MYCAMPANIA.TRAVEL: Leveraging Generative AI to Enhance Digital Travel Experiences

Aniello Somma[1,2,*], Sergio Di Meglio[1,2], Fabio Scippacercola[1], Ermanno Battista[1], Sergio Di Martino[2] and Luigi Libero Lucio Starace[2]

[1]Fervento srl, Naples, Italy
[2]Dept. of Electrical Engineering and Information Technology, Università degli Studi di Napoli Federico II, Naples, Italy

Abstract

The digital transformation of the tourism sector is reshaping traditional business models by integrating advanced technologies such as artificial intelligence (AI), the Internet of Things (IoT), and big data. Within this landscape, Tourism 4.0 emphasizes interoperability, personalization, and seamless service integration. However, legacy travel platforms often fall short of delivering the cohesive and tailored experiences that modern travelers demand. This paper introduces MyCampania.Travel, an ongoing project developed by the Fervento srl innovative startup and aimed at developing an AI-driven digital ecosystem to enhance tourism in Italy's Campania region through improved data harmonization, automated content generation, and personalized user experiences. The system employs a cloud-native microservice architecture integrated with large language models (LLMs) to address fragmentation in travel data and enrich content. Results from a preliminary evaluation of two representative use cases (room classification, and description refactoring) suggest that LLMs can enhance the clarity, consistency, and completeness of tourism content, enabling more intuitive and trustworthy booking experiences.

Keywords

Tourism, Generative AI, LLMs, AI in Industry, Campania

1. Introduction and Industrial Context

The travel and tourism sector is undergoing rapid digital transformation, driven by technologies such as Artificial Intelligence (AI), the Internet of Things (IoT), and Big Data [1]. This shift extends beyond workflow digitization, aiming to reshape business models and deliver smarter, personalized experiences. The Tourism 4.0 paradigm exemplifies this evolution, emphasizing interoperability, data integration, and tailored value creation across the tourism value chain [2]. However, the industry remains highly fragmented and dynamic [3], with persistent challenges in data consistency, particularly for room types, amenities, and policies across multiple providers. While third-party catalogs often offer outdated or generic content, building proprietary solutions demands significant ongoing effort [4].

Fervento, an innovative startup focused on building software products for various domains, with a strong specialization in travel and tourism software, is addressing these challenges through MyCampania.Travel, a cloud-native [5], AI-powered platform currently in development. This initiative, carried out also within the framework of the Italian National Recovery and Resilience Plan and the FAIR (Future Artificial Intelligence Research) project, aims to enhance tourism in Italy's Campania region. By leveraging Generative AI, the platform tackles data fragmentation and inconsistency, enabling seamless integration, enriched content, and personalized discovery at scale.

The platform's design centers on three strategic objectives: (1) **Data Harmonization**, standardizing heterogeneous data into coherent formats; (2) **Traveler-Centric Content Generation**, aggregating fragmented inputs into transparent, trust-building documents; and (3) **Content Completion and**

Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy
*Corresponding author.
aniello@fervento.com (A. Somma); sergio.dimeglio@unina.it (S. Di Meglio); fabio@fervento.com (F. Scippacercola); ermanno@fervento.com (E. Battista); sergio.dimartino@unina.it (S. Di Martino); luigiliberolucio.starace@unina.it (L. L. L. Starace)
https://www.linkedin.com/in/sergio-di-meglio (S. Di Meglio); https://luistar.github.io/ (L. L. L. Starace)
0000-0002-1019-9004 (S. Di Martino); 0000-0001-7945-9014 (L. L. L. Starace)

Enrichment, using AI-driven insights to fill gaps (e.g., contextualizing proximity to landmarks). To achieve these, FERVENTO explores two complementary LLM strategies: prompt engineering [6] with high-capacity models for few-shot generation, and parameter-efficient fine-tuning for domain-specific adaptation [7].

Recent research underscores the growing interest in LLMs for tourism. Wei et al. [8] propose domain-specific adaptations like *TourLLM*, while Hsu et al. [9] outline ethical and technical frameworks for sector-specific LLMs. In [10], Starace and Di Martino showed that LLMs, even with zero-shot learning approaches, can effectively reason about spatial data, generating coherent descriptions of spatial phenomena of interest. Others, such as Dogan Gursoy and Song [11], examine generative AI's disruptive potential in hospitality, and Rather [12] analyze trust and customer experience in AI-driven services.

However, most studies remain conceptual, with limited real-world deployments. The MYCAMPANIA.TRAVEL project aims to fill this gap, by defining a reference architecture, AI workflows, and comparatively analyzing the performance of different LLM-based strategies such as fine-tuning versus prompt engineering, highlighting trade-offs in scalability, adaptability, and performance in a real-world setting.

Preliminary evaluations on room classification and description refactoring demonstrate promising improvements in content consistency, completeness, and quality. The following sections elaborate on the platform's development, technical foundations, and roadmap, offering a practical case study for AI integration in tourism.

2. MyCampania.Travel

MYCAMPANIA.TRAVEL is a B2B (Business-to-Business) digital platform designed to support incoming tourism in Italy's Campania region. It streamlines the discovery and booking of accommodations, local services, and culturally rich experiences, from archaeological sites and museums to boat tours and guided excursions. By offering personalized, context-aware recommendations, the platform enhances the visibility and appeal of regional offerings. More than a booking tool, the platform aims to strengthen the local tourism ecosystem by improving the digital presence and competitiveness of small operators, ultimately driving sustainable growth and broader access to regional resources.

2.1. Platform architecture

The platform is built on a cloud-native microservice architecture, emphasizing performance, scalability, and reliability. Powered byFERVENTO's proprietary framework, *Medesimo*, the architecture supports modular expansion, robust orchestration, and fault-tolerant deployments. Key architectural features include:

1. **Extensible Data Exchange APIs**: A flexible interoperability layer supports real-time integration with third-party systems, enabling bookings and data retrieval for hotels, museums, tours, transfers, and more.
2. **Territory-Specific Experiences**: In addition to catalogued services, the platform integrates unique, locally-sourced experiences that are often overlooked by standard distribution channels, discovered by analyzing open spatial data [13]. This boosts market differentiation and promotes underrepresented operators.
3. **AI-Driven Content Intelligence**: Advanced natural language processing and classification techniques are used to harmonize and enrich provider data (e.g., room types, amenities, and travel documents), resulting in consistent and high-quality content.

These features not only enhance the B2B offering but also establish the foundation for future B2C (Business-to-Consumer) expansions and API-based resale models, opening up new regional distribution opportunities.

2.1.1. Core modules

MyCampania.Travel is built around four key modules designed to enhance the visibility of Campania's cultural heritage while ensuring adequate performance [14, 15], scalability, interoperability, and intelligent content processing.

The **presentation module** provides a responsive user interface optimized for cross-device compatibility, enabling unified discovery and booking of regional tourism offerings through intuitive navigation and filtering capabilities. Underlying this interface, the extensible **service integration module** establishes standardized connectivity with external providers via RESTful APIs, incorporating validation protocols to maintain data consistency across hotel, museum, and tour booking systems. At the core of the platform, the **data processing module** employs LLMs for automated content harmonization, including classification of accommodation attributes and generation of standardized descriptions from heterogeneous inputs, complemented by quality assessment modules. Supporting these functions, the **resource management module** implements optimization strategies through multi-level caching and computational reuse, monitored through dedicated infrastructure analytics to balance performance and operational costs. This architecture's modular design facilitates both current operational requirements and future extensibility while addressing sector-specific challenges of data heterogeneity and service integration.

2.2. Introduction of Generative AI

To address the challenges of fragmented and inconsistent data, the platform integrates Generative AI to automate content harmonization, enhance quality, and deliver personalized user experiences. LLMs are applied to the following core tasks:

- **Room Classification**: Standardizes room names and types across multiple providers, resolving inconsistencies and ambiguous labels. This enhances platform usability by presenting users with a clear, harmonized view of available accommodations, improving decision-making and trust.
- **Description Refactoring**: Consolidates fragmented facility data from disparate sources into coherent, well-structured descriptions. By unifying the presentation and filling in missing details, this process supports content completeness, clarity, and platform reliability.
- **Local Heritage Enrichment**: Enhances facility listings with contextual references to nearby historical, cultural, and natural landmarks. This promotes regional identity and encourages travelers to explore lesser-known but valuable attractions in Campania.
- **Onsite Fee Detection**: Identifies and highlights hidden or under-reported fees, such as local taxes or surcharges, based on implicit signals in the input data. This improves transparency and reduces the likelihood of negative surprises for travelers.

2.3. Road Map: Development phases

The project roadmap is structured into three sequential phases:

- **Phase I – Design & AI Experimentation**: Establish the foundational architecture and validate the feasibility of AI components. This phase includes prompt design, fine-tuning experiments, and evaluation of model performance using tourism-specific data to determine their readiness for integration.
- **Phase II – Implementation & Integration**: Build the production infrastructure and incorporate AI pipelines into real-world workflows. This includes deploying microservices, integrating third-party APIs, and embedding LLMs into tasks such as classification, content enrichment, and itinerary curation. Moreover, this phase will also include a thorough testing process, including three tiers of validation: unit tests [16, 17] for individual components, integration tests for subsystem interactions, and end-to-end tests to validate the system from the perspective of end users [18, 19, 20, 21].

- **Phase III – Launch, Promotion & Sustainability Strategy**: Launch the platform publicly, execute targeted outreach and onboarding efforts, and implement a sustainability strategy. This includes adopting green computing practices [15], autoscaling infrastructure, and planning for long-term feature evolution such as B2C extensions and open APIs.

3. Preliminary Evaluation

As part of Phase 1, we conducted an initial evaluation focusing on two key tasks: *room classification* and *description refactoring*.

For room classification, we evaluated three LLMs, namely LLaMA 3.1 8B [22], Mistral 7B [23], and Mixtral 8x7B[24] on a dataset of 30 hotel facilities, each containing inconsistent room naming conventions. The models generated standardized classifications, which were reviewed by human annotators to assess quality, accuracy, and reliability. Table 1 shows the results expressed as average accuracy and F1 scores with standard deviations, showing that Mixtral 8x7B achieved the highest performance and exhibited strong consistency.

Table 1
Room Classification Performance (Prompt Engineering)

Model	Accuracy	F1 Score
LLaMA 3.1 8B	0.8918% ± 0.1095%	0.9390% ± 0.0664%
Mistral 7B	0.8666% ± 0.1069%	0.9249% ± 0.0635%
Mixtral 8x7B	**0.9211**% ± 0.1085%	**0.9553**% ± 0.0641%

Regarding description refactoring, we compared a QLoRA fine-tuned [25] Mistral 7B model with a prompt-engineered Mixtral 8x7B. The evaluation was based on a dataset of 20 hotel descriptions from various providers, often characterized by incomplete or inconsistent content. The models were tasked with generating coherent and structured descriptions to enhance the clarity and appeal of listings. Human reviewers assessed the outputs in terms of completeness, precision, length, and hallucinations [26]. Table 2 shows that Mixtral 8x7B outperforms the fine-tuned Mistral 7B across all evaluation criteria, particularly in minimizing hallucinations and maximizing content completeness.

Table 2
Evaluation of Description Refactoring Models

Model	Completeness	Precision	Length	Hallucinations
Mistral 7B-FT	93% ± 8.8%	96% ± 3.2%	277 ± 70	4% ± 3.8%
Mixtral 8x7B	**99.6%** ± 1.4%	**98.8%** ± 3.2%	249.2 ± 28	**1.2%** ± 3.2%

4. Conclusion and Future Work

In this work, we introduced MyCampania.Travel, a digital platform designed to enhance regional tourism in Campania by combining a cloud-native infrastructure with Generative AI technologies. The platform addresses critical challenges such as fragmented data, inconsistent content, and the limited visibility of local experiences, aiming to offer travelers a coherent, engaging, and culturally rich planning experience.

We presented an overview of this ongoing project, including a preliminary evaluation of different LLM-based strategies for content enrichment, and showed how these components integrate into a scalable system tailored to the tourism domain. Preliminary results suggest significant benefits, including improved content quality, greater traveler engagement, and enhanced visibility for regional operators.

Future developments will focus on expanding the platform's reach and capabilities. These include exposing public APIs to support third-party integrations, transitioning toward a B2C model for direct

traveler access, and further optimizing fine-tuned LLMs for key subtasks such as classification and summarization. Additionally, we plan to strengthen multilingual support, enabling more inclusive and accessible content, and exploring agent-based orchestration for modular and context-aware AI workflows. To ensure reliability and relevance, we will also develop rigorous evaluation frameworks to assess the accuracy, fluency, and consistency of AI-generated content, based on benchmarks and evaluation protocols introduced in prior work [27].

Through these advancements, MyCampania.Travel aims not only to support regional tourism growth but also to offer a replicable blueprint for the integration of AI in domain-specific platforms.

Acknowledgements

This work has been partially supported by the Italian PNRR MUR project PE0000013-FAIR.
Authors from Fervento inform that MyCampania.travel is a project co-financed by Unione Europea, Stato Italiano, Regione Campania, with the PR CAMPANIA FESR 2021-2027 - Asse I - Obiettivo Specifico 1.1 – Azione 1.1.3 – Avviso pubblico CAMPANIA STARTUP 2023 (CUP B68I23005640007).

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT (GPT-4o) and Grammarly in order to: perform automated grammar and spelling checks; paraphrase and reword. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Imtiaz, D. J. Kim, Digital transformation: Development of new business models in the tourism industry, Culinary Science & Hospitality Research 25 (2019) 91–101. doi:10.20878/cshr.2019.25.4.010.

[2] T. Pencarelli, The digital revolution in the travel and tourism industry, Information Technology & Tourism 22 (2020) 455–476.

[3] S. Angeloni, C. Rossi, Online search engines and online travel agencies: A comparative approach, Journal of Hospitality & Tourism Research 45 (2021) 720–749. doi:10.1177/1096348020980101.

[4] Z. Xiang, K. Wöber, D. R. Fesenmaier, Representation of the online tourism domain in search engines, Journal of Travel Research 47 (2008) 137–150.

[5] V. Casola, A. De Benedictis, S. Di Martino, N. Mazzocca, L. L. L. Starace, Security-aware deployment optimization of cloud–edge systems in industrial iot, IEEE Internet of Things Journal 8 (2020) 12724–12733.

[6] G. Marvin, N. Hellen, D. Jjingo, J. Nakatumba-Nabende, Prompt engineering in large language models, in: International conference on data intelligence and cognitive informatics, Springer, 2023, pp. 387–402.

[7] X. Lin, W. Wang, Y. Li, S. Yang, F. Feng, Y. Wei, T.-S. Chua, Data-efficient fine-tuning for llm-based recommendation, in: Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval, 2024, pp. 365–374.

[8] Q. Wei, M. Yang, J. Wang, W. Mao, J. Xu, H. Ning, Tourllm: Enhancing llms with tourism knowledge, 2024. URL: https://arxiv.org/abs/2407.12791. arXiv:2407.12791.

[9] C. H. Hsu, G. Tan, B. Stantic, A fine-tuned tourism-specific generative ai concept, Annals of Tourism Research 104 (2024) 103723. Annals of Tourism Research: 50th Anniversary Issue.

[10] L. L. L. Starace, S. Di Martino, Can large language models automatically generate gis reports?, in: International Symposium on Web and Wireless Geographical Information Systems, Springer, 2024, pp. 147–161.

[11] Y. L. Dogan Gursoy, H. Song, Chatgpt and the hospitality and tourism industry: an overview of current trends and future research directions, Journal of Hospitality Marketing & Management 32 (2023) 579–592.
[12] R. A. Rather, Ai-powered chatgpt in the hospitality and tourism industry: benefits, challenges, theoretical framework, propositions and future research directions, Tourism Recreation Research 0 (2024) 1–11. doi:10.1080/02508281.2023.2287799.
[13] S. Di Martino, E. Landolfi, N. Mazzocca, F. R. di Torrepadula, L. L. L. Starace, A visual-based toolkit to support mobility data analytics, Expert Systems with Applications 238 (2024) 121949.
[14] S. Di Meglio, L. L. L. Starace, S. Di Martino, Starting a new rest api project? a performance benchmark of frameworks and execution environments., in: IWSM-Mensura, 2023.
[15] S. Di Meglio, L. L. L. Starace, Evaluating performance and resource consumption of rest frameworks and execution environments: Insights and guidelines for developers and companies, IEEE Access (2024).
[16] F. Altiero, A. Corazza, S. Di Martino, A. Peron, L. L. L. Starace, Inspecting code churns to prioritize test cases, in: IFIP International Conference on Testing Software and Systems, Springer, 2020, pp. 272–285.
[17] F. Altiero, A. Corazza, S. Di Martino, A. Peron, L. L. L. Starace, Regression test prioritization leveraging source code similarity with tree kernels, Journal of Software: Evolution and Process 36 (2024) e2653.
[18] A. Corazza, S. Di Martino, A. Peron, L. L. L. Starace, Web application testing: Using tree kernels to detect near-duplicate states in automated model inference, in: Proceedings of the 15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM), 2021.
[19] E. Battista, S. Di Martino, S. Di Meglio, F. Scippacercola, L. L. L. Starace, E2E-Loader: A framework to support performance testing of web applications, in: 2023 IEEE Conference on Software Testing, Verification and Validation (ICST), IEEE, 2023, pp. 351–361.
[20] S. Di Meglio, L. L. L. Starace, Towards predicting fragility in end-to-end web tests, in: Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering, 2024, pp. 387–392.
[21] S. Di Meglio, L. Libero Lucio Starace, S. Di Martino, E2e-loader: A tool to generate performance tests from end-to-end gui-level tests, in: 2025 IEEE Conference on Software Testing, Verification and Validation (ICST), 2025, pp. 747–751. doi:10.1109/ICST62969.2025.10989035.
[22] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, et al., The llama 3 herd of models, arXiv preprint arXiv:2407.21783 (2024).
[23] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de Las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mistral 7b, CoRR abs/2310.06825 (2023). URL: https://doi.org/10.48550/arXiv.2310.06825. doi:10.48550/ARXIV.2310.06825. arXiv:2310.06825.
[24] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. d. l. Casas, E. B. Hanna, F. Bressand, et al., Mixtral of experts, arXiv preprint arXiv:2401.04088 (2024).
[25] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, Qlora: Efficient finetuning of quantized llms, 2023. URL: https://arxiv.org/abs/2305.14314.
[26] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang, Q. Chen, W. Peng, X. Feng, B. Qin, et al., A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions, ACM Transactions on Information Systems 43 (2025) 1–55.
[27] S. Di Meglio, L. L. L. Starace, V. Pontillo, R. Opdebeeck, C. De Roover, S. Di Martino, E2egit: A dataset of end-to-end web tests in open source projects, in: 2025 IEEE/ACM 22nd International Conference on Mining Software Repositories (MSR), IEEE/ACM, 2025, pp. 10–15.