

AI-Supported E-Commerce Management in Jewelry Retail: A Design-Oriented Case Study on ERP, Ad Tracking, and Product Monitoring

Ata Sani Kılıc¹ and Ercan Erkalkan²

¹ The MacDuffie School, Granby, MA, USA

² Department of Computer Technologies, Vocational School of Technical Sciences,
Marmara University, Istanbul, Türkiye
`ercan.erkalkan@marmara.edu.tr`

Abstract. E-commerce operations increasingly depend on the coordinated management of enterprise resource planning (ERP), digital advertising, catalog governance, and product-level performance monitoring. This coordination problem is particularly pronounced in jewelry retail, where visual merchandising, product variety, stock sensitivity, and campaign responsiveness directly affect conversion and margin protection. This study presents a design-oriented case study for an AI-supported e-commerce management framework that integrates ERP records, storefront analytics, ad-tracking signals, and product monitoring into a unified decision-support workflow. The framework combines explainable scoring, anomaly flags, and action-oriented dashboards to connect back-office and growth data. The paper contributes a practical architecture, an interpretable KPI and alert structure, and an evaluation protocol suitable for small and medium-sized digital retailers.

Keywords: e-commerce management · jewelry retail · enterprise resource planning · ad tracking · product monitoring · artificial intelligence · decision support

1 Introduction

E-commerce management has evolved from a storefront-centered activity into a coordination problem spanning product information management, inventory visibility, campaign execution, customer interaction monitoring, and post-order control. In practice, these layers are frequently distributed across separate tools: ERP systems store stock, order, and product master data; ad platforms expose impressions, clicks, and spend; and storefront analytics capture browsing, cart, and purchase behavior. When these systems are weakly integrated, decisions become delayed, inconsistently justified, and difficult to audit.

The challenge becomes sharper in jewelry retail. Product presentation strongly influences buying behavior, collections change frequently, and small mismatches between promotion, availability, and product visibility can create lost revenue or customer dissatisfaction. Publicly visible evidence from

MIONZA indicates a continuing family jewelry tradition and large cumulative sales volume, which together suggest a non-trivial need for synchronized catalog, campaign, and operational governance [1]. The same storefront’s privacy policy states that device information, cookies, and log files are collected, implying the availability of digitally traceable interaction signals for operational analytics and product monitoring [2].

AI is increasingly used in e-commerce for recommendation, customer-service automation, personalization, review analytics, and marketing optimization [3,4,6,5]. In many small and medium-sized retail settings, however, the central problem is not model scarcity but the absence of an interpretable workflow linking product, campaign, stock, and content signals.

This paper develops an AI-supported e-commerce management framework for jewelry retail with three operational pillars: ERP integration, advertisement tracking, and product monitoring. The study is design-oriented and case-driven, with emphasis on implementation feasibility and managerial explainability.

The contributions are fourfold:

1. A domain-grounded architecture for integrating ERP, storefront, and campaign data around product-centered entity resolution.
2. An explainable KPI and scoring layer for stock-aware promotion planning, anomaly detection, and product prioritization.
3. An operational workflow that turns analytics into campaign, content, and stock actions.
4. An evaluation protocol for future deployment studies in digital retail.

The remainder of the paper is organized as follows. Section 2 positions the study within prior work. Section 3 describes the case context and design requirements. Section 4 presents the framework and analytical logic. Section 5 provides illustrative operational scenarios. Section 6 defines the evaluation protocol. Section 7 discusses implications and limitations. Section 8 concludes the paper.

2 Related Work and Positioning

AI-enabled e-commerce research has expanded rapidly across recommendation, customer interaction, marketing analytics, and enterprise integration. A review by Pillarisetty and Mishra highlights how AI tools in online fashion retail affect personalization, engagement, and customer experience, but also shows that many deployments are organized around front-end interaction rather than end-to-end operational control [3]. Chau et al. similarly show that AI-powered customer-service functions shape user experience and decision-making behavior in e-commerce settings [4]. These studies establish the strategic relevance of AI, yet they do not fully resolve how managerial operations should connect ad efficiency, stock data, and product monitoring.

Recommendation research is highly mature. Alhijawi et al. survey recommender-system objectives, evaluation methodology, and solution families, while Boka et al. review sequential recommendation techniques that capture evolving user

preferences over time [5,6]. Nguyen et al. demonstrate that retrieval-driven personalized recommendation can improve engagement and ranking quality in e-commerce environments [7]. These works are valuable for product suggestion and personalization, but the managerial need addressed here is broader: ad efficiency, stock readiness, content adequacy, and campaign-product alignment must be coordinated jointly.

A second stream of literature concerns attribution and advertising effectiveness. Kannan et al. define attribution modeling as the allocation of credit across multiple touchpoints in a customer journey [8]. Anderl et al. advance the field by modeling online attribution with graph-based Markov structures [9]. Ben Mrad and Hnich further argue that more intelligent attribution logic is required to improve trust and actionability in digital marketing decisions [10]. For a retail manager, however, attribution is only one layer of actionability. Even well-attributed campaign performance remains insufficient if promoted products are weakly converting, out of stock, poorly described, or misaligned with catalog priorities.

A third stream studies retail analytics and enterprise systems. Brackmann et al. identify 21 machine-learning application areas in retail, emphasizing decision-oriented and economically operative tasks [11]. Micus et al. review methods for analyzing customer usage data in product decision processes [12]. In enterprise systems, ERP integration and post-implementation quality remain central issues [13,14,15]. ERP responsiveness is further strengthened when big-data contextualization and integration capability are addressed explicitly [16]. Complementary evidence on ERP-CRM integration and omni-channel retail shows that value is generated when back-office systems and customer-facing processes are coordinated rather than treated as independent silos [17,18].

This study is positioned at the intersection of these streams. The proposed contribution is a design-oriented framework that unifies ERP integration, ad tracking, and product monitoring into a single explainable management workflow for a jewelry retail case.

3 Case Context and Design Requirements

The case context considered in this paper is a jewelry e-commerce environment with a live public storefront, diversified product catalog, digital advertising activity, and back-office operational requirements. MIONZA publicly presents itself as a women-owned fine jewelry brand continuing a family tradition since 1980 and reports more than 28,000 pieces sold [1]. These public signals motivate the design requirements defined next.

3.1 R1: ERP-Centered Consistency

Product identifiers, stock quantities, prices, returns, and order states should remain consistent across storefront and back-office records. Prior ERP literature shows that integration quality and post-implementation design decisions

strongly influence downstream operational value [14,15,13]. For retail use, this requirement means that campaign-facing product records should never drift from operational truth.

3.2 R2: Campaign-to-Product Traceability

Campaign spend must be analyzable at product or product-group level whenever technically possible. Attribution literature demonstrates that customer journeys are multi-touch and that credit assignment requires structured modeling [8,9,10]. In a retail management setting, this requirement implies that impressions and clicks should be linked to landing pages, products, or collections in a way that supports budget correction and content review.

3.3 R3: Product-Level Monitoring

Managers need more than store-level KPIs. They need to know which items are visible but not converting, which items convert well but lack exposure, and which promoted items are operationally risky because of low stock, high returns, or weak content. Research on customer-usage analytics and ML application areas in retail supports this product-centered analytical framing [11,12].

3.4 R4: Explainable AI Assistance

The AI layer should produce interpretable rankings, flags, and recommendations rather than opaque black-box outputs. For small and medium-sized retailers, trust in the decision pipeline is often more important than marginal model complexity [10,16]. Therefore, the framework privileges explainable scores, threshold rules, and evidence-linked alerts.

3.5 R5: Actionability for Daily Operations

The output must support actions such as pausing campaigns, boosting selected products, revising product content, or adjusting stock-aware promotion plans. This requirement positions the system as an operational decision-support layer rather than a passive analytics dashboard.


Table 1 maps these requirements to capabilities provided by the proposed framework.

4 Proposed Framework

The proposed framework consists of four coordinated layers: data integration, analytics, AI-supported monitoring, and management action support. Figure 1 shows the overall architecture.

Table 1. Mapping of design requirements to framework capabilities

Req.	Operational concern	Framework capability
R1	ERP–storefront sync	Shared product key, entity resolution, stock and order synchronization
R2	Campaign traceability	Campaign landing-page linkage, product-group mapping, attribution-ready KPI panel
R3	Product monitoring	Visibility, conversion, stock risk, and content adequacy indicators
R4	Explainable AI assistance	Score-based ranking, threshold flags, auditable alert logic
R5	Daily actionability	Dashboard outputs tied to campaign, content, and stock interventions

**Fig. 1.** Proposed AI-supported e-commerce management architecture linking ERP, storefront, ad-platform, analytics, AI monitoring, governance, and decision-support layers.

4.1 Data Integration and Entity Resolution

The first layer harmonizes ERP data, storefront behavior data, and advertising signals into a product-centered analytical model. The core design choice is explicit entity resolution. Each product is associated with a stable internal key used to connect SKU-level ERP records, public product pages, campaign landing destinations, and monitoring outputs. This decision is consistent with ERP-integration literature emphasizing technical and business integration as prerequisites for downstream value [13,17].

The integration layer performs four functions:

1. **Identifier mapping:** reconciliation of SKUs, variant identifiers, category labels, and campaign landing entities.
2. **Temporal alignment:** synchronization of ERP, storefront, and advertising timestamps into common reporting windows.
3. **Quality checks:** detection of missing prices, duplicated SKUs, negative stock anomalies, or unlinked campaign destinations.

Table 2. Illustrative KPI families used by the framework

KPI family	Illustrative metrics	Primary use
Traffic/exposure	Sessions, product views, impressions, click-through rate	Visibility assessment
Commercial performance	Conversion rate, revenue per visit, return on ad spend, cost per acquisition	Campaign/product efficiency
Inventory readiness	Stock coverage, stock-out frequency, reorder horizon	Promotion risk control
Content quality proxy	Bounce rate, low-engagement pages, image/title completeness, add-to-cart drop-off	Content revision priority
Operational exceptions	Unlinked campaigns, stale prices, inconsistent categories, abrupt KPI shifts	Anomaly review

4. **Evidence logging:** storage of source provenance so that dashboard outputs remain auditable.

4.2 Analytics Layer

The analytics layer computes operational and commercial indicators from the integrated dataset. Representative metrics include traffic volume, add-to-cart rate, conversion rate, ad click-through rate, cost per acquisition, return on ad spend, stock coverage, and product contribution to sales. The purpose is not merely reporting. The purpose is to construct comparable product-level evidence for ranking and intervention.

Table 2 summarizes the KPI families used in the framework.

4.3 Explainable AI-Supported Monitoring

The AI layer supports prioritization, anomaly detection, and recommendation. In contrast to opaque end-to-end automation, the proposed design uses explainable score composition to preserve managerial trust and operational auditability.

For each product i , a *promotion readiness score* (PRS) can be defined as

$$\begin{aligned} \text{PRS}_i = & w_1 z(\text{CVR}_i) + w_2 z(\text{ENG}_i) \\ & + w_3 z(\text{MARGIN}_i) + w_4 z(\text{STOCK}_i) \\ & - w_5 z(\text{RET}_i) \end{aligned} \quad (1)$$

where $z(\cdot)$ denotes normalized scaling, CVR is conversion rate, ENG denotes engagement, MARGIN is margin or profitability proxy, STOCK represents stock readiness, and RET captures return-risk or post-purchase friction. Products with higher PRS values are stronger candidates for promotion.

Similarly, a *campaign-product alert score* (CAS) can be defined as

$$\begin{aligned} \text{CAS}_{ij} = & a_1 z(\text{SPEND}_{ij}) + a_2 z(\text{IMP}_{ij}) \\ & - a_3 z(\text{REV}_{ij}) - a_4 z(\text{STOCK}_i) \\ & + a_5 z(\Delta_{ij}) \end{aligned} \quad (2)$$

where SPEND_{ij} , IMP_{ij} , and REV_{ij} correspond to spend, impressions, and attributed revenue for product i under campaign j , and Δ_{ij} represents a deviation term against historical baseline. When $\text{CAS}_{ij} > \tau$, the pair is flagged for managerial review.

The scores in Eqs. 1 and 2 are not intended as immutable formulas. They are implementation templates that make decision logic explicit, adjustable, and inspectable. This explainable design is consistent with the operational emphasis found in AI-enabled retail and attribution literature [11,10].

Three monitoring functions are emphasized:

- **Promotion candidate ranking:** identify high-potential products with sufficient stock and favorable behavioral evidence.
- **Underperformer detection:** surface items or campaign-product pairs consuming spend without acceptable downstream performance.
- **Catalog health review:** detect products with strong visibility but weak click-through, strong traffic but weak conversion, or content-related friction indicators.

4.4 Governance and Management Action Support

The governance layer stores thresholds, weighting configurations, and explanation traces. The action layer transforms monitored evidence into decisions such as pausing spend on weak campaign-product pairs, revising product pages, boosting well-performing items, or excluding low-stock products from active promotion. This design aligns with the practical need to transform analytics into action rather than to produce dashboards without intervention logic.

Figure 2 summarizes the analytical procedure.

5 Illustrative Operational Scenarios

This section presents three illustrative scenarios showing how the framework supports routine managerial reasoning without claiming measured effects from undisclosed proprietary data.

5.1 Scenario 1: High Visibility, Weak Conversion

Suppose that a necklace category receives strong impressions and page traffic after a paid social campaign, yet add-to-cart and purchase conversion remain weak. The analytics layer identifies a large visibility–conversion gap, while the AI layer keeps the campaign-product pair on the alert list because spend is increasing faster than attributed revenue. The manager can then inspect whether the issue is caused by weak product titles, misleading imagery, pricing friction, or poor landing-page-product match.




Fig. 2. Analytical procedure used in the proposed decision-support workflow.

5.2 Scenario 2: Good Conversion, Insufficient Exposure

A ring product may show strong organic conversion and high engagement but receive limited paid exposure. In this case, a high PRS indicates strong promotion readiness. The system recommends prioritizing the product or its category group in upcoming campaigns, while the stock-readiness term prevents promotion if coverage falls below the defined threshold.

5.3 Scenario 3: Campaign Efficiency Blocked by Operational Risk

A bracelet collection may appear attractive from an ad-efficiency viewpoint but carry low inventory, stale price synchronization, or repeated return signals. The system should not rank such items as high-priority promotion candidates without qualification. Instead, governance rules downgrade the candidate until operational risks are resolved.

These scenarios reflect a central design principle: e-commerce management should not reward isolated channel metrics if operational readiness is insufficient.

6 Evaluation Protocol

The framework is designed to be evaluated as a design-and-implementation study in a live or pilot retail environment. Because the present paper does not disclose proprietary internal datasets, the evaluation section specifies how future empirical work should be conducted.

6.1 Research Questions

The empirical evaluation should address the following questions:

- **RQ1.** To what extent can ERP, storefront, and ad-platform data be integrated into a product-centered management view with minimal linkage loss?
- **RQ2.** Can explainable alerting improve managerial visibility into inefficient campaign-product pairs and underexposed products?
- **RQ3.** Does the framework reduce operational friction such as promoting low-stock items, acting on stale data, or overlooking content-related conversion problems?

6.2 Data Categories

A future deployment study should collect four categories of data:

1. ERP data: product master records, stock levels, prices, orders, returns.
2. Storefront data: views, sessions, add-to-cart events, transactions, referrers.
3. Advertising data: impressions, clicks, spend, campaign identifiers, attributed conversions.
4. Monitoring metadata: alert logs, manager interventions, overridden recommendations, and review timestamps.

Table 3. Suggested evaluation dimensions for a future deployment study

Dimension	Illustrative measures
Integration quality	share of products with complete ERP–storefront–campaign linkage; synchronization latency; missing-value rate
Alert usefulness	percentage of alerts judged actionable; precision of flagged underperformers; time-to-review
Promotion governance	share of promoted products with acceptable stock coverage; reduction in low-stock promotions
Managerial efficiency	dashboard review time; intervention rate; number of manual cross-tool checks avoided
Commercial utility	conversion lift of prioritized products; spend reduction on paused underperformers; change in return on ad spend

6.3 Evaluation Dimensions and Metrics

The evaluation should combine integration quality, analytical utility, and managerial actionability. Table 3 summarizes suggested dimensions.

6.4 Analytical Procedure

A practical deployment can follow the six-step flow shown earlier in Fig. 2: data extraction, entity resolution, KPI computation, anomaly detection, ranking, and managerial review. Weekly baselines or rolling windows should be used to distinguish normal fluctuation from genuine anomalies. Where sufficient historical data exist, more advanced models such as sequential recommendation or uplift scoring can be introduced. Where data remain limited, the explainable score-based approach remains appropriate.

6.5 Threats to Validity

Several validity threats should be anticipated: imperfect platform attribution, simultaneous content and campaign changes, sparse data for low-volume products, and evolving operational thresholds during deployment. These issues should be documented explicitly during empirical testing.

7 Discussion

The immediate value of AI in e-commerce operations may lie less in end-user novelty and more in managerial coherence. By linking ERP, storefront, and ad-platform evidence at product level, the proposed framework transforms fragmented observations into operationally useful intelligence. Although motivated by jewelry retail, the same design can be adapted to other catalog-driven categories such as fashion, beauty, and accessories.

The paper has three main limitations. First, it does not report proprietary quantitative outcomes because the current objective is to provide a transparent

design-and-implementation blueprint without overstating undisclosed evidence. Second, the case context is partially motivated by public storefront information, which is valuable for design grounding but insufficient for causal performance claims. Third, the scoring equations are intentionally modular and configurable; they should be calibrated during real deployment rather than treated as universal constants.

Future work should implement the framework on live operational data and report quantitative results on linkage quality, alert precision, conversion lift for prioritized products, and waste reduction in paid campaigns. Additional extensions include segment-aware promotion logic, explainable large-language-model assistance for content revision, and controlled campaign experiments.

8 Conclusion

This paper presented an AI-supported e-commerce management framework for jewelry retail that integrates ERP data, ad-tracking signals, and product monitoring into an explainable decision-support workflow. The contribution is design-oriented and operational rather than purely algorithmic. The framework combines product-centered data integration, KPI computation, score-based monitoring, and action-oriented governance to support stock-aware promotion, anomaly detection, and product prioritization. The proposed structure addresses a practical gap between isolated analytics tools and the daily managerial need for interpretable, coordinated, and auditable retail intelligence. The framework is sufficiently concrete for pilot implementation and sufficiently modular for extension in future empirical studies.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. MIONZA Jewelry: About the brand. <https://mionza.com/pages/mionza-fine-jewelry-about-us>. accessed 24 Apr 2026
2. MIONZA Jewelry: Privacy policy. <https://mionza.com/pages/privacy-policy>. accessed 24 Apr 2026
3. Pillarisetty, R., Mishra, P.: A review of AI (Artificial Intelligence) tools and customer experience in online fashion retail. *International Journal of E-Business Research* **18**(2), 1–12 (2022). <https://doi.org/10.4018/IJEER.294111>
4. Chau, H.K.L., Ngo, T.T.A., Bui, C.T., Tran, N.P.N.: Human-AI interaction in e-commerce: The impact of AI-powered customer service on user experience and decision-making. *Computers in Human Behavior Reports* **19**, 100725 (2025). <https://doi.org/10.1016/j.chbr.2025.100725>
5. Alhijawi, B., Awajan, A., Fraihat, S.: Survey on the objectives of recommender systems: Measures, solutions, evaluation methodology, and new perspectives. *ACM Computing Surveys* **55**(5), 93 (2023). <https://doi.org/10.1145/3527449>

6. Boka, T.F., Niu, Z., Neupane, R.B.: A survey of sequential recommendation systems: Techniques, evaluation, and future directions. *Information Systems* **125**, 102427 (2024). <https://doi.org/10.1016/j.is.2024.102427>
7. Nguyen, D.-N., Nguyen, V.-H., Trinh, T., Ho, T., Le, H.-S.: A personalized product recommendation model in e-commerce based on retrieval strategy. *Journal of Open Innovation: Technology, Market, and Complexity* **10**, 100303 (2024). <https://doi.org/10.1016/j.joitmc.2024.100303>
8. Kannan, P.K., Li, H.A., Sridhar, S.: The path to purchase and attribution modeling: Introduction to special section. *International Journal of Research in Marketing* **33**(3), 449–456 (2016). <https://doi.org/10.1016/j.ijresmar.2016.04.001>
9. Anderl, E., Becker, I., von Wangenheim, F., Schumann, J.H.: Mapping the customer journey: Lessons learned from graph-based online attribution modeling. *International Journal of Research in Marketing* **33**(3), 457–474 (2016). <https://doi.org/10.1016/j.ijresmar.2016.03.001>
10. Ben Mrad, A., Hnich, B.: Intelligent attribution modeling for enhanced digital marketing performance. *Intelligent Systems with Applications* **21**, 200337 (2024). <https://doi.org/10.1016/j.iswa.2024.200337>
11. Brackmann, C., Huetsch, M., Wulfert, T., et al.: Identifying application areas for machine learning in the retail sector: A literature review and interview study. *SN Computer Science* **4**, 426 (2023). <https://doi.org/10.1007/s42979-023-01888-w>
12. Micus, C., Schramm, S., Boehm, M., Krcmar, H.: Methods to analyze customer usage data in a product decision process: A systematic literature review. *Operations Research Perspectives* **10**, 100277 (2023). <https://doi.org/10.1016/j.orp.2023.100277>
13. Tobon Valencia, E., Lamouri, S., Pellerin, R., Dubois, P., Moeuf, A.: The integration of ERP and inter-intra organizational information systems: A literature review. *IFAC-PapersOnLine* **51**(11), 1212–1217 (2018). <https://doi.org/10.1016/j.ifacol.2018.08.425>
14. Zhu, Y., Li, Y., Wang, W., Chen, J.: What leads to post-implementation success of ERP? An empirical study of the Chinese retail industry. *International Journal of Information Management* **30**(3), 265–276 (2010). <https://doi.org/10.1016/j.ijinfomgt.2009.09.007>
15. Malhotra, R., Temponi, C.: Critical decisions for ERP integration: Small business issues. *International Journal of Information Management* **30**(1), 28–37 (2010). <https://doi.org/10.1016/j.ijinfomgt.2009.03.001>
16. Bandara, F., Jayawickrama, U., Subasinghage, M., Olan, F., Alamoudi, H., Alharthi, M.: Enhancing ERP responsiveness through big data technologies: An empirical investigation. *Information Systems Frontiers* **26**(1), 251–275 (2024). <https://doi.org/10.1007/s10796-023-10374-w>
17. Ruivo, P., Mestre, A., Johansson, B., Oliveira, T.: Defining the ERP and CRM integrative value. *Procedia Technology* **16**, 704–709 (2014). <https://doi.org/10.1016/j.protcy.2014.10.019>
18. Hubner, A., Holzapfel, A., Kuhn, H.: Distribution systems in omni-channel retailing. *Business Research* **9**(2), 255–296 (2016). <https://doi.org/10.1007/s40685-016-0034-7>