extant digital-marketing and customer experience (CX) capability frameworks by targeting the selection and reasoning layers of AI systems, not only the interface or the channel (Aggarwal et al., 2024).

This reconceptualization is timely for at least three reasons. First, agentic surfaces are now mainstream retail entry points (e.g., Rufus and Ask Instacart), compressing the path from need recognition to action. Second, the mixture of platform claims, and third-party audits indicates that model behavior can be volatile and uneven, raising the stakes for brand accuracy and safety. Third, regulatory baselines (e.g., the EU AI Act) are crystallizing obligations around transparency and risk management, increasing the need for explicit governance over how brand data are exposed to and reused by AI systems.

Although scholars have mapped digital customer journeys and reviewed AI’s role in marketing and CX, the literature has only recently begun to formalize optimization for generative engines and has not yet articulated an integrated, managerially actionable capability for brand-safe participation in AI answers. Addressing these research gaps, the paper answers the following questions:

RQ1. How do model-mediated, answer-first experiences restructure consumer attention, trust, and choice mechanisms across journey stages?

RQ2. What specific capabilities constitute AAO for brands (content/evidence design, technical signals, measurement, and governance), and how do they operationalize visibility and fidelity in AI answers?

RQ3. How can firms measure and manage outcomes such as share-of- answer, citation recall, and downstream conversion in environments where clicks and referrals are partially disintermediated by agents?

Specifically, *the paper is structured as follows*: The first section decodes the concept of the contemporary digital customer journey. This section integrates the scholarly principles that link the increasing significance of AI (especially agentic systems) in branding and customer experience, and the rising necessity for AI-agent optimization (AAO) as a supplement to (rather than a substitute for) SEO. The section also enhances and refines the framework presented in the paper while ensuring that the argumentation is consistent with recent peer-reviewed studies. The research methodology is organized in accordance with the strict PRISMA guidelines, which are designed to guarantee transparency, thoroughness, and precision in the synthesis of existing literature. The paper presents a structured narrative review that integrates peer-reviewed studies from marketing, information systems, human-computer interaction, and retailing to develop an integrative conceptual framework (AOJ) alongside the capability construct of AI-agent optimization (AAO). The next section seeks to emphasize cross-stage regularities from the literature, emphasizing capabilities, outcomes and brand task for each of the customer journey's stages. The paper also proposes useful metrics for AI-mediated brand visibility. The practical instrumentation and the managerial implications are also considered. Finally, the paper takes into account the brand-specific limitations and suggests some agenda for future research with an emphasis on brand impact. The concluding section highlights key insights obtained from the research and discusses the contributions made in this paper.

Literature review

This section synthesizes the academic foundations that connect (a) the modern digital customer journey, (b) the accelerating role of AI (particularly agentic systems) in branding and customer experience, and (c) the emerging need for AI-agent optimization (AAO) as a complement to (not a replacement for) SEO. It extends and tightens the framing introduced in the manuscript while aligning the argumentation with recent peer-reviewed research.

Customer-Journey Foundations in a Digital, Omnichannel Context

Marketing research has moved from funnel metaphors toward journey-centric views that map dynamic, cross-channel episodes of cognition, affect, and behavior over time (e.g., touchpoints before, during, and after purchase). Lemon and Verhoef's (2016) synthesis remain an anchor, formalizing customer experience spanning multiple media and organizational functions. Later work in Journal of Retailing stresses customer-journey management as a strategic capability in technology-intensive retail, linking journey design to performance outcomes (Grewal & Roggeveen, 2020). At the method level, Rosenbaum et al. (2017) clarify how journey mapping should be grounded in customer research rather than managerial assumptions, resolving common pitfalls when journeys become complex. Meanwhile, Kannan and Li (2017) provide a digital-marketing framework that explicitly situates touchpoints as loci where data, analytics, and automation reshape decisions – an early signal that AI would systematically alter journeys long before generative systems arrived.

Implication for AAO: If journeys are multi-episode, then optimization cannot target only single search moments; instead, AAO should be conceived as episode-aware, i.e., ensuring that

brand facts, claims, and evidence the agents retrieve are consistent and useful across pre-purchase triage, purchase facilitation, and post-purchase service.

AI in Marketing and Branding: From Augmentation to Agency

Reviews in leading outlets outline how AI alters marketing strategy, structures, and governance. Davenport et al. (2020) propose a three-dimensional lens (intelligence level, task type, embodiment) to anticipate where AI will remake marketing. Huang and Rust (2018) theorize AI's impact on service by decomposing tasks into mechanical, analytical, intuitive, and empathetic "intelligences", explaining task substitution vs. augmentation in service contexts. Libai et al. (2020) examine AI's consequences for customer-relationship management and call for safeguards as automation scales personalization. Field-level syntheses (Mustak et al., 2021; Verma et al., 2021; Vlačić et al., 2021) document the surge of AI-in-branding scholarship and map research needs (e.g., governance, explainability, cross-functional deployment). Collectively, these works shift attention from isolated tools to systemic, AI-mediated market interfaces – the conceptual bridge to agentic systems that curate, recommend, negotiate, or transact on customers' behalf.

Implication for AAO: As AI takes over interpretive and coordinative work formerly done by consumers, optimization must target the information diets of agents (what they ingest, trust, and surface), not merely human clickstreams.

Conversational and Embodied Agents: Anthropomorphism, Trust, and Compliance

At the interaction layer, research shows that conversational agents (CAs) and chatbots shape outcomes through social-cognitive mechanisms. Adam, Wessel, and Benlian (2021) demonstrate that design choices in AI chatbots influence user compliance, highlighting the role of anthropomorphism and control cues in service encounters. A multi-study meta-analysis synthesizing 82 samples finds that chatbot anthropomorphism tends to improve journey outcomes (e.g., engagement, satisfaction), with moderators such as service result valence and social-presence cues (Yanxia et al., 2024). Complementary work identifies when and why consumers trust chatbots (e.g., cognitive-appraisal pathways and task framing) and how flow experiences emerge in AI-mediated service (Baabdullah et al., 2022; Wang et al., 2023). The findings outlined in this study have the capacity to generalize beyond single touchpoints. Indeed, it has been demonstrated that agent characteristics could re-weight the salience of brand attributes, alter risk perceptions, and thereby redirect choice architectures across the journey.

Implication for AAO: Because anthropomorphic and social-presence cues change how agents frame brands, AAO must address not only content availability but also presentation constraints (e.g., structured, answer-oriented content that agent user interfaces (UIs) can render credibly and consistently).

Recommender Systems and the Economics of Visibility

Recommend-and-assist technologies are now central to discovery and evaluation. A large, cross-domain survey (Pappalardo et al., 2024) catalogs behavioral impacts of AI recommenders (concentration, diversity shifts, echo chambers, and polarization) across social, retail, urban, and generative-AI ecosystems. In retail and e-commerce specifically, recent reviews document rapid growth and methodological advances in AI-driven recommendation and in how review text is integrated into models that affect exposure and sales (e.g. Hasan et al., 2025; Valencia-Arias et al., 2024). Taken together, the literature implies that visibility is algorithmically endogenous: where

and how items appear is a function of agent models and training corpora, not just consumer search terms.

Implication for AAO: Since recommender exposure is model-mediated, brand optimization must include evidence curation (high-quality reviews, credible benchmarks, and structured product facts) that agents can ingest and reuse in explainable ways.

From SEO to answer-first AI search: the case for AI-agent optimization (AAO)

A new, research-backed thread in information retrieval is Generative Engine Optimization (GEO) – methodologies for increasing content visibility in LLM-based “answer engines” that synthesize multi-source responses (Aggarwal et al., 2024). Controlled experiments show that structured, machine-readable signals can lift citation/visibility in generative answers, underscoring that LLMs do source selection and synthesis rather than rank-ordered link retrieval. Complementary work demonstrates that strategic text sequences on product pages can influence LLM recommendations – evidence that answer engines can be steered by content design, for better or worse (Kumar & Lakkaraju, 2024). Managerial literatures describe parallel shifts in LLM-powered search behavior and the consequent need to optimize for agents as primary readers, not only for humans (Puntoni et al., 2024).

Building on that stream, AI-agent optimization (AAO) for branding can be defined as the systematic preparation of brand knowledge, evidence, and structures so that autonomous or semi-autonomous agents can (a) identify the brand entity, (b) verify claims against trusted corpora, and (c) select the brand in answer-first recommendations across journey stages. Conceptually, AAO extends GEO beyond “getting cited” to journey-stage adequacy (e.g., diagnostic facts for pre-purchase triage, risk- and price- justification content at purchase, policy and troubleshooting knowledge post-purchase), explicitly aligning optimization with agent tasks and user intents.

Scholarly guardrails: Two research strands warn against overclaiming: (1) AI answer engines can be susceptible to manipulation, implying governance and auditing are integral to AAO; and (2) philosophical work on artificial agency argues that optimizing for agents adds normative and design questions distinct from optimizing for intelligence per se (Floridi, 2025¹). These strands reinforce that AAO should remain evidence-first and ethics-aware (Kumar & Lakkaraju, 2024).

Gaps of the Present Study Addresses

Three persistent gaps motivate the AAO lens for online retail branding:

1. *Holism across the journey.* Much work isolates single touchpoints (e.g., chatbot service; dynamic pricing) rather than modeling agent-mediated continuity from pre-purchase to post-purchase (cf. reviews in AMJ and IJEM). AAO reframes optimization as stage-contingent and agent-aware across episodes (Chen & Prentice, 2025; Rana et al., 2022).
2. *Source preferences and machine-readability.* LLMs and agentic systems privilege structured, trusted, and well-linked sources when composing answers. The brand literature rarely specifies what evidence formats travel well into agent responses (schemas, FAQs, benchmarks). GEO research begins this specification; AAO extends it to branding tasks (Aggarwal et al., 2024; Kannan & Li, 2017).

¹ For definitional grounding in artificial agency, see Floridi (2025), Philosophy & Technology, who argues that optimizing for agents (not only intelligence) raises distinctive design and governance considerations relevant to AAO.

3. *Governance and trust in agent mediation.* Findings on anthropomorphism, compliance, and trust suggest design and disclosure choices affect outcomes; yet brand-level guidelines for agent-ready content and post-answer validation are nascent. AAO explicitly requires provenance, comparators, and counter-evidence to earn agent-level recommendations responsibly (Adam et al., 2021; Wang et al., 2023; Yanxia et al., 2024).

Conceptual Framework: The Agent-Orchestrated Journey (AOJ)

The paper conceptualizes AI agents as an orchestration layer that mediates intent, information, and action across the customer journey. Building on journey theory (Lemon & Verhoef, 2016) and digital-marketing frameworks that locate technology's influence at touchpoints (Kannan & Li, 2017), the Agent-Orchestrated Journey (AOJ) reframes how discovery, evaluation, choice, and post-purchase memory are produced when model-mediated systems summarize, recommend, and sometimes transact on customers' behalf. Under AOJ, four stages become analytically distinct yet interdependent:

Inquiry (need recognition and goal formation). Consumers increasingly articulate needs conversationally; answer engines fuse retrieval and generation to produce answers rather than ranked links. Visibility thus hinges on inclusion/citation within the synthesized response rather than position on a SERP. Early technical evidence from generative-engine optimization (GEO) indicates that content design and structure measurably affect whether sources are cited by LLMs, establishing a new optimization regime adjacent to SEO (Aggarwal et al., 2024).

Agentic mediation (planning and pruning). Agents compare alternatives, apply constraints, and compress consideration sets using preference, context, and reputation signals. Work on AI-enabled CX and CRM anticipates both personalization gains and discrimination risks as systems prioritize high-value customers (Huang & Rust, 2018; Libai et al., 2020).

Action (execution and transaction). Conversational agents shape compliance and conversion through social presence, anthropomorphism, and control cues; effects are moderated by task type and result valence (Adam et al., 2021), and meta-analytic synthesis in Yanxia et al. (2024).

Relationship & memory (post-purchase service and retention). Agents manage service episodes (tracking, returns, troubleshooting), influencing satisfaction, loyalty, and the locus of relationship (brand ↔ customer vs. agent-mediated). Reviews of AI recommenders show that algorithmic curation re-weights exposure and can reshape long-run behaviors – implications for brand memory and equity (Pappalardo et al., 2024).

Propositions: From this synthesis, the paper advances testable propositions to guide empirical work:

- P1 (Answer visibility). Under answer-first interfaces, citation/inclusion share within generated answers is a stronger mediator of awareness than rank on traditional SERPs (Aggarwal et al., 2024).
- P2 (Data sufficiency). Agentic pruning increases the premium on structured, trustworthy product and policy data for inclusion in recommendations (Kannan & Li, 2017).
- P3 (Design × trust). Anthropomorphic and transparency cues interact to shape compliance and satisfaction in agentic recommendations (Adam et al., 2021; Yanxia et al., 2024).
- P4 (Attribution shift). As agents manage post-purchase episodes, loyalty attributions shift from brand-customer ties to agent-mediated ties, complicating LTV models (Libai et al., 2020).

AI-Agent Optimization (AAO): Definition, Capabilities, and an Operational Contrast with SEO

How will brands respond to the ongoing structural transformations occurring in online retail? At present, brands are striving to achieve the highest possible rankings in search engines and online marketplaces. Empirical studies indicate that consumers are more inclined to choose options located near the top of a list of results, solely due to their position, regardless of relevance, price, or quality of the alternatives (Ursu, 2018). Consequently, securing the top position in a ranking provides online retailers with a considerable advantage. Likewise, the clickthrough rates of advertisements on search engines diminish as their positions lower (Agarwal, et al., 2011). Therefore, due to this ranking effect and position bias, brands allocate resources SEO and ranking optimization as essential tools for boosting brand sales.

In the new AI dominated world this could change. Should AI agents' purchasing decisions be less swayed by ranking effects and position bias and instead rely on an objective evaluation of various factors, it may become less rational for brands and retailers to allocate resources towards SEO or to secure a higher ranking on online platforms. Consequently, alternative strategies might need to be implemented. Instead of focusing on enhancing their visibility for search engines, brands could consider refining their marketing strategies to suit agentic consumption. Although conventional SEO methods remain important, they typically emphasize indicators tailored for traditional search engines (such as keyword density and meta tags), which may not be adequate for agentic AI that depends on alternative (and more profound) types of information quality. To maintain machine-retrievability in generative results, brands must integrate AAO strategies alongside traditional SEO.

The paper defines AI-agent optimization (AAO) as the cross- functional capability that makes a brand correctly findable, faithfully represented, and action-ready inside model- mediated experiences across the journey. AAO extends GEO's focus on getting cited by generative engines to the full AOJ: (a) Content and evidence design (concise claims, verifiable facts, comparative benchmarks), (b) Technical signals (entity hygiene, schemas, knowledge-graph alignment, accessible product feeds/APIs), (c) Measurement (share-of- answer, citation recall, agent-recommendation share), and (d) Governance (provenance, disclosure, safety).

Table 1 contrasts traditional SEO with AAO along the dimensions most relevant to answer-first, agent- mediated journeys. Where possible, rows with mechanisms empirically noted in GEO experiments, are aligned (content structure → inclusion/citation) and with touchpoint logic in digital-marketing scholarship (Kannan & Li, 2017; Lemon & Verhoef, 2016).

Table 1. SEO vs. AAO: An operational contrast for answer-first, agent- mediated journey

Dimension	Traditional SEO (human-first SERPs)	AAO (agent-/answer-first systems)
Primary unit of visibility	Rank-ordered URL on SERP	Inclusion/citation within generated answer or agent recommendation (GEO).
Optimization target	Ranking signals (authority, relevance, backlinks)	Citation/selection signals (structured claims, evidence density, entity clarity) that improve answer inclusion.
Content design	Keyword targeting; readability for humans	Evidence-rich, structured facts (tables, FAQs, specs) optimized for model parsing/synthesis (e.g., schema/knowledge-graph alignment).
Measurement	Rank, CTR, impressions, dwell time	Share-of- Answer (SoA), Citation Recall, Agent Recommendation Share; answer-layer A/B tests (prompt- set audits).

Dimension	Traditional SEO (human-first SERPs)	AAO (agent-/answer-first systems)
Retrieval pathway	Index → query → ranked list	Retrieval-augmented generation → synthesis → cited snippets; agent policy filters.
Failure/risks	Link schemes; keyword stuffing	Hallucination, omission, prompt/answer manipulation; need for provenance and counter-evidence.
Governance emphasis	Webmaster guidelines; spam policies	Provenance and auditability of brand facts; disclosure and safety guardrails for agentic use. (Conceptual alignment with service/CRM governance.)

Source: Author's own work.

Implications: AAO does not replace SEO; rather, it extends optimization to the selection and reasoning layers that determine whether and how a brand appears in generated answers and agentic flows. In operational terms, AAO requires (i) entity hygiene (consistent names, IDs, attributes), (ii) fact packaging (concise, verifiable claims with citations the models can attribute), and (iii) agent-ready interfaces (feeds/APIs enabling agents to fetch specs, availability, and post-purchase policies reliably).

Methodology

Review type and rationale

This study is a structured narrative review that synthesizes peer-reviewed research across marketing, information systems, human-computer interaction, and retailing to build an integrative conceptual framework (AOJ) and the capability construct AI-agent optimization (AAO). A structured narrative approach is appropriate when the goal is theory integration rather than effect-size estimation and when bodies of evidence are heterogeneous across methods and disciplines (Snyder, 2019; Tranfield et al., 2003). To enhance transparency and reproducibility, the research aligned reporting with the PRISMA 2020 guidance and, where relevant to scope-mapping, with PRISMA- ScR conventions (Page et al., 2021).

Sources and search strategy

The paper targeted leading bibliographic databases that index marketing and adjacent fields: Scopus, Web of Science Core Collection, Business Source Complete, PsycINFO, and ACM Digital Library. Searches covered January 2016–October 2025 to capture contemporary AI developments while retaining foundational pre-2016 works (e.g., Lemon & Verhoef, 2016) for theoretical anchoring. Representative query templates combined journey, agent, and branding terms (Boolean operators and truncations used per database rules), for example:

"customer journey" AND (AI OR "artificial intelligence" OR agent* OR chatbot*) AND (retail OR e-commerce OR branding)
(generative OR LLM OR "answer engine") AND (search OR discovery) AND (brand* OR retail)
(recommender OR "decision support") AND (consumer OR customer) AND (brand equity OR loyalty)

Given reporting expectations for literature searches, search strings, date stamps, and database-specific limits are documented in a supplementary file following PRISMA- S (Rethlefsen et al., 2021).

Eligibility criteria

Inclusion criteria prioritized peer-reviewed journal articles and scholarly reviews (English language) that: (a) theorize or empirically examine AI/agent systems in consumer journeys, branding, retail, or service, or (b) provide methodological/technical foundations directly shaping how AI systems retrieve, synthesize, or recommend brand information (e.g., recommender-system surveys; generative-engine optimization). Seminal pre- 2016 pieces were retained when they provided canonical journey theory or conceptual guidance. Exclusion criteria removed non-scholarly commentary and practitioner pieces unless used to contextualize platform changes (kept separate from the analytic corpus). When emergent topics lacked archival coverage (e.g., GEO), the research tracked peer-reviewed or preprinted sources cautiously and labeled them explicitly.

Screening and selection

Titles/abstracts were screened for topical relevance to AI-mediated discovery, evaluation, transaction, or post-purchase phenomena. Full texts were examined for construct clarity (definitions of “agent,” “assistant,” “chatbot,” “recommender,” “answer engine”), research design, and direct link to branding or journey outcomes. The research favored journals with established peer review in marketing, information systems, and service (e.g., Journal of Marketing, Journal of Retailing, Journal of Service Research), while including adjacent outlets for technical depth (e.g., recommender-system surveys). The resulting corpus underpins the synthesis in §§2–4 and the propositions in §3. (Procedural reporting follows PRISMA 2020 headings to facilitate replication by other scholars.) (Page et al., 2021).

Data extraction and coding

The research extracted bibliographic data, focal constructs, context (industry/setting), methods, measures, and reported mechanisms. A thematic synthesis procedure was applied to code findings into first-order concepts (e.g., anthropomorphism effects, agent transparency, evidence density), second-order themes (e.g., trust formation in agentic interactions), and aggregate dimensions aligned with AOJ stages (Inquiry, Agentic Mediation, Action, Relationship & Memory). Thematic synthesis is appropriate for integrating heterogeneous qualitative and quantitative insights while preserving proximity to primary claims (Thomas & Harden, 2008).

Synthesis logic and conceptual integration

Following guidance on developing conceptual contributions in marketing (MacInnis, 2011) and the value of conceptual integration to knowledge advancement (Yadav, 2010), the research used abductive reasoning to iteratively map coded evidence onto the Agent-Orchestrated Journey (AOJ) and to specify AI-agent optimization (AAO) as a capability construct. This process connected micro-level findings (e.g., chatbot social- presence effects) to meso- level journey mechanisms (e.g., compliance, pruning of consideration sets) and macro-level platform shifts (e.g., answer-first search).

Quality assurance and bias mitigation

To reduce search/reporting bias, the research (a) queried multiple databases spanning marketing and computing, (b) used iteratively refined Boolean strings, and (c) documented information-source details per PRISMA- S. To mitigate selective-citation risk, the research

triangulated findings across disciplines and prioritized archival, peer-reviewed sources for claims – reserving preprints for nascent technical topics and labeling them accordingly (e.g., GEO). The research also distinguished platform announcements/audits from scholarly evidence and used the former only to contextualize 2024–2025 changes described in §1 (Rethlefsen et al., 2021).

Limitations

As a conceptual synthesis, this review does not compute pooled effects nor exhaustively cover all domains affected by agents. Selection is limited to English-language scholarship and peer-reviewed outlets (with targeted exceptions for foundational preprints). Rapidly evolving platform practices (e.g., search answer interfaces) introduce recency bias risks; the research therefore treats those developments as context rather than as inferential evidence. Future empirical work should test §3 propositions using field experiments and agent-exposure measurements (e.g., share-of- answer ; citation recall).

Findings

AI agents are set to transform how consumers search for branded products. Currently, the customer journey frequently initiates with a search engine, such as Google, or a major online marketplace like Amazon, which has effectively evolved into a product search engine. In the foreseeable future, AI agents may act as novel access points. As a result, brands need to leverage the capabilities of AAO to provide engaging results to their intended audience, thus improving the overall consumer experience.

The integration of AAO into marketing and branding strategies has significantly transformed the conventional consumer decision-making process. Rather than adhering to the linear path usually linked with decision-making, where consumers navigate through a set sequence of steps, AAO has introduced a non-linear approach. This shift from traditional methods suggests that consumers can engage in the decision-making process at different stages (Sharma et al., 2023), directly influenced by AAO. The widespread impact of AAO highlights the changing dynamics of consumers' purchasing behavior as a reaction to technological progress.

The commonly accepted stages of the customer journey encapsulate the brand shopping experiences of customers through three distinct stages: pre-purchase, purchase, and post-purchase. This straightforward framework is versatile enough to apply to a diverse array of products and purchasing contexts. Considering the previously proposed four staged conceptual framework about the AOJ, it should be noted that the stages are connected as follows: inquiry (need recognition and goal formation) and agentic mediation (planning and pruning) are related to the pre-purchase stage of the customer journey, action (execution and transaction) correspond to the purchase stage and relationship & memory (post-purchase service and retention) to the post -purchase. a rough outline of the AI related *cross-stage regularities* that can impact a customer's brand decision journey is presented below.

Cross-stage regularities from the literature

Pre-purchase (discovery and evaluation). Studies consistently show that conversational/agentic interfaces compress consideration sets by filtering and summarizing alternatives, with outcomes mediated by anthropomorphic and social-presence cues (Adam et al., 2021; Yanxia et al., 2024). Journey work in marketing implies that this compression occurs against a multi-episode background in which perceptions formed early shape downstream evaluations (Kannan & Li, 2017;

Lemon & Verhoef, 2016). Evidence from AI-in- marketing reviews further indicates that personalization benefits coexist with selection risks (e.g., omission or bias), reinforcing the need for auditable facts and sources during discovery (Davenport et al., 2020; Mustak et al., 2021; Verma et al., 2021).

Capabilities. Conversational retrieval, entity linking, and recommender systems compress search costs by matching needs to brand entities, summarizing pros/cons, and exposing tradeoffs.

Outcomes. Higher perceived relevance, faster narrowing of consideration sets, but greater dependence on third-party evidence (reviews, media, knowledge bases).

Brand task. Ensure the brand is disambiguated as an entity with answer-ready evidence and transparent policies.

Purchase (decision and transaction/action). In transactional contexts, chatbots and virtual agents can reduce friction by providing just-in- time assistance, with compliance and conversion sensitive to agent design and transparency (Adam et al., 2021; Baabdullah et al., 2022; Wang et al., 2023). Conceptual and empirical work suggests that AI shifts price/assortment navigation from manual comparison to machine-assisted curation, aligning with the AOJ claim that agents increasingly orchestrate choice (Huang & Rust, 2018; Libai et al., 2020).

Capabilities. Dynamic offers and assisted checkout reduce friction; agents handle clarifications on specs, policies, and shipping.

Outcomes. Lower abandonment when agents provide precise, context-aware help; increased cross-sell via constraint-satisfying bundles.

Brand task. Align pricing/availability data feeds and policy explanations with agent-consumable formats.

Post-purchase (use, support, loyalty relationship & memory). Agentic tools manage returns, troubleshooting, and proactive recommendations, with potential to strengthen or displace brand–customer ties depending on perceived competence and care (Araújo et al., 2022; Libai et al., 2020). Large-scale surveys of recom mender systems indicate that algorithmic exposure patterns shape long-run behaviors (diversity/concentration), implying brand- equity consequences for how agents “remember” and resurface a brand (Pappalardo et al., 2024).

Capabilities. Agents orchestrate support, nudge registration/returns, and propose add-ons based on observed use.

Outcomes. Higher service responsiveness, opportunity for loyalty reinforcement, but visibility of failure modes if evidence contradicts claims.

Brand task. Close the loop: expose updated manuals, troubleshooting trees, and feedback channels in structured, crawlable formats.

Technical evidence from generative-engine optimization (GEO) demonstrates that content design and structure can measurably increase inclusion/citation within generated answers, reframing discoverability from link rank to answer presence (Aggarwal et al., 2024). This corroborates the need for AI-agent optimization (AAO) that targets selection and reasoning layers across the journey, not merely SERP rank.

Operational metrics for an answer-/agent- first marketplace

To manage brand performance under AOJ and implement AAO, the research defines three metrics that adapt classic information-retrieval ideas (precision/recall) to the answer layer (cf. Manning et al., 2008) and are grounded in well-established evaluation frameworks from information retrieval and recommender-systems research:

SoA (Share-of-Answer) - framed as a weighted inclusion rate across prompts/engines; conceptually tied to IR's precision/recall setup (aggregating binary inclusion over a test set). See Manning, Raghavan, & Schütze (2008) for the classical treatment of precision/recall.

CR (Citation Recall) - a recall-style metric over a curated "ground-truth" fact set; directly analogous to IR recall computed at the claim level (again, Manning et al.).

ARS (Agent Recommendation Share) - a top-k selection rate across agent tasks; the rank-weighted variant corresponds to DCG/NDCG-style utility weighting long used in ranking/recommender evaluation (Järvelin & Kekäläinen (2002) for DCG/NDCG; Herlocker et al. (2004) and Shani & Gunawardana (2011) for recommender hit/recall@k and ranking evaluation practice).

The notation that is used (common to all three indicators) is as follows:

B : the focal brand (or SKU family).

P : set of prompts tested; weight $w_p \geq 0$ for prompt $p \in P$.

E : set of answer engines (e.g., AI Overviews, Perplexity); weight $v_e \geq 0$ for engine $e \in E$.

T : set of agent tasks/templates (e.g., "recommend 3 laptops under €1000"); weight $\beta_t \geq 0$ for task $t \in T$.

A : set of agents (e.g., retailer/vertical agents); weight $u_a \geq 0$ for agent $a \in A$.

$\mathbb{1}[\cdot]$: indicator function (1 if statement true, else 0).

k : top-k cutoff for inclusion (default $k = 3$ unless specified).

All metrics are *bounded in* $[0,1]$ by construction.

(1) Share-of- Answer (SoA)

Definition. Probability-weighted share of prompts whose generated answer (not the link list) includes or cites B .

Let $I_{p,e}^{(k)}(B) = 1$ if B is (a) named, (b) linked, or (c) cited as a source *within the top-level generated answer* for prompt p on engine e (or appears among the engine's *top-k* recommendations), and 0 otherwise.

$$SoA_k(B) = \frac{\sum_{p \in P} \sum_{e \in E} w_p v_e I_{p,e}^{(k)}(B)}{\sum_{p \in P} \sum_{e \in E} w_p v_e}. \quad (1)$$

Unweighted variant: set $w_p = v_e = 1$.

Single-engine SoA: drop the sum over e .

Purpose: Quantify brand visibility inside generated answers (pre /mid journey inquiry).

(2) Citation Recall (CR)

Definition. Fraction of *curated, ground-truth facts* about B that are correctly surfaced at least once across the evaluated prompts/engines.

Let F be the vetted set of brand facts; optional importance weights $\alpha_f \geq 0$ with $\sum_{f \in F} \alpha_f = 1$.

Let $J_{f,p,e}(B) = 1$ if fact f is *correctly* expressed in the generated answer for (p, e) (verbatim or faithful paraphrase), else 0.

Any-hit (global) recall (a fact counts if it appears correctly *at least once* across all (p, e)):

$$CR_{any}(B) = \sum_{f \in F} \alpha_f \mathbb{1}[\sum_{p \in P} \sum_{e \in E} J_{f,p,e}(B) \geq 1]. \quad (2a)$$

Average per-prompt/engine recall (recall computed per answer, then averaged with (w_p, v_e) weights):

$$CR_{avg}(B) = \frac{\sum_{p \in P} \sum_{e \in E} w_p v_e (\sum_{f \in F} \alpha_f J_{f,p,e}(B))}{\sum_{p \in P} \sum_{e \in E} w_p v_e} \quad (2b)$$

Optional companion metric (precision):

If $S_{p,e}(B)$ is the number of *brand-related facts* the answer asserts for (p, e) , then:

$$\text{Precision}(B) = \frac{\sum_{p,e} w_p v_e \sum_{f \in F} \alpha_f J_{f,p,e}(B)}{\sum_{p,e} w_p v_e S_{p,e}(B)}$$

Purpose: Assess fidelity of the brand's representation (facts captured by engines).

(3) Agent Recommendation Share (ARS)

Definition. Share of agent tasks for which B is *selected* by an agent among its top- k recommended items (or appears at/near the top).

Let $r_{a,t}(B) \in \{1, 2, \dots\} \cup \{\infty\}$ be the *rank* of B in agent a 's recommendation list for task t (∞ if not recommended). The *top- k inclusion* indicator is:

$$R_{a,t}^{(k)}(B) = 1[r_{a,t}(B) \leq k].$$

Top- k ARS:

$$ARS_k(B) = \frac{\sum_{t \in T} \sum_{a \in A} \beta_t u_a R_{a,t}^{(k)}(B)}{\sum_{t \in T} \sum_{a \in A} \beta_t u_a} \quad (3a)$$

Rank-weighted ARS (more credit for higher ranks): choose a decreasing weight function $g: \mathbb{N} \rightarrow [0, 1]$, (e.g., $g(r) = 1/\log_2(1 + r)$; set $g(\infty) = 0$):

$$ARS^g(B) = \frac{\sum_{t \in T} \sum_{a \in A} \beta_t u_a g(r_{a,t}(B))}{\sum_{t \in T} \sum_{a \in A} \beta_t u_a g(1)} \quad (3a)$$

Purpose: Track selection by agents during planning/transaction (mid-journey)

Implementation notes (brief, for full implementation toolkit see Appendix A):

Modality for $I_{p,e}^{(k)}(B)$: pre-register what counts as “included” (named mention, linked URL, attributed citation) and use that rule consistently when scoring SoA.

Fact scoring for $J_{f,p,e}(B)$: evaluate at the *claim level* against a ground-truth table (with DOIs/manuals) and record error types (omission, distortion, hallucination).

Reliability: for CR and SoA, run multiple sessions per (p, e) and compute inter-rater agreement for manual scoring.

Practical instrumentation

Prompt/test set design. Derive prompts from consumer language (site search logs, FAQs) and public corpora; cluster by intent and vertical. Maintain balance and update quarterly.

Engine/agent coverage. Audit across major answer engines and relevant retailer/vertical agents; report per-engine metrics and an exposure-weighted composite.

Reliability. Use repeated runs and session randomization; compute test-retest reliability and inter-rater agreement for CR coding.

Benchmarking. Track peers' SoA/ARS where feasible via neutral descriptors (“top-selling X under Y”) to avoid brand-name prompt bias.

Governance linkage. When CR identifies systematic omissions or errors, file provenance tickets and adjust content packaging/structured data before pursuing appeal mechanisms (cf. governance guidance in Libai et al., 2020; Huang & Rust, 2018).

Managerial implications (Brand-centric AAO playbook)

The objective of AAO is brand visibility and fidelity inside agent-mediated experiences. The unit of optimization is not a page but the brand entity – its facts, claims, and proofs as consumed by answer engines and shopping agents. The following actions translate AAO into brand management practice (owners and artifacts are suggested to aid execution).

Brand entity hygiene and knowledge scaffolding. Normalize the brand’s legal names, product lines, sub brands, and identifiers across sites, feeds, knowledge graphs, and retailer catalogs. Publish machine readable ground truth tables (SKUs, specs, compatibilities, policies) and link them to canonical brand IDs. This reduces ambiguity during entity resolution and raises Citation Recall and Share-of-Answer.

Owner: Brand management × Data engineering × Legal.

Artifacts: Master brand/architecture registry; schema.org markup; Wikidata/GS1 entries; retailer feed specs.

Success measures: CR↑, SoA↑; reduction in misattribution/aliasing cases.

Evidence-rich brand content packaging. Recast positioning claims into compact, verifiable statements paired with tabular specifications, FAQs, and third-party benchmarks. Package differentiators (safety, sustainability, performance) with citations that agents can surface. Avoid ambiguous phrasing and duplicate descriptions that degrade synthesis.

Owner: Brand comms × Product marketing.

Artifacts: Claim library with source links; comparison tables; concise “answer cards.”

Success measures: SoA↑ on informational/comparative prompts; CR↑ for core claims.

Conversational answerability of the brand voice. Write page-level copy to answer the intents consumers actually express (e.g., “Which size fits me for ...?”, “Is it compatible with ...?”), using the brand’s tone but making constraints and examples explicit. Provide canonical phrasings that agents can quote faithfully to preserve voice and legal accuracy.

Owner: Brand editorial × CX content.

Artifacts: Intent → answer map; tone/voice guardrails; microcopy library.

Success measures: SoA↑ on “how/which/compare” prompts; reduction in mis-quotes in agent answers.

Review and UGC stewardship as brand evidence. Curate authentic, diverse reviews and maintain post-purchase Q&A that clarifies fit, use, and service. Encourage credible third-party tests and media coverage; respond to inaccuracies with sourced corrections. Because recommenders and agents reuse UGC as evidence, stewardship directly affects long-run exposure and consideration.

Owner: Brand advocacy × Community/PR.

Artifacts: Review seeding guidelines; expert/test protocols; correction playbook.

Success measures: ARS↑ in agent shortlists; balanced sentiment mix; fewer omission errors in CR.

Trust-by-design in owned brand agents. Design the brand’s chat/voice agents with transparent rationales (why-this-recommendation), controllable hand-offs to humans, and calibrated anthropomorphism appropriate to task risk. Align scripts with brand values (care,

safety, inclusivity) so the agent itself strengthens brand meaning rather than merely transacting.

Owner: CX/Service × Brand × Compliance.

Artifacts: Explanation templates; escalation matrix; risk-tiered dialogue policies.

Success measures: Resolution and CSAT↑; complaint rate↓; conversion lift in assisted flows.

Measurement integration for brand equity under mediation. Incorporate SoA, CR, and ARS into brand dashboards next to awareness, consideration, and preference. Treat answer exposure as an upstream event in media mix models and brand-lift studies, recognizing that clicks may be sparse when agents summarize.

Owner: Brand insights/Analytics.

Artifacts: Answer-layer audit suite; SoA/CR/ARS panels by market and engine; MMM inputs.

Success measures: Statistically linked improvements from SoA/CR/ARS to aided/unaided awareness and consideration.

Compliance, provenance, and claims substantiation. Maintain living documentation of sources behind brand claims, update cadences, takedown/appeal logs, and regional policy variants. Align disclosures and data handling with service/consumer-protection standards so that provenance becomes a positive optimization signal as well as a safeguard.

Owner: Legal/Regulatory × Brand comms.

Artifacts: Claims registry; provenance ledger; regional policy matrix.

Success measures: CR accuracy by policy facts; audit pass rates; reduction in flagged hallucinations.

Test–learn cadence and brand governance. Run quarterly GEO/AAO audits; pair content/structure tests with live ABs (e.g., FAQ rewrites → SoA/CR lift → conversion). Establish a cross-functional Brand–Agent Governance Board (Brand, CX, Data, Legal) to prioritize experiments, monitor risks, and retire patterns that inadvertently bias recommendations.

Owner: CMO staff office.

Artifacts: Quarterly audit reports; experiment backlog; risk register.

Success measures: Sustained SoA/CR/ARS gains; time-to-fix for misrepresentation incidents.

Limitations and Future Research

Limitations (brand-specific)

First, as a structured narrative review, this work integrates heterogeneous literatures without estimating pooled effects; causal magnitudes linking answer-layer metrics (SoA, CR, ARS) to brand outcomes (awareness, consideration, equity) remain to be established. Second, fast-evolving platform behavior introduces recency and opacity risks: engines change answer policies and citation practices, and proprietary agent logic limits observability—both complicate brand inference. Third, instrumentation for SoA/CR/ARS depends on prompt/task design and session variability; although reliability procedures mitigate noise, construct validity requires continued refinement with brand-appropriate gold standards. Fourth, brand heterogeneity (size, distribution, pricing power, category norms) may confound relationships between answer visibility and equity; findings may not generalize across B2C/B2B or regulated categories. Finally, adversarial dynamics (e.g., knowledge-graph poisoning, astroturfed reviews) can distort answer layers in ways not fully addressed here.

Future research agenda (brand-impact focus)

Several additional issues, worthy of future research, are summarized below.

Causal identification of brand lift. Field experiments that randomize answer exposure (SoA) or agent shortlist inclusion (ARS) by market or time to estimate incremental effects on aided/unaided awareness, consideration, and sales.

Brand equity under mediation. Longitudinal studies connecting SoA/CR/ARS changes to memory structures (distinctive assets, salience) and to post-purchase loyalty; integration with brand-tracking models.

Fairness, visibility, and market structure. Tests for differential inclusion/omission by brand size, margin structure, or geography; interventions that improve visibility for challenger brands without compromising quality.

Design science for AAO. Controlled trials of content/structure patterns that improve CR without gaming, plus evaluations of counterfactual and contrastive explanations that protect brand meaning while increasing trust.

Governance and contestability. Mechanisms for transparent provenance, appeal protocols when agents misrepresent brand facts, and audits that detect knowledge-base poisoning or review manipulation.

Cross-cultural and regulatory moderation. Comparative studies of agent adoption, trust, and disclosure effectiveness across jurisdictions; interaction with the EU AI Act and consumer-protection regimes.

Attribution models for the answer layer. New MMM variants and causal graphs that explicitly include answer-layer exposures and agent selections, mapping SoA/CR/ARS to conversion and retention.

Conclusion

The digital search environment is experiencing a profound transformation due to the swift progress of AI. Traditional search engines, which previously depended on keyword matching and link-based rankings, are increasingly incorporating AI-driven technologies. These technologies offer direct answers, summarize content, and facilitate conversational search, thereby redefining the manner in which users access information. Often, they provide accurate responses without necessitating clicks to external websites. This evolution signifies the emergence of AI-agent optimization (AAO) as a new paradigm that complements established SEO practices, compelling marketers and brands to adapt. For many years, SEO strategies have been founded on principles such as keyword optimization, backlink acquisition, and on-page technical improvements. However, the rise of AI-powered answer engines is reducing the significance of traditional organic search results. Research indicates an increase in zero-click searches, where users receive answers directly from SERPs. This trend raises critical questions regarding how brands can sustain visibility when search engines bypass their websites, what new optimization strategies should accompany traditional SEO, and what the future may hold for online retail in an AI-centric search ecosystem.

This paper advances the field by (1) synthesizing how agentic, answer-first systems reshape the customer journey; (2) defining AI-agent optimization (AAO) as an integrative capability framework grounded in recent platform developments and technical research; and (3) proposing measurement constructs and managerial design principles that enable brands to compete for answers, not just rankings.

Agents already mediate how needs become choices. AAO gives brands a pragmatic path to remain present in this mediation: be easy to identify, easy to cite, and easy to recommend. Competing on evidence quality and machine legibility is the strategic hedge in an answer-first web.

References

- Adam, M., Wessel, M., & Benlian, A. (2021). AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets*, 31(2), 427–445.
- Agarwal, A., Hosanagar, K., & Smith, M. D. (2011). Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6), 1057–1073.
- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). GEO: Generative Engine Optimization.
- Araújo, C., Gonçalves, R., Costa, R. L. Da, Dias, Á., & Pereira, L. (2022). Artificial intelligence in the digital customer journey. *International Journal of Electronic Customer Relationship Management*, 13(3), 248.
- Baabdullah, A. M., Alalwan, A. A., Algharabat, R. S., Metri, B., & Rana, N. P. (2022). Virtual agents and flow experience: An empirical examination of AI-powered chatbots. *Technological Forecasting and Social Change*, 181, 121772.
- Chen, Y., & Prentice, C. (2025). Integrating Artificial Intelligence and Customer Experience. *Australasian Marketing Journal*, 33(2), 141–153.
- Court, D., Elzinga, D., Mulder, S., & Vetvik, O. J. (2009). The consumer decision journey. *McKinsey Quarterly*, 3, 97–107.
- Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42.
- Floridi, L. (2025). AI as Agency without Intelligence: On Artificial Intelligence as a New Form of Artificial Agency and the Multiple Realisability of Agency Thesis. *Philosophy & Technology*, 38(1), 30.
- Grewal, D., & Roggeveen, A. L. (2020). Understanding Retail Experiences and Customer Journey Management. *Journal of Retailing*, 96(1), 3–8.
- Hasan, E., Rahman, M., Ding, C., Huang, J. X., & Raza, S. (2025). Based recommender systems: a survey of approaches, challenges and future perspectives. *ACM Computing Surveys*, 58(1), 1–41.
- Herlocker, J. L., Konstan, J. A., Terveen, L. G., & Riedl, J. T. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1), 5–53.
- Huang, M.-H., & Rust, R. T. (2018). Artificial Intelligence in Service. *Journal of Service Research*, 21(2), 155–172.
- Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, 20(4), 422–446.
- Kannan, P. K., & Li, H. “Alice.” (2017). Digital marketing: A framework, review and research agenda. *International Journal of Research in Marketing*, 34(1), 22–45.
- Kumar, A., & Lakkaraju, H. (2024). Manipulating large language models to increase product visibility. *arXiv preprint arXiv:2404.07981*.

- Lemon, K. N., & Verhoef, P. C. (2016). Understanding Customer Experience Throughout the Customer Journey. *Journal of Marketing*, 80(6), 69–96.
- Libai, B., Bart, Y., Gensler, S., Hofacker, C. F., Kaplan, A., Kösterheinrich, K., & Kroll, E. B. (2020). Brave New World? On AI and the Management of Customer Relationships. *Journal of Interactive Marketing*, 51(1), 44–56.
- MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of Marketing*, 75(4), 136–154.
- Manning, C. D., Raghavan, P., & Schütze, H. (2008). Introduction to Information Retrieval. In *Introduction to Information Retrieval*.
- Mustak, M., Salminen, J., Plé, L., & Wirtz, J. (2021). Artificial intelligence in marketing: Topic modeling, scientometric analysis, and research agenda. *Journal of Business Research*, 124, 389–404.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *Systematic Reviews*, 10(1), 89.
- Pappalardo, L., Ferragina, E., Citraro, S., Cornacchia, G., Nanni, M., Rossetti, G., ... & Pedreschi, D. (2024). A survey on the impact of AI-based recommenders on human behaviours: methodologies, outcomes and future directions. *arXiv preprint arXiv:2407.01630*.
- Puntoni, S., Ensing, M., & Bowers, J. (2024). How marketers can adapt to LLM-powered search. *Harvard Business Review*.
- Rana, J., Gaur, L., Singh, G., Awan, U., & Rasheed, M. I. (2022). Reinforcing customer journey through artificial intelligence: a review and research agenda. *International Journal of Emerging Markets*, 17(7), 1738–1758.
- Rethlefsen, M. L., Kirtley, S., Waffenschmidt, S., Ayala, A. P., Moher, D., Page, M. J., Koffel, J. B., Blunt, H., Brigham, T., Chang, S., Clark, J., Conway, A., Couban, R., de Kock, S., Farrah, K., Fehrmann, P., Foster, M., Fowler, S. A., Glanville, J., ... Young, S. (2021). PRISMA-S: an extension to the PRISMA Statement for Reporting Literature Searches in Systematic Reviews. *Systematic Reviews*, 10(1), 39.
- Shani, G., & Gunawardana, A. (2010). Evaluating recommendation systems. In *Recommender systems handbook* (pp. 257–297). Boston, MA: Springer US.
- Sharma, P., Ueno, A., Dennis, C., & Turan, C. P. (2023). Emerging digital technologies and consumer decision-making in retail sector: Towards an integrative conceptual framework. *Computers in human Behavior*, 148, 107913.
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339.
- Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8(1), 45.
- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a Methodology for Developing Evidence-Informed Management Knowledge by Means of Systematic Review. *British Journal of Management*, 14(3), 207–222.
- Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4), 530–552.

- Valencia-Arias, A., Uribe-Bedoya, H., González-Ruiz, J. D., Santos, G. S., Ramírez, E. C., & Rojas, E. M. (2024). Artificial intelligence and recommender systems in e-commerce. Trends and research agenda. *Intelligent Systems with Applications*, 24, 200435.
- Verma, S., Sharma, R., Deb, S., & Maitra, D. (2021). Artificial intelligence in marketing: Systematic review and future research direction. *International Journal of Information Management Data Insights*, 1(1), 100002.
- Vlačić, B., Corbo, L., Costa e Silva, S., & Dabić, M. (2021). The evolving role of artificial intelligence in marketing: A review and research agenda. *Journal of Business Research*, 128, 187–203.
- Wang, C., Li, Y., Fu, W., & Jin, J. (2023). Whether to trust chatbots: Applying the event-related approach to understand consumers' emotional experiences in interactions with chatbots in e-commerce. *Journal of Retailing and Consumer Services*, 73, 103325.
- Yadav, M. S. (2010). The decline of conceptual articles and implications for knowledge development. *Journal of Marketing*, 74(1).
- Yang, K. C. (2025). News Source Citing Patterns in AI Search Systems.
- Yanxia, C., Shijia, Z., & Yuyang, X. (2024). A meta-analysis of the effect of chatbot anthropomorphism on the customer journey. *Marketing Intelligence & Planning*, 42(1), 1–22.

Appendix

Appendix A. Metric Implementation Toolkit (Template Prompt Set & CR Scoring Rubric)

Scope. This appendix provides (1) a parameterized prompt set to audit answer-layer performance across engines/agents and (2) a claim-level rubric for scoring **Citation Recall (CR)** as defined in §6.2.

A1. Template Prompt Set (balanced by intent and constraint)

Tokens: {CATEGORY}, {BRAND}, {SKU}, {ATTRIBUTE}, {PRICE_CAP}, {USE_CASE}, {REGION}, {COMPETITOR}.

Design: 12 prompts (4 informational, 4 comparative, 4 transactional), randomized order; run across engines/agents and two sessions per (engine × prompt).

Informational (I)

1. *What is the {ATTRIBUTE} of the {BRAND} {SKU} {CATEGORY} in {REGION}?*
Ex. battery life; warranty length; return window.
2. *Does the {BRAND} {SKU} {CATEGORY} support {USE_CASE}?*
Ex. airline carry-on size; waterproof to IPX rating; HDMI 2.1.
3. *What materials/specifications define the {BRAND} {SKU} {CATEGORY}?*
Ex. fabric composition; processor; dimensions; weight.
4. *What is the official returns/warranty policy for {BRAND} {CATEGORY} in {REGION}?*

Comparative (C)

5. *Compare {BRAND} {SKU} vs. {COMPETITOR} on {ATTRIBUTE} under {PRICE_CAP}.*
6. *Which {CATEGORY} best for {USE_CASE}: {BRAND} {SKU} or {COMPETITOR}?*
7. *Is {BRAND} {SKU} compatible with {ACCESSORY/PLATFORM} while {COMPETITOR} is not?*
8. *Rank {BRAND} {SKU} and {COMPETITOR} by {ATTRIBUTE} with references.*

Transactional (T)

9. Find the {BRAND} {SKU} {CATEGORY} under {PRICE_CAP} in {REGION} with delivery by {DATE}.
10. Where to buy {BRAND} {SKU} {CATEGORY} with {FEATURE} included?
11. Show bundles for {BRAND} {SKU} {CATEGORY} that add {ACCESSORY} under {PRICE_CAP}.
12. If {BRAND} {SKU} is out of stock, recommend the closest alternative meeting {ATTRIBUTE}.

Fact set (F) for CR auditing. Pre-register a ground-truth table of ~15–30 atomic facts per brand/SKU (e.g., battery_life_hours, weight_g, dimensions_mm, warranty_years, return_days, compatibility_list, materials, price_msrp, certifications). Each fact includes: canonical value, unit, valid region, effective date, source/DOI/URL.

A2. Citation Recall (CR) Scoring Rubric (claim-level)

Unit of analysis. A fact $f \in F$ scored within a specific generated answer to a prompt–engine (or prompt–agent) pair. CR can be computed as **any-hit (global)** or **average per answer** (see §6 equations).

A2.1 Coding categories (strict binary for CR; optional soft credit shown)

Code	Label	Decision rule (apply in order)
C	Correct & complete	Value matches canonical and units/constraints (region/date) are correct; faithful paraphrase permitted.
CI	Correct but imprecise	Semantically correct but rounded or unit converted without material change (e.g., 9.8 → 10 hrs; inches ↔ mm) and constraint context preserved.
P	Partially correct	Mixed claim: includes the correct value and an ancillary inaccuracy or omits a required qualifier that changes interpretability (e.g., “waterproof” without IP code).
O	Omitted	Fact not stated.
H	Hallucinated/incorrect	Value contradicts canonical table, wrong unit, wrong region/date, or invented attribute.

*Use soft credit only for *exploratory* analysis; report strict CR as the primary metric.

Normalization rules.

Numbers/units. Accept conversions if numerically equivalent ($\pm 1\%$ rounding for consumer specs) and unit is explicit.

Dates/regions. A claim valid only in {REGION/DATE} is **incorrect** if presented as universal.

Comparatives. If a prompt requests comparison and an answer embeds a brand fact (e.g., “{BRAND} warranty 2 years”), code that fact against the table for the relevant region.

Coder instructions.

1. Read the answer *only* (ignore external links unless explicitly quoted in the answer).
2. Mark each $f \in F$ as C/CI/P/O/H.
3. Add brief rationale for any P/H.
4. Submit codes independently; reconcile disagreements by rulebook.

Reliability.

Double-code $\geq 25\%$ of items; report **Cohen's κ** (binary) or **Krippendorff's α** (multi-category).

Target $\kappa/\alpha \geq .80$ for deployment; revise rulebook if $< .70$.

A3. Computation & Reporting (concise)

Strict CR (primary). Use *binary* mapping ($C, CI \rightarrow 1; P, O, H \rightarrow 0$) in Eq. (2a)/(2b) of §6.2.

Error profile. Report shares of P and H by attribute family (e.g., specs vs. policies).

Sensitivity. Optionally report *soft CR* using the weights in the table (exploratory only).

Documentation. Archive the ground-truth table (with sources), prompt set, engine/agent versions, session seeds, and coder logs.

Implementation note. Keep prompts stable quarter-to-quarter; rotate {CATEGORY}/{SKU} panels to prevent overfitting and mirror your regional mix when setting {REGION}. For AAO interventions, pair CR audits with **Share-of-Answer (SoA)** and **Agent Recommendation Share (ARS)** to trace visibility $\rightarrow$ fidelity $\rightarrow$ selection through the AOJ.