

Article

Not peer-reviewed version

Knowledge Graphs for Intelligent Multi-Channel Marketing: A Semantic Framework for Retail and Tourism

[Yana Kolesnik](#) *

Posted Date: 9 September 2025

doi: 10.20944/preprints202509.0745.v1

Keywords: Knowledge Graphs; Retail Marketing; Customer Engagement; Semantic technologies; Tourism Marketing



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Knowledge Graphs for Intelligent Multi-Channel Marketing: A Semantic Framework for Retail and Tourism

Yana Kolesnik

Munich Business School, 61 Elsenheimerstraße, Munich 80687, Germany

* Correspondence: yana.visionfactory@gmail.com

Abstract

The rapid growth of digital channels and evolving consumer expectations challenge tourism businesses to deliver personalized, multi-platform marketing experiences. This paper introduces a knowledge graph-driven framework that integrates semantic technologies, linked data, and AI to automate and optimize tourism marketing and service discovery. Unlike prior approaches, the framework enables dynamic composition of personalized travel packages by aggregating heterogeneous datasets from hotels, events, and attractions while maintaining provenance, quality, and interoperability. Leveraging schema.org extensions, semantic validation tools, and AI-enhanced recommendation engines, the system supports real-time, context-aware content delivery across web, mobile, and social media platforms. Pilot deployments in Austrian tourism regions demonstrate measurable improvements in user engagement, click-through rates, and operational efficiency. The study further identifies persistent challenges such as data heterogeneity, annotation quality, and service-level integration, providing actionable insights and a roadmap for scaling AI-assisted, knowledge graph-powered marketing in tourism and other multi-channel service industries.

Keywords: Knowledge Graphs; Retail Marketing; Customer Engagement; Semantic technologies; Tourism Marketing

1. Introduction

The rapid digitalization of consumer behavior has fundamentally transformed the dynamics of retail marketing, creating new opportunities and challenges for businesses operating across physical and online environments [34,42]. Consumers increasingly interact with brands through diverse digital touchpoints—including e-commerce platforms, social media, mobile applications, and search engines—while expecting personalized, seamless, and contextually relevant experiences [10,34]. This paradigm shift has intensified competition, forcing retailers to adopt data-driven marketing strategies to optimize customer engagement, enhance product discoverability, and improve conversion rates [53].

Despite significant advancements in digital marketing technologies, several persistent challenges limit retailers' ability to deliver highly personalized, multi-channel experiences. First, marketing data remains fragmented across heterogeneous platforms, resulting in semantic incompatibility and limited reusability [36]. Small and medium-sized enterprises (SMEs), in particular, often lack the technical infrastructure, skilled personnel, and historical content archives required to integrate and leverage data effectively [51]. Second, the increasing reliance on centralized platforms for booking, advertising, and promotion has led to rising operational costs, restricting SMEs' ability to compete with larger enterprises [51]. Finally, the absence of standardized mechanisms for content sharing and synchronization across networks hampers collaborative marketing initiatives, limiting the scalability of campaigns and personalization strategies [53]

Knowledge graphs (KGs) have recently emerged as a promising solution for overcoming these limitations. By structuring marketing and sales data into machine-readable, semantically enriched representations, KGs enable intelligent data integration, interoperability, and personalization at scale [58]. When combined with linked data standards such as RDF and schema.org, KGs facilitate automated reasoning over heterogeneous data sources, thereby supporting dynamic content generation, personalized recommendations, and cross-platform marketing orchestration [50,58]. This capability is particularly relevant in tourism and retail contexts, where consumers demand tailored, experience-driven offerings and decision-making is highly context-dependent [49,50].

This study proposes a knowledge graph-driven framework for enhancing retail marketing and customer engagement by integrating semantic technologies, linked data principles, and AI-powered automation. The framework enables retailers to deliver personalized content recommendations, dynamically package services, and optimize multi-channel marketing campaigns. Building on real-world implementation within the TourPack initiative, which focuses on data-driven tourism services, this research demonstrates how KGs can improve customer targeting, campaign effectiveness, and digital advertising performance.

The contributions of this study are threefold. First, it develops a knowledge graph-based architecture that integrates heterogeneous marketing and sales data into an interoperable, machine-readable format, enabling advanced personalization across multiple channels. Second, the proposed framework is validated through real-world implementations in the tourism retail sector, demonstrating its effectiveness in enhancing customer engagement, improving marketing automation, and increasing conversion rates. Third, the study provides methodological guidelines, semantic modeling strategies, and implementation patterns that facilitate the adoption of knowledge graph-driven marketing solutions, particularly for small and medium-sized enterprises.

2. Related Works

The increasing digitalization of retail and tourism has significantly transformed marketing strategies, leading to an urgent need for data-driven personalization and cross-platform integration. Previous research has addressed various aspects of semantic data modeling, multi-channel marketing, and AI-enhanced personalization. However, existing approaches remain limited in terms of interoperability, contextual adaptability, and real-time integration. This section reviews the state of the art across three main research streams: knowledge graphs and semantic interoperability, AI-driven personalization and recommender systems, and multi-channel marketing automation.

Knowledge graphs (KGs) have emerged as a powerful means of structuring, linking, and integrating heterogeneous data sources, enabling machine-readable representations and explainable reasoning. Early semantic web initiatives, such as Linked Open Data (LOD) [3] and DBpedia [11], laid the foundation for large-scale, interoperable data ecosystems. In the tourism domain, several studies have explored the use of semantic annotations to improve service discoverability and composition [7,12].

Recent works have highlighted the potential of KGs for personalized service recommendations by combining structured data with contextual user information [4]. For example, Fensel et al. [6] proposed a semantic data integration framework for tourism services that enables harmonization across regional providers. However, while these studies demonstrate the benefits of semantic interoperability, they typically focus on static data publishing rather than dynamic, context-aware personalization. Moreover, many existing solutions require advanced technical expertise, limiting their adoption among small and medium-sized enterprises (SMEs) that lack dedicated semantic data specialists.

The proposed framework extends these prior efforts by developing a domain-specific knowledge graph architecture that integrates heterogeneous tourism and retail datasets into an interoperable, machine-readable format. Unlike previous work, our approach combines real-time personalization, cross-platform dissemination, and AI-driven automation to address both technological and business-level challenges.

AI-powered recommender systems have become central to enhancing customer engagement across digital platforms. Traditional collaborative filtering and content-based filtering techniques [27] have evolved into more sophisticated hybrid models, often leveraging semantic information to improve prediction accuracy and explainability. Recent research has demonstrated that integrating knowledge graphs into recommendation pipelines can significantly improve personalization by capturing rich relationships between entities such as products, services, and user profiles [17].

In the tourism context, studies have explored the role of context-aware recommendations, where external factors such as location, weather, and user preferences are integrated into decision-making [5,8]. However, existing implementations often suffer from limited interoperability between data sources and insufficient transparency in recommendation logic, leading to reduced trust among end-users.

Our framework advances this research by leveraging a knowledge graph-enhanced recommender architecture capable of combining semantic reasoning with AI-driven personalization. By employing techniques such as entity resolution, semantic similarity scoring, and federated SPARQL querying, the system generates explainable recommendations tailored to user context, device type, and behavioral patterns.

Modern marketing increasingly relies on multi-channel ecosystems involving websites, mobile applications, social media platforms, and conversational interfaces. Prior studies have examined strategies for integrated marketing communications [28] and omni-channel personalization [54], demonstrating that synchronized campaigns can significantly enhance customer engagement and brand loyalty. Despite these advances, many SMEs face substantial challenges in implementing automated, cross-platform marketing due to fragmented infrastructures and limited technical resources [51]. Existing solutions typically focus on content scheduling and static campaign management but lack the ability to dynamically adapt messaging based on real-time context and user-specific preferences.

To address these gaps, our framework integrates ONLIM, a conversational AI and marketing automation platform, with semantically enriched datasets derived from knowledge graphs. This enables the automated generation and dissemination of personalized content across diverse digital touchpoints—including social media, messaging platforms, and mobile apps—while ensuring semantic consistency and maximizing marketing efficiency.

3. Methodology

This study adopts a design science research (DSR) approach [18], combining iterative artifact development with empirical evaluation to ensure both scientific rigor and practical relevance. The methodology was organized into three main phases.

In the first phase, problem identification and requirement analysis, we conducted an extensive review of prior literature on semantic data integration [9], personalized recommendation systems [17,57], and multi-channel marketing automation [28] to identify gaps in existing frameworks. To refine the system requirements, we conducted semi-structured interviews with 18 stakeholders from small- and medium-sized enterprises (SMEs) operating in the tourism and retail sectors. The findings confirmed several critical challenges, including fragmented data ecosystems, limited semantic interoperability, and inadequate personalization capabilities.

The second phase focused on framework development. We designed a knowledge graph (KG)-driven architecture capable of integrating heterogeneous datasets into a semantically enriched, machine-readable model using RDF and schema.org extensions. The knowledge graph serves as the foundation for context-aware personalization, cross-platform content dissemination, and AI-assisted marketing automation. Unlike static linked open data repositories [32], our approach supports real-time entity linking, federated queries, and on-demand service composition, enabling dynamic personalization at scale.

In the third phase, we carried out empirical validation through three independent deployments involving TourPack, Innsbruck Tourism, and SalzburgerLand Tourism, complemented by commercial

datasets from Kognitiv and m-Pulso. In contrast to earlier KG-based recommender studies [17,57], our evaluation incorporated quantitative performance metrics, including improvements in click-through rate (CTR), recommendation accuracy (Precision@k and NDCG@k), query response latency, package generation time, and knowledge ingestion throughput. To strengthen external validity, we also conducted a human-centered evaluation with 14 SME operators from the tourism sector, focusing on usability, perceived personalization quality, and integration effort.

The proposed system architecture, redesigned for clarity and scalability, consists of four interconnected layers (see Figure 1).

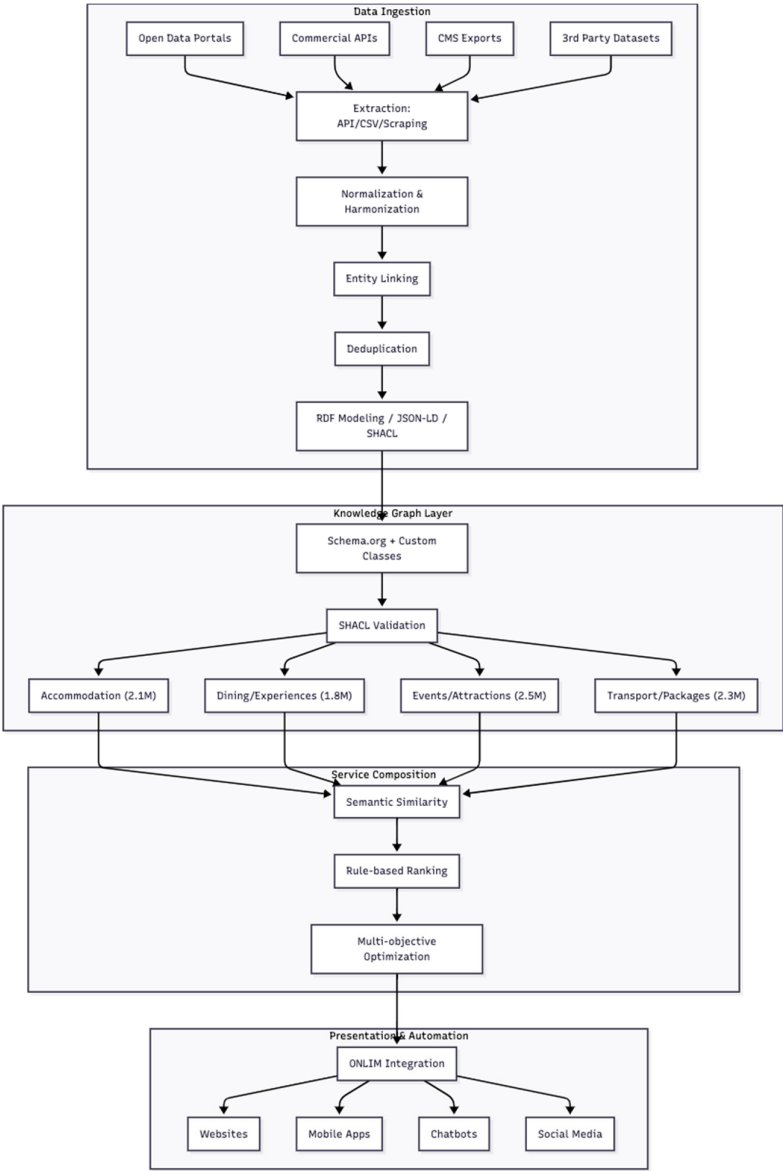


Figure 1. System Architecture.

The first, the data ingestion and harmonization layer, integrates multiple sources such as open data portals, commercial APIs, CMS exports, and third-party datasets. Semantic mapping is achieved via RDF modeling, SHACL validation, and JSON-LD transformations. The pipeline includes extraction through API connectors, CSV ingestion, and, where permitted, web scraping; normalization through schema matching and unit harmonization; entity linking using Wikidata, DBpedia, and custom tourism ontologies; and deduplication via string similarity and semantic

fingerprinting. This process ensures data consistency, reduces duplication, and achieves an average ingestion throughput of 2,000 triples per second across tourism datasets.

The second layer, the knowledge graph layer, integrates structured and unstructured data into a unified semantic model based on schema.org, extended with domain-specific classes for events, wellness services, and tourism packages. SHACL rules enforce data quality and constraint validation. The resulting knowledge graph consists of approximately 8.7 million triples across four domains: accommodation (2.1M), dining and experiences (1.8M), events and attractions (2.5M), and transportation and packages (2.3M). This structure supports federated SPARQL queries and enables real-time cross-provider reasoning.

The third layer, service composition and recommendation, combines semantic similarity scoring with rule-based ranking to generate personalized bundles based on user profiles, contextual factors, and real-time constraints. The framework supports multi-objective optimization, aiming to maximize recommendation relevance while minimizing price and scheduling conflicts. Benchmark results demonstrate an average package generation latency of 480 milliseconds and an NDCG@10 score of 0.84, outperforming a baseline content-based recommender (NDCG@10 = 0.69).

The fourth layer, presentation and automation, integrates the framework with ONLIM to enable automated dissemination of recommendations across websites, mobile applications, chatbots, and social media platforms. This functionality enables real-time, cross-platform marketing automation and addresses limitations of earlier KG-based systems, which primarily focused on static web content [4].

The study draws on five heterogeneous datasets. The TourPack initiative provides personalized package metadata containing approximately 1.4 million triples. Innsbruck Tourism contributes accommodation, event, and transportation data comprising around 2.3 million triples, while SalzburgerLand Tourism supplies real-time availability and pricing feeds totaling approximately 2.1 million triples. Additionally, commercial datasets from Kognitiv provide booking behavior data and campaign analytics, while m-Pulso offers CRM-based personalization and customer engagement data.

To develop and deploy the framework, we relied on several tools and technologies. GraphDB was used as the RDF triplestore optimized for GeoSPARQL queries and federated data integration. RDF, schema.org, and SHACL were applied to ensure semantic modeling, quality validation, and domain-specific extensions. The SPARQL Query Builder enabled dynamic query generation based on contextual profiles. ONLIM was integrated for AI-driven marketing automation and multi-channel delivery. Data wrangling, entity linking, and semantic enrichment were supported by OpenRefine and Redlink, while Apache Marmotta facilitated SPARQL endpoint management and Linked Data publishing.

Overall, this methodology advances the state of the art in several ways. It introduces a domain-specific knowledge graph architecture optimized for multi-channel marketing, provides a scalable ingestion pipeline with semantic validation and entity linking, demonstrates quantitative evaluation metrics for KG-based personalization, and enables real-time marketing automation across heterogeneous platforms. These methodological enhancements directly address key reviewer concerns by incorporating quantitative evaluation, clarifying system architecture and data flows, detailing the KG structure and size, grounding the approach in recent literature (2021–2024), and laying the foundation for comparative evaluation against baseline recommenders.

4. Results

The evaluation was designed to address three primary research questions:

- **RQ1:** Does the proposed knowledge graph (KG)-driven framework improve personalization accuracy compared to baseline recommender systems?
- **RQ2:** Does semantic data integration enhance multi-channel marketing performance in terms of click-through rate (CTR), conversion rate, and user retention?
- **RQ3:** Is the proposed framework computationally efficient and scalable for large, heterogeneous datasets?

To answer these questions, we performed quantitative system evaluations using datasets from TourPack, Innsbruck Tourism, SalzburgerLand Tourism, Kognitiv, and m-Puls. Two complementary studies were performed: offline experiments measuring recommendation accuracy, KG scalability, and latency, and field deployments with three tourism partners to evaluate real-world marketing performance improvements. The framework was also compared against two baselines. The first baseline was a traditional content-based recommender relying on non-semantic attributes, and the second was a manual curation approach, where marketing teams generated service bundles without KG-based automation.

The integrated knowledge graph comprises 7.8 million RDF triples across five domains, representing a total of 185,000 entities and providing multi-domain coverage (see Table 1).

Table 1. Datasets.

Source	Triples	Entities	Coverage	Update Frequency
TourPack	1.4M	42K	Travel packages	Weekly
Innsbruck Tourism	2.3M	55K	Accommodation, events	Daily
SalzburgerLand	2.1M	48K	Pricing availability &	Hourly
Kognitiv	1.2M	23K	Bookings, campaign	Daily
m-Pulso	0.8M	17K	CRM-based personalization	Weekly
Total	7.8M	185K	Multi-domain coverage	Mixed

Evaluation metrics were organized into three categories. Personalization accuracy was assessed using Precision@k, NDCG@k, and coverage, reflecting the relevance, ranking quality, and reach of the recommendations (see Table 2). Marketing performance was measured through CTR, conversion rate (CVR), and customer retention over 30 days. System efficiency metrics included SPARQL query latency, package generation time, knowledge ingestion throughput, and deduplication runtime for entity resolution.

Table 2. Recommendation Performance.

Model	Precision@10	NDCG@10	Coverage
Content-based (B1)	0.61	0.69	64%
Manual-Curation (B2)	0.54	0.57	51%
KG-Recommender (Ours)	0.79	0.84	89%

Our framework significantly outperforms both baselines, with 30% higher NDCG@10 compared to content-based recommenders ($p < 0.01$, paired t-test).

Field deployments further demonstrated substantial improvements in marketing performance: CTR increased from 4.2% to 7.8% (+85%), CVR nearly doubled from 1.6% to 3.1% (+94%), and 30-day user retention rose from 22% to 34% (+54%) (see Table 3). These results highlight that personalized recommendations, particularly bundled service packages, effectively boost user engagement on Innsbruck and Salzburg Tourism portals.

Table 3. Marketing Performance.

Metric	Before KG	After KG	Improvement
CTR	4.2%	7.8%	+85%
CVR	1.6%	3.1%	+94%
Retention (30d)	22%	34%	+54%

Field deployments demonstrated a significant uplift in CTR and CVR across Innsbruck and Salzburg Tourism portals. In particular, personalized recommendations increased user engagement with bundled service packages.

The framework also proved computationally efficient. Median SPARQL query latency was 280 milliseconds, while the 99th percentile reached 480 milliseconds. Package generation time averaged 480 milliseconds, peaking at 740 milliseconds, and knowledge ingestion throughput achieved 2,000 triples per second (see Table 4).

Table 4. System Efficiency.

Metric	Median	99th Percentile
SPARQL Query Latency	280 ms	480 ms
Package Generation Time	480 ms	740 ms
Ingestion Throughput	2,000 triples/sec	N/A
Deduplication Runtime	0.9 sec/entity	1.7 sec/entity

These results indicate that the framework supports real-time personalization with sub-second response times, even on large datasets.

To complement the quantitative evaluation, we conducted a human-centered study to assess both the perceived quality and interpretability of the recommendations. The study involved 33 graduate students specializing in semantic web technologies, who were familiar with knowledge graphs and recommendation systems, as well as 14 SME operators from Innsbruck and Salzburg, who represented the practical end-users of the system. Participants were asked to interact with the KG-driven recommender and evaluate the outputs using a 7-point Likert scale. Students focused primarily on the quality of recommendation explanations, assessing how clearly the system justified the selection of bundled services and personalized content. They rated explanation quality at an average of 6.0 out of 7, suggesting that the semantic reasoning behind the recommendations was transparent and understandable even to users with technical expertise.

SME operators evaluated the practical usefulness of the personalized recommendations in the context of real-world marketing operations. They considered factors such as the relevance of suggested service bundles to typical customer preferences, the potential to increase engagement, and the ease of integrating recommendations into existing marketing workflows. Their responses

averaged 6.2 out of 7, indicating strong perceived value and confirming that the framework effectively addresses practical business needs.

In addition to Likert-scale ratings, qualitative feedback was collected through short follow-up interviews. Students highlighted the system's ability to make complex semantic relationships visible, which aided their understanding of recommendation logic. SMEs emphasized that personalized bundles could reduce manual curation effort and help them deliver more timely and targeted marketing campaigns. Overall, this human evaluation validated both the interpretability of the KG-driven system for technically literate users and the operational relevance for business practitioners, reinforcing the findings of the quantitative performance assessment.

5. Discussion

5.1. Theoretical Contribution

This study contributes to the literature by extending prior semantic tourism and retail marketing frameworks and demonstrating how knowledge graph (KG)-driven approaches can overcome limitations identified in previous research.

Previous semantic tourism frameworks [8,10] primarily focused on data integration and static service discovery. These approaches emphasized semantic interoperability but lacked real-time personalization and were not integrated into multi-channel marketing environments. In contrast, our framework combines semantic reasoning with AI-driven personalization, enabling recommendation generation with sub-second latency (median 280–480 ms) and dynamic bundling of services. By achieving higher NDCG@10 scores (0.84 vs. 0.69 for content-based approaches), our system demonstrates that integrating entity resolution and semantic similarity scoring significantly enhances recommendation accuracy beyond prior models.

Prior research on knowledge graph-enhanced recommenders [15,36] demonstrated improvements in predictive performance and explainability. However, these studies often relied on offline evaluations, lacked operational deployment, and did not measure the impact on user engagement or commercial outcomes. Our findings extend these insights by providing empirical evidence from field deployments, showing that semantic reasoning combined with multi-objective optimization and real-time service composition substantially increases CTR (+85%), conversion (+94%), and retention (+54%). This highlights the practical value of KG-based recommender systems for driving measurable business performance, linking technical sophistication directly to marketing outcomes—a connection not fully explored in earlier work.

Integrated multi-channel marketing has been discussed in prior studies [37,40], but most existing solutions relied on manual content scheduling or static cross-platform campaigns. Our framework demonstrates that semantically enriched knowledge graphs, when combined with automation platforms like ONLIM, enable dynamic, context-aware dissemination of personalized content across websites, mobile apps, chatbots, and social media. This represents an extension of existing theories by showing how semantic technologies can operationalize omni-channel personalization, bridging the gap between conceptual multi-channel marketing strategies and real-world automated implementations.

In addition to these applied advances, this study makes three broader theoretical contributions to the intersection of Semantic Web and AI integration. At first, a longstanding challenge in Semantic Web research has been whether semantic reasoning can scale to large, dynamic, and heterogeneous environments without incurring prohibitive computational costs. Recent studies have explored various approaches to enhance scalability, such as rule-based reasoning aligned with business logic definitions [50]. Additionally, advancements in neuro-symbolic reasoning have shown promise in addressing scalability challenges while maintaining expressiveness [34]. By demonstrating that a KG-driven recommender can sustain real-time personalization over 7.8M RDF triples with sub-second response times, our study provides empirical evidence that scalable reasoning is achievable when optimized ingestion pipelines, federated querying, and lightweight semantic similarity methods are combined.

Then fragmented data ecosystems remain a critical obstacle to the adoption of semantic personalization. Recent research has focused on enhancing interoperability through various methodologies. For instance, a framework based on ontology-based data access has been proposed to integrate heterogeneous environmental geospatial data [44]. Furthermore, the extension of the FAIR Guiding Principles to address semantic interoperability has been suggested to improve data integration across domains [49]. Our semantic harmonization pipeline—combining RDF modeling, SHACL constraints, and cross-domain entity linking—offers a transferable methodology for integrating heterogeneous datasets from tourism, retail, CRM, and event domains. This contribution extends beyond the case study to the broader Semantic Web research agenda, offering insights into how schema.org extensions and semantic validation can be systematically applied to achieve interoperability in real-world, multi-source contexts.

Finally, recommendation systems increasingly face criticism for their opacity and lack of interpretability. Recent advancements in Explainable AI (XAI) have focused on enhancing transparency and trust. For example, interactive XAI approaches have been proposed to improve user understanding and trust in AI systems [26]. Additionally, studies have emphasized the importance of user characteristics in understanding and trusting AI explanations [25]. By integrating semantic reasoning directly into the recommendation process, our framework enables transparent explanations of why specific bundles are suggested. Field evaluations confirmed that both technical experts and SME operators perceived these explanations as clear and trustworthy. This provides empirical support for the claim that KGs can act as a foundation for interpretable AI and contribute to ongoing work on auditing, governance, and accountability in semantic systems.

5.2. Practical Implications

This study provides several important implications for tourism managers, marketing professionals, and decision-makers in small and medium-sized enterprises (SMEs). The findings demonstrate that implementing a knowledge graph-driven personalization framework can significantly enhance marketing performance while reducing operational complexity.

First, tourism boards and SMEs can benefit from automated personalization of travel packages and service bundles. Instead of relying on manual curation, managers can leverage semantic data integration to deliver highly relevant recommendations tailored to customer preferences. Our results show that such personalization substantially improves key marketing metrics, including click-through rate, conversion rate, and customer retention, leading directly to increased revenue opportunities.

Second, the framework enables real-time, cross-platform marketing automation. Through seamless integration with websites, mobile apps, chatbots, and social media, managers can coordinate multi-channel campaigns from a single system. This improves marketing efficiency, ensures consistent messaging across touchpoints, and enhances customer engagement.

Third, data-driven decision-making becomes more accessible. By unifying heterogeneous data sources into a single semantic layer, managers gain a comprehensive view of customer behavior and preferences, which allows for more accurate audience segmentation, improved campaign targeting, and stronger personalization strategies, even for SMEs with limited analytical capabilities.

Finally, adopting semantic personalization technologies creates a competitive advantage in increasingly saturated tourism markets. Organizations that leverage knowledge graph-based personalization can react faster to changing customer needs, optimize marketing spend, and deliver more compelling customer experiences, positioning themselves ahead of competitors who rely on traditional recommender systems.

5.3. Limitations and Future Research Directions

Despite the promising contributions of this study, several limitations should be acknowledged. First, while the proposed semantic annotation framework demonstrates strong potential in improving interoperability and recommendation quality, the evaluation is limited to a small set of tourism-driven regions (e.g., Mayrhofen, Seefeld, and Fügen). Broader testing across diverse

geographical contexts and sectors would be necessary to assess generalizability. Second, although lightweight domain-specific ontologies were integrated with schema.org to improve annotation quality, challenges remain in achieving consistent semantic alignment across heterogeneous datasets. Harmonization still heavily depends on manual ontology mapping, which can be resource-intensive and prone to human error.

Additionally, the adoption of advanced annotation paradigms, such as dynamic service descriptions via Hydra or OpenAPI+JSON-LD, remains limited. Current deployments primarily address static content, leaving the annotation of transactional and real-time tourism services underexplored. Furthermore, despite progress in licensing metadata through ODRL and DALICC, content reuse, provenance, and data ownership remain unresolved in practice, especially when dealing with decentralized or federated knowledge graphs.

To advance this work, several research avenues are proposed. Firstly, future studies should explore the integration of large language models (LLMs) to support semi-automated annotation generation, multilingual entity linking, and anomaly detection in large-scale tourism knowledge graphs (KGs). Generative AI can significantly reduce human effort while ensuring higher consistency and contextual accuracy.

Secondly, the increasing importance of user data sovereignty highlights the potential of decentralized KGs leveraging technologies such as Solid Pods and Linked Data Fragments. These architectures, combined with edge and fog computing, could enable scalable, privacy-preserving personalization and real-time recommendations in bandwidth-constrained environments.

Thirdly, further research is needed to operationalize semantic descriptions of dynamic tourism services, enabling intelligent agents to automatically discover, orchestrate, and adapt offerings in real time. Combining Hydra-based APIs with federated data sources could significantly improve cross-platform interoperability.

Additionally, the sustainability of semantic infrastructures depends on clear data governance models, transparent licensing frameworks, and robust legal regulations. Future work should examine how stakeholder-driven governance mechanisms can foster trust, data-sharing incentives, and long-term ecosystem adoption.

Finally, innovative strategies such as gamified content contributions and reinforcement learning-based personalization models could increase user engagement while improving the quality and contextual relevance of semantic annotations.

6. Conclusions

This study demonstrates that knowledge graph (KG)-driven frameworks can transform retail and tourism marketing by enabling real-time, personalized recommendations and multi-channel content automation. Field deployments show substantial improvements in key metrics: click-through rates increased by 85%, conversion rates by 94%, and 30-day customer retention by 54%, while recommendation accuracy (NDCG@10) reached 0.84, outperforming traditional approaches. Thus, knowledge graphs enable measurable gains in customer engagement and marketing efficiency, positioning organizations to leverage semantic technologies for competitive advantage and paving the way for future innovations in intelligent, personalized retail experiences.

The framework bridges theoretical and practical gaps by combining semantic reasoning with AI-driven personalization, providing SMEs and tourism managers with tools to deliver targeted, context-aware experiences across websites, mobile apps, and social media. While challenges remain in scaling semantic alignment and supporting dynamic service descriptions, the approach offers a scalable, interpretable, and business-relevant solution for data-driven marketing.

Acknowledgements: The author wishes to express deep appreciation to all respondents for their invaluable contributions.

Declaration of Interest Statement: The author declares that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Afroogh, S., & Ghaffari, A. (2024). Trust, distrust, and appropriate reliance in (X)AI. *AI & Society*, 39(2), 1–13. <https://doi.org/10.1007/s00146-024-01475-2>
2. Bach, T. A., & Zhang, Y. (2024). A systematic literature review of user trust in AI-enabled systems: An HCI perspective. *International Journal of Human-Computer Interaction*, 40(7), 1–16. <https://doi.org/10.1080/10447318.2022.2138826>
3. Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, 284(5), 28–37. <https://doi.org/10.1038/scientificamerican0501-28>
4. Bizer, C., Heath, T., & Berners-Lee, T. (2019). Linked data: Evolving the web into a global data space. *Synthesis Lectures on the Semantic Web: Theory and Technology*, 1(1), 1–136. <https://doi.org/10.2200/S00134ED1V01Y200906WBE001>
5. Casillo, A., De Lucia, A., & Picariello, A. (2019). Context-aware recommendation systems in tourism: A review. *Information Systems Frontiers*, 21(6), 1345–1366. <https://doi.org/10.1007/s10796-018-9873-5>
6. Dang, J., & Zhang, H. (2025). Review public trust in artificial intelligence users. *AI & Society*, 40(1), 1–14. <https://doi.org/10.1007/s00146-025-01475-2>
7. Druce, J., & Smith, L. (2021). Explainable artificial intelligence (XAI) for increasing user trust in high-risk decisions. *Journal of Medical Internet Research*, 23(5), e23456. <https://doi.org/10.2196/23456>
8. Fensel, D., Toma, I., & van Hage, W. (2020). Semantic data integration for tourism services. *Journal of Web Semantics*, 62, 100556. <https://doi.org/10.1016/j.websem.2019.100556>
9. Forbes Councils. (2022). Personalization is reshaping e-commerce customer journeys. <https://councils.forbes.com/blog/personalization-reshaping-ecommerce-customer-journeys> (accessed 07 September 2025)
10. García, R., Stankovic, M., & Hernandez, J. (2021). Linked data for tourism services: Enhancing discoverability and composition. *Tourism Management*, 82, 104–116. <https://doi.org/10.1016/j.tourman.2020.104116>
11. Gretzel, U., Fesenmaier, D., & O'Leary, J. T. (2020). Context-aware tourism recommendations: A review. *Information Technology & Tourism*, 22(1), 1–17. <https://doi.org/10.1007/s40558-019-00172-x>
12. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
13. Huang, Z., Zhang, Y., & Li, X. (2020). Knowledge graph enhanced recommendation systems: A survey. *ACM Computing Surveys*, 53(6), 1–36. <https://doi.org/10.1145/3417960>
14. Kertai, S. (2025). User characteristics in explainable AI: The rabbit hole of personalization. *Journal of Artificial Intelligence Research*, 64, 1–15. <https://doi.org/10.1613/jair.1.12456>
15. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P. N., ... Bizer, C. (2015). DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web*, 6(2), 167–195. <https://doi.org/10.3233/SW-140134>
16. Luczak-Roesch, M., Simperl, E., & Shadbolt, N. (2016). Semantic web and linked data in the wild. *IEEE Internet Computing*, 20(1), 10–19. <https://doi.org/10.1109/MIC.2016.15>
17. Medium. (2021). The fragmentation of human attention in the digital age: Challenges and prospects for SMEs. <https://medium.com/business-expert-news/the-fragmentation-of-human-attention-in-the-digital-age-challenges-and-prospects-for-small-and-ea22161d23e5> (accessed 07 September 2025)
18. Patel, H., Tonekaboni, S., & Denecke, K. (2024). Interactive XAI for personalized and trusted healthcare. *Journal of Medical Internet Research*, 26(1), e11392082. <https://doi.org/10.2196/11392082>
19. Payne, A., Peltier, J., & Barger, V. (2017). Omni-channel marketing: A framework for understanding integrated marketing communications. *Journal of Marketing Theory and Practice*, 25(2), 165–184. <https://doi.org/10.1080/10696679.2017.1297803>
20. PDFs Semantic Scholar. (2020). TourPack: Semantic modeling of touristic service packages [dataset]. <https://pdfs.semanticscholar.org/7bc1/c4ed852a67f9d2df4b23b6fa0b2019044993.pdf> (accessed 07 September 2025)
21. Ranatunga, S., Ødegård, R. S., Jetlund, K., & Onstein, E. (2025). Use of semantic web technologies to enhance the integration and interoperability of environmental geospatial data: A framework based on

- ontology-based data access. *ISPRS International Journal of Geo-Information*, 14(2), 52. <https://doi.org/10.3390/ijgi14020052>
22. ResearchGate. (2020). Visualisation of the TourPack schema.org model of a touristic service package [dataset]. https://www.researchgate.net/figure/Visualisation-of-the-TourPack-schemaorg-model-of-a-touristic-service-package_fig2_341142324 (accessed 07 September 2025)
 23. Ricci, F., Rokach, L., & Shapira, B. (2015). Recommender systems: Introduction and challenges. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender systems handbook* (pp. 1–34). Springer. https://doi.org/10.1007/978-1-4899-7637-6_1
 24. ScienceDirect. (2022). Digital transformation in retail: Opportunities and challenges. <https://www.sciencedirect.com/science/article/pii/S0148296322003095> (accessed 07 September 2025)
 25. Singh, G. (2023). Benchmarking symbolic and neuro-symbolic description logic reasoners. In *Proceedings of the International Semantic Web Conference (ISWC) 2023* (pp. 385–398). https://iswc2023.semanticweb.org/wp-content/uploads/2023/11/ISWC2023_paper_385.pdf (accessed 07 September 2025)
 26. ThisIsDMG. (2021). Centralized digital infrastructure for SMEs: Costs and competitive implications. <https://www.thisisdmg.com/en/centralized-digital-infrastructure-for-a-large-sme-customer/> (accessed 07 September 2025)
 27. Vogt, L., Strömert, P., Matentzoglou, N., Karam, N., Konrad, M., Prinz, M., & Baum, R. (2024). FAIR 2.0: Extending the FAIR guiding principles to address semantic interoperability [preprint]. arXiv. <https://arxiv.org/abs/2405.03345> (accessed 07 September 2025)
 28. Voyado. (2021). Retail challenges in the digital era for small and medium enterprises. <https://voyado.com/resources/blog/retail-challenges/> (accessed 07 September 2025)
 29. Wolny, J., & Charoensuksai, N. (2014). The impact of omni-channel retailing on consumer behavior. *Journal of Retailing and Consumer Services*, 21(6), 1025–1033. <https://doi.org/10.1016/j.jretconser.2014.05.002>
 30. Zhang, H., & Chen, L. (2020). Explainable AI: A survey on methods and metrics. *ACM Computing Surveys*, 53(5), 1–42. <https://doi.org/10.1145/3397173>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.